

Symantec Backup Exec 2014

Administrator's Guide



Documentation version: 2014

Legal Notice

Copyright © 2014 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this Symantec product for more information on the Third Party Programs.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our website at the following URL:

www.symantec.com/business/support/

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information

- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs, DVDs, or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Contents

Technical Support	4
Chapter 1 Introducing Backup Exec 2014	37
About Backup Exec	37
How Backup Exec works	38
What's new in Backup Exec 2014	39
What's new in Backup Exec 2014 agents and options	48
Chapter 2 Installation	53
About the Backup Exec installation process and licensing options	54
Backup Exec preinstallation checklist	57
Running the Environment Check before installing or upgrading Backup Exec	58
Microsoft SQL Server 2008 Express Edition components installed with Backup Exec	60
Backup Exec's standard features	61
System requirements for Backup Exec	63
Installing Backup Exec by using the Installation Wizard	66
Installing additional Backup Exec options to the local Backup Exec server	76
Push-installing Backup Exec to remote computers	78
Methods for installing the Agent for Windows	86
Push-installing the Agent for Windows to remote computers	86
Installing updates to the Agent for Windows on remote computers	93
Installing the Agent for Windows in an Active Directory network	94
Using a command prompt to install the Agent for Windows on a remote computer	98
Using a command script to install the Agent for Windows	101
Installing the Remote Administrator	102
Installing the Remote Administrator using the command line	103
Running the Remote Administrator	105
Installing Backup Exec using the command line (silent mode)	108

	Command line switches for silent mode installation of Backup Exec	
	Exec	109
	Creating and using installation parameter files	117
	Installing a trial version of Backup Exec agents and options	118
	About the installation log	119
	Viewing the Installation Summary Report	119
	Repairing Backup Exec	120
	Updating Backup Exec with LiveUpdate	121
	Viewing installed updates	123
	Viewing license information	123
	Backup Exec maintenance contract information	124
	Updating expired maintenance contracts	124
	Managing maintenance contract customer numbers	125
	About upgrading from previous versions of Backup Exec to Backup Exec 2014	125
	Upgrade checklist for Backup Exec 2010 and earlier	127
	Using the Migration Report to determine how existing jobs changed after an upgrade to Backup Exec 2014	129
	Post-installation tasks	129
	Uninstalling Backup Exec	130
	Uninstalling Backup Exec using the command line	131
	Uninstalling Backup Exec options from the local Backup Exec server	131
Chapter 3	Getting Started	133
	About the Backup Exec Administration Console	133
	Displaying the version information for Backup Exec	136
	How to sort, filter, and copy information on the Backup Exec Administration Console	136
	Customizing views on the Backup Exec Administration Console	138
	Configuring the Home tab	139
	Configuring the Symantec RSS Reader	143
	Creating a disaster preparation plan (DPP)	144
	Getting started with backups	146
Chapter 4	Backups	149
	How to prepare for your first backup with Backup Exec	150
	Required user rights for backup jobs	151
	About the list of servers on the Backup and Restore tab	152
	Adding servers that you want to back up to the list of servers on the Backup and Restore tab	153

Removing servers from the list of servers on the Backup and Restore tab	154
Creating a server group	154
Hiding or viewing server groups on the Backup and Restore tab	155
Editing a server group	156
Retiring servers from Backup Exec	157
Removing server groups from Backup Exec	158
Backing up data	158
How Backup Exec backs up and restores volumes that are enabled for deduplication in Windows	168
How Backup Exec backs up and restores Microsoft Virtual Hard Disk (vhd) files	169
About selecting data to back up	169
Changing the order in which backup sources are backed up	175
Excluding files from backups	178
Including specific files with a backup job's backup selections	181
About backing up critical system components	183
Backup Exec Shadow Copy Components file system	187
Backup methods in Backup Exec	188
Configuring backup methods for backup jobs	196
How Backup Exec determines if a file has been backed up	199
Configuring Backup Exec to automatically delete files after a backup	202
Configuring network options for backup jobs	204
Running the next scheduled backup job before its scheduled time	205
Editing backup definitions	206
Testing credentials for backup sources	211
Replacing the credentials for a backup source	212
Creating new credentials for a backup source	213
Deleting retired or unused backup sources from the Credentials pane	214
How job scheduling works in Backup Exec	215
Including a specific date in the schedule for a backup job	216
Preventing backup jobs from running on a specific date	218
Viewing all scheduled backup jobs on a calendar	219
Adding a stage to a backup definition	219
Editing a stage	221
Duplicating backup sets or a job history manually	221
Running a test run job manually	226
Verifying backed up data manually	227
Copying data from a virtual tape library to a physical tape device using DirectCopy to tape	229

Chapter 5	Restores	231
	Methods for restoring data in Backup Exec	232
	Searching for data to restore	234
	Restoring data from a server, a backup set, a backup job, or a storage device	234
	Restoring file system data	238
	Performing a complete online restore of a Microsoft Windows computer	239
	Restoring System State	240
	Installing a new Windows Server domain controller into an existing domain by using a redirected restore	242
	Restoring Backup Exec Shadow Copy Components	244
	Restoring utility partitions or Unified Extensible Firmware Interface system partitions	245
	About restoring encrypted data	245
	About restoring NetWare SMS volume backups to non-SMS volumes with Backup Exec	246
	Restoring data from ARCserve tapes	246
	Canceling a restore job	247
	How Backup Exec catalogs work	248
	Configuring default options for catalogs	249
	Cataloging backup sets	251
Chapter 6	Job management and monitoring	253
	How to monitor and manage jobs in Backup Exec	253
	About the Job Monitor	255
	Viewing job activity details for active jobs	257
	Canceling an active job	257
	Holding jobs	258
	Removing the hold on jobs	259
	Holding the job queue	260
	Removing the hold on the job queue	260
	Running a scheduled job immediately	261
	Changing the priority for a scheduled job	262
	Deleting scheduled jobs	263
	About the Job History	264
	Viewing the history of a job	265
	Deleting a job from the Job History	265
	Running a job from the Job History	266
	Viewing the job log	267
	Finding text in the job log	267
	Printing the job log	269

Saving the job log	269
Linking from the job log to the Symantec Technical Support website	270
How to use job logs with vertical applications	270
Configuring default job log options	271
Error-handling rules for failed or canceled jobs	272
Creating a custom error-handling rule	273
Enabling or disabling error-handling rules	274
Deleting a custom error-handling rule	274
Enabling an error-handling rule for a failed job	275
Custom error-handling rule for recovered jobs	275
About the cluster failover error-handling rule	276
List of job statuses in Backup Exec	277
Setting job status and recovery options	282

Chapter 7	Alerts and notifications	285
	Alerts and notifications in Backup Exec	286
	Where to find alerts in Backup Exec	287
	Enabling active alerts and alert history to display on the Home tab	289
	Active alerts properties	290
	Viewing the alert history for a server or a storage device	291
	Deleting an alert from alert history	291
	Copying alert text to a document or email	292
	Filtering alerts	292
	Viewing the job log from an alert	293
	Responding to active alerts	293
	Clearing all informational alerts manually	294
	Setting up notification for alerts	294
	Configuring email or text message notification for alerts	295
	Adding recipients for notification	297
	Adding a recipient group for alert notifications	299
	Removing a recipient from a group	300
	Disabling email or text message alert notification for a recipient	300
	Editing recipient notification properties	301
	Deleting recipients	301
	Configuring alert categories	302
	Assigning recipients to receive notifications for specific alert categories	304
	Sending a notification when a job completes	304
	Notification options for jobs	305
	Disabling notifications for a specific alert category	305

Configuring default alert settings	306
Enabling or disabling pop-up alerts	307
SNMP traps for Backup Exec alerts	308
Installing and configuring the SNMP system service	311
Installing the Windows Management Instrumentation performance counter provider	311
Installing the Windows Management Instrumentation provider for SNMP	312
Uninstalling the Windows Management Instrumentation performance counter provider	312
Uninstalling the Windows Management Instrumentation provider for SNMP	312

Chapter 8 Disk-based and network-based storage 313

Features and types of disk-based storage and network-based storage	313
Storage trending statuses for disk storage and virtual disks	316
Setting low disk space thresholds on disk-based storage	317
Configuring disk storage	318
Changing the location of a disk storage device	320
Editing disk storage properties	321
How to restore data from a reattached or reinserted disk-based storage device	327
Configuring disk cartridge storage	328
Editing disk cartridge properties	329
Editing disk cartridge media properties	333
How data lifecycle management (DLM) deletes expired backup sets on disk-based storage	335
Backup sets	342
Changing the expiration date of backup sets on disk-based storage	345
Retaining backup sets on disk-based storage to prevent them from expiring	346
Releasing retained backup sets on disk-based storage	347
Viewing the contents or properties of backup sets	348

Chapter 9 Legacy backup-to-disk folders 350

Legacy backup-to-disk folders	350
Editing legacy backup-to-disk folder properties	351
Changing the name or description of a legacy backup-to-disk folder	351

Changing the buffered read setting for a legacy backup-to-disk folder	352
Changing the location of a legacy backup-to-disk folder	352
Recreating a legacy backup-to-disk folder and its contents	353
Restoring data from a legacy backup-to-disk folder	354
Chapter 10 Tape storage	356
Support for tape drives and robotic libraries	357
About the Virtual Tape Library Unlimited Drive Option	358
About the Library Expansion Option	358
Adding or replacing devices with the Hot-swappable Device Wizard	358
Installing Symantec tape device drivers	360
Editing tape drive properties	361
Viewing tape drive statistics	366
Robotic libraries in Backup Exec	366
Requirements for setting up robotic library hardware	367
Inventorying robotic libraries when Backup Exec services start	368
Configuring barcode rules for a robotic library	368
Initializing a robotic library when the Backup Exec service starts	372
Defining a cleaning slot	373
Editing robotic library properties	373
Creating robotic library partitions	373
Changing the name of a robotic library partition	375
Adding or removing tape drives in a robotic library partition	375
Reassigning a slot base number for robotic libraries	376
Editing robotic library slot properties	376
Removing or reconfiguring robotic library partitions	377
Managing tapes	378
Default media sets	379
Overwrite protection periods and append periods in media sets	382
Creating media sets for tapes	386
Changing the overwrite protection period or the append period for a media set	391
Changing the name and description of a media set	393
Changing the media vault or vaulting periods for a media set	393
Deleting a media set	394
Renaming a media set	394

	Media overwrite protection levels for tape media	394
	Overwriting allocated or imported tape media	395
	How Backup Exec searches for overwritable media in tape drives	395
	Viewing tapes that are used by a job	398
	Labeling tape media	398
	About labeling imported tape media	400
	How barcode labels become media labels for tapes	400
	Renaming a tape media label	401
	How WORM media is used in Backup Exec	401
	Default media vaults	402
	Changing the name or description of a media vault	403
	Creating media vault rules to move tape media to and from media vaults	404
	Updating the tape media location in media vaults	405
	Deleting a media vault	405
	Moving tape media to a vault	406
	Retiring damaged tape media	406
	Deleting tape media	407
	Erasing tape or disk cartridge media	407
	About cataloging tape or disk cartridge media that contains encrypted backup sets	409
	Associating tape media with a media set	409
	Editing tape media properties	410
	Tape media rotation strategies	412
Chapter 11	Storage device pools	416
	Creating storage device pools	416
	Specifying a default selection method for a device in a storage device pool	419
	Viewing jobs, job histories, and active alerts for a storage device pool	420
	Adding or removing devices in a storage device pool	420
	Storage device pool properties	421
Chapter 12	Storage operations	423
	About storage operation jobs	424
	Storage tab overview in Backup Exec	425
	Using the Configure Storage wizard	428
	Viewing details for multiple storage devices	431
	Sending a notification when a scheduled storage operation job completes	432

Scheduling a storage operation job	432
Editing global settings for storage	439
Sharing storage devices	448
Shared Storage options	449
Deleting a storage device	450
Changing the state of a storage device to online	450
Renaming a storage device	451
Viewing jobs, job histories, backup sets, and active alerts for storage devices	451
Cataloging a storage device	452
Scanning a storage device	452
Inventorying a storage device	453
Inventorying and cataloging a storage device	454
Pausing and unpausing a storage device	455
Disabling and enabling a storage device	455
Initializing a robotic library	456
Formatting a tape as a WORM tape	456
Retensioning a tape	456
Formatting a tape in a tape drive	457
Ejecting media from a disk cartridge or tape drive	458
Cleaning a robotic library drive	458
Importing media to Backup Exec	459
Exporting media and expired media	466
Locking and unlocking the robotic library's front portal	469
Backup Exec server and storage device states	475

Chapter 13 Conversion to virtual machines 479

How conversion of physical computers to virtual machines works in Backup Exec	479
Requirements for conversion to virtual machine jobs	487
Converting to a virtual machine simultaneously with a backup job	488
Converting to a virtual machine after a backup job	492
Adding a conversion to virtual machine stage to a backup job	498
Converting to a virtual machine from a point in time	503
Creating a one-time conversion to a virtual machine	505
Setting default options for conversion to virtual machine jobs	509

Chapter 14 Configuration and settings 513

Changing default backup job settings	515
Configuring schedules for backup jobs	519
Configuring storage options for backup jobs	527
Configuring automatic test run jobs for backup jobs	533

Configuring automatic verify operations for backup jobs	534
Configuring delayed catalog operations to improve backup performance	536
Configuring Advanced Open File options for backup jobs	540
Configuring checkpoint restart	543
Configuring Backup Exec to work with Symantec Endpoint Protection	545
Configuring pre/post commands for backup jobs	547
Configuring file and folder options for backup jobs	552
Setting default schedule options for rule-based jobs and run now jobs	559
Excluding dates from the backup schedule for all backups	562
Removing dates from the list of excluded dates	564
Exporting a list of dates that are excluded from all backups to another server	565
Changing the default preferences	565
Configuring the default setting for backing up multiple servers or applications	566
Configuring database maintenance	567
Scheduling Backup Exec to check logon accounts	570
Configuring Backup Exec to discover data to back up	571
Adding discovered servers to the list of servers in Backup Exec	573
Backup networks	574
Changing network and security options for Backup Exec	576
Using Backup Exec with firewalls	578
Backup Exec ports	580
Backup Exec listening ports	582
About enabling a SQL instance behind a firewall	583
Using encryption with Backup Exec	583
Encryption key management	586
Creating encryption keys	587
Replacing an encryption key	588
Deleting encryption keys	590
Granular Recovery Technology	590
Setting default Granular Recovery Technology (GRT) options	596
DBA-initiated job templates	597
Creating DBA-initiated job templates	598
Editing DBA-initiated job templates	599
Deleting DBA-initiated job templates	599
Storage options for DBA-initiated jobs	600
General options for DBA-initiated jobs	605
Network options for DBA-initiated jobs	605
Duplicate job settings for DBA-initiated jobs	606

Backup Exec logon accounts	610
Creating a Backup Exec logon account	612
Editing a Backup Exec logon account	614
Changing the password for a Backup Exec logon account	616
Replacing a Backup Exec logon account	616
Deleting a Backup Exec logon account	617
Changing your default Backup Exec logon account	618
Creating a new Backup Exec System Logon Account	619
Copying logon account information to another Backup Exec server	620
Testing logon accounts	620
Starting and stopping Backup Exec services	621
Changing the credentials for a service account	621
Changing startup options for Backup Exec services	623
Configuring audit logs	624
Viewing the audit log	625
Removing entries from the audit log	625
Saving an audit log to a text file	626
Copying configuration settings to another Backup Exec server	626
Viewing server properties	627

Chapter 15

Reports	629
Reports in Backup Exec	629
Running a report now	631
Scheduling a report	631
Creating a custom report	633
Adding or removing fields on a custom report	637
Changing filters for a custom report	638
Changing the way data is grouped or sorted in a custom report	640
Changing graph options in custom reports	641
Previewing custom reports	641
Copying a custom report	642
Saving a report	642
Printing a report from the Backup Exec Report Viewer	643
Viewing completed reports	643
Editing a report	644
Re-running a completed report	644
Deleting a report	644
Setting defaults for standard and custom reports	645
Viewing report properties	646
List of Backup Exec standard reports	647

Alert History report	652
Alert History By Backup Exec Server report	653
Archive Job Success Rate report	653
Archive Selections by Archive Rules and Retention Categories report	654
Audit Log report	655
Backup Job Success Rate report	655
Backup Recommendations report	655
Backup Resource Success Rate report	656
Backup Sets by Media Set report	656
Backup Size By Resource report	657
Daily Device Utilization report	658
Deduplication Device Summary report	659
Deduplication Summary report	660
Device Summary report	660
Disk Storage Summary report	661
Error-Handling Rules report	662
Event Recipients report	663
Exchange Mailbox Group Archive Settings report	664
Failed Archive Jobs report	664
Failed Backup Jobs report	665
File System Archive Settings report	666
Jobs Summary report	666
Managed Backup Exec Servers report	667
Media Audit report	669
Media Errors report	670
Media Required for Recovery report	670
Media Summary report	671
Media Vault Contents report	671
Move Media to Vault report	672
Operations Overview report	673
Overnight Archive Summary report	675
Overnight Summary report	676
Problem Files report	677
Recently Written Media report	678
Resource Protected Recently report	678
Resource Risk Assessment report	679
Restore Set Details by Resource report	680
Retrieve Media from Vault report	680
Robotic Library Inventory report	681
Scheduled Server Workload report	682
Scratch Media Availability report	682
Test Run Results report	683

	Vault Store Usage Details report	684
	Vault Store Usage Summary report	685
Chapter 16	Troubleshooting Backup Exec	686
	Troubleshooting hardware-related issues in Backup Exec	687
	How to get more information about alerts and error messages	691
	Troubleshooting backup issues in Backup Exec	691
	Troubleshooting failed components in the SAN	693
	Troubleshooting offline storage devices in a SAN	694
	Finding hardware errors in a SAN	696
	Resetting the SAN	696
	Bringing storage devices online after an unsafe device removal event in a SAN	697
	How to improve Backup Exec's performance	698
	Accessing Symantec Online	698
	Searching the Symantec Knowledge Base	699
	Contacting Backup Exec Technical Support	700
	Using Symantec Remote Assistance	700
	Managing your Backup Exec support cases	701
	About Backup Exec diagnostic tools	701
	Running the Symantec Help Tool (SymHelp)	702
	Generating a diagnostic file for troubleshooting Backup Exec	702
	Command line switches for a diagnostic file	704
	Running the begather utility to troubleshoot Backup Exec components on Linux servers	706
	Using the Backup Exec Debug Monitor for troubleshooting	707
	About the Backup Exec debug tool	707
Chapter 17	Using Backup Exec in cluster environments	709
	About Backup Exec and clusters	710
	Requirements for clustering Backup Exec in a Microsoft Cluster Server environment	710
	How Backup Exec works in a Microsoft Cluster Server environment	711
	Requirements for installing Backup Exec on a Microsoft Cluster Server	712
	Installing Backup Exec on a Microsoft Cluster Server	713
	Upgrading Backup Exec on a Microsoft cluster	714
	Installing additional Backup Exec options on a Microsoft cluster	715
	Uninstalling Backup Exec from a Microsoft cluster	716
	Creating storage device pools for Microsoft Cluster Servers	716
	Specifying a different failover node	717

Designating a new central administration server in a Microsoft Cluster Server	718
Configurations for Backup Exec and Microsoft Cluster Servers	719
Two-node cluster with locally attached storage devices	720
Two-node cluster with tape devices on a shared SCSI bus	721
Configuring a shared SCSI bus for tape devices	722
Multi-node clusters on a fibre channel SAN with the Central Admin Server Option	724
Using the Central Admin Server Option with Microsoft clusters and a storage area network	726
About backing up Microsoft Cluster Servers	727
About restoring data to a Microsoft cluster	727
Disaster recovery of a cluster	728
Using Simplified Disaster Recovery to prepare for disaster recovery of a cluster	729
Recovering nodes on the cluster using Simplified Disaster Recovery	729
Recovering Backup Exec on a Microsoft cluster using Simplified Disaster Recovery	730
Recovering the entire cluster using a manual disaster recovery procedure	731
Restoring the Microsoft Cluster Server data files	732
Recovering all shared disks in a Microsoft cluster	732
Recovering Backup Exec in a Microsoft cluster	733
Changing the Quorum disk signature for clusters on Windows Server 2003	733
Manually joining two cluster disk groups and resynchronizing volumes	734
Troubleshooting clusters	734

Chapter 18	Simplified Disaster Recovery	736
	About Simplified Disaster Recovery	736
	Requirements for using Simplified Disaster Recovery	737
	Preparing computers for use with Simplified Disaster Recovery	739
	How to ensure that backups are enabled for Simplified Disaster Recovery	743
	How Simplified Disaster Recovery uses disaster recovery information files	745
	Setting or changing the alternate location for the disaster recovery information file	747
	Changing the default path for the disaster recovery information files	748

Disaster recovery information file data paths	748
Creating a Simplified Disaster Recovery disk image	750
Contents of the Simplified Disaster Recovery disk image	760
Preparing to recover from a disaster by using Simplified Disaster	
Recovery	761
Hardware replacement during disaster recovery	762
Prepare to recover IBM computers with Simplified Disaster	
Recovery	763
Recovering a computer with Simplified Disaster Recovery	764
Recovery notes for using Simplified Disaster Recovery with storage	
pools and storage spaces	769
Recovery notes for using Simplified Disaster Recovery with	
Exchange, SQL, SharePoint, CASO, Hyper-V hosts, and the	
Deduplication Option	771
Advanced Disk Configuration on the Recover This Computer	
Wizard	773

Appendix A	Symantec Backup Exec Agent for Windows	777
	About the Agent for Windows	777
	Requirements for the Agent for Windows	778
	Stopping and starting the Agent for Windows	779
	Establishing a trust between the Backup Exec server and a remote	
	computer	780
	About the Backup Exec Agent Utility for Windows	780
	Starting the Backup Exec Agent Utility	781
	Viewing the activity status of the remote computer in the Backup	
	Exec Agent Utility	782
	Viewing the activity status of the remote computer from the system	
	tray	782
	Starting the Backup Exec Agent Utility automatically on the remote	
	computer	783
	Setting the refresh interval on the remote computer	783
	About publishing the Agent for Windows to Backup Exec	
	servers	784
	Adding Backup Exec servers that the Agent for Windows can	
	publish to	784
	Editing Backup Exec server information that the Agent for Windows	
	publishes to	785
	Removing Backup Exec servers that the Agent for Windows can	
	publish to	786
	Configuring database access for Oracle operations	787

Removing a security certificate for a Backup Exec server that has a trust with the Agent for Windows	789
Using the Backup Exec Agent Utility Command Line Applet	789
Backup Exec Agent Utility Command Line Applet switches	790

Appendix B Symantec Backup Exec Deduplication Option 794

About the Deduplication Option	795
Deduplication methods for Backup Exec agents	797
Requirements for the Deduplication Option	798
Installing the Deduplication Option	803
Configuring an OpenStorage device	804
Editing the properties of an OpenStorage device	806
Specifying a Backup Exec server that has proximity to a shared OpenStorage device	808
Creating or importing deduplication disk storage	809
Editing the properties of a deduplication disk storage device	811
Changing the password for the logon account for deduplication disk storage	816
Selecting storage devices for direct access sharing	817
Editing server properties for direct access	818
Changing the location of a deduplication disk storage device	819
Sharing a deduplication device between multiple Backup Exec servers	821
How to use client-side deduplication	821
How to set up backup jobs for deduplication	822
Using optimized duplication to copy deduplicated data between OpenStorage devices or deduplication disk storage devices	823
Copying deduplicated data to tapes	826
Using deduplication with encryption	826
Restoring a deduplication disk storage device or deduplicated data	827
Disaster recovery of deduplication disk storage devices	827
Disaster recovery of OpenStorage devices	828

Appendix C Symantec Backup Exec Agent for VMware 830

About the Agent for VMware	830
Requirements for using the Agent for VMware	831
Using the Agent for VMware with Windows Server 2012	832
Using the Agent for VMware with vSphere 5.5	833
About installing the Agent for VMware	833
Adding VMware vCenter servers and ESX/ESXi hosts to the list of servers on the Backup and Restore tab	834

Viewing details about VMware resources	835
Installing the Agent for Windows on VMware virtual machines	836
Push-installing the Agent for Windows to VMware virtual machines	837
Backing up VMware data	838
Setting default backup options for virtual machines	843
How Backup Exec automatically backs up new VMware virtual machines during a backup job	847
Using Granular Recovery Technology (GRT) with the Agent for VMware	847
How delayed cataloging works with VMware virtual machine backups	851
Restoring VMware virtual machines and vmdk files	852

Appendix D

Symantec Backup Exec Agent for Microsoft Hyper-V	860
About the Agent for Microsoft Hyper-V	860
Requirements for using the Agent for Microsoft Hyper-V	862
About installing the Agent for Microsoft Hyper-V	864
Using the Agent for Hyper-V with Windows Server 2012/2012 R2	865
Disk space optimization with the Agent for Hyper-V	868
Adding a Hyper-V host to the list of servers on the Backup and Restore tab	869
Viewing details about Hyper-V resources	870
Installing the Agent for Windows on Hyper-V virtual machines	871
Push-installing the Agent for Windows to Hyper-V virtual machines	872
Backing up Microsoft Hyper-V virtual machines	873
Setting default backup options for Hyper-V	878
How Backup Exec automatically protects new virtual machines during a backup job	880
Using Granular Recovery Technology (GRT) with the Agent for Hyper-V	881
About backing up and restoring highly available Hyper-V virtual machines	883
How delayed cataloging works with Hyper-V virtual machine backups	883
Restoring Microsoft Hyper-V virtual machines	885

Appendix E	Symantec Backup Exec Agent for Microsoft SQL Server	890
	About the Agent for Microsoft SQL Server	890
	Requirements for using the SQL Agent	892
	About installing the SQL Agent	893
	Backup strategies for SQL	893
	Adding SQL Servers to the list of servers on the Backup and Restore tab	895
	Configuring Backup Exec to run a consistency check before every SQL backup	895
	Using snapshot technology with the SQL Agent	896
	Using database snapshots for SQL Server	897
	Backing up SQL databases and transaction logs	899
	Setting default backup options for SQL Server	911
	Restoring SQL databases and transaction logs	921
	Restoring the SQL master database	923
	Restarting SQL using database copies	924
	Disaster recovery of a SQL Server	926
	Manual recovery of a SQL Server	928
Appendix F	Symantec Backup Exec Agent for Microsoft Exchange Server	929
	About the Backup Exec Exchange Agent	930
	Requirements for using the Exchange Agent	931
	Granting permissions on the Exchange Server to enable database backups and restores, and Granular Recovery Technology operations	935
	About installing the Exchange Agent	939
	Adding Exchange Servers and database availability groups to the list of servers on the Backup and Restore tab	939
	Managing preferred server configurations for Microsoft Exchange Database Availability Groups	940
	Recommended configurations for Exchange	943
	Requirements for accessing Exchange mailboxes	944
	Backup strategies for Exchange	945
	How Granular Recovery Technology works with the Exchange Information Store	947
	Snapshot and offhost backups with the Exchange Agent	948
	Backing up Exchange data	950
	Setting default backup options for Exchange Server	957
	Restoring Exchange data	962

	Disaster recovery of an Exchange Server	968
Appendix G	Symantec Backup Exec Agent for Microsoft SharePoint	970
	About the Agent for Microsoft SharePoint	971
	About installing the Agent for Microsoft SharePoint	971
	Requirements for the Agent for Microsoft SharePoint	971
	Using the Agent for Microsoft SharePoint with SharePoint Server 2010/2013 and SharePoint Foundation 2010/2013	972
	Using the Agent for Microsoft SharePoint with SharePoint Server 2007 and Windows SharePoint Services 3.0	973
	Adding a Microsoft SharePoint server farm to the list of servers on the Backup and Restore tab	974
	Backing up Microsoft SharePoint data	974
	Setting default backup options for SharePoint	978
	Restoring Microsoft SharePoint data	981
	Disabling or enabling communication between a Microsoft SharePoint web server and Backup Exec	983
	Viewing or changing SharePoint farm properties	984
	Disaster recovery of Microsoft SharePoint 2010/2013 data	985
	Disaster recovery of Microsoft SharePoint 2007 data	989
Appendix H	Symantec Backup Exec Agent for Oracle on Windows or Linux Servers	992
	About the Backup Exec Oracle Agent	992
	About installing the Oracle Agent	993
	Configuring the Oracle Agent on Windows computers and Linux servers	994
	Configuring an Oracle instance on Windows computers	995
	Viewing an Oracle instance on Windows computers	997
	Editing an Oracle instance on Windows computers	998
	Deleting an Oracle instance on Windows computers	999
	Enabling database access for Oracle operations on Windows computers	999
	Configuring an Oracle instance on Linux servers	1001
	Viewing an Oracle instance on Linux servers	1002
	Editing an Oracle instance on Linux servers	1003
	Deleting an Oracle instance on Linux servers	1003
	Enabling database access for Oracle operations on Linux servers	1004
	About authentication credentials on the Backup Exec server	1005

Setting authentication credentials on the Backup Exec server for Oracle operations	1006
Deleting an Oracle server from the Backup Exec server's list of authentication credentials	1007
About Oracle instance information changes	1008
About backing up Oracle databases	1008
About backing up Oracle RAC databases	1009
About performing a DBA-initiated backup job for Oracle	1010
Oracle backup options	1011
About restoring Oracle resources	1012
About DBA-initiated restore for Oracle	1014
About redirecting a restore of Oracle data	1014
Requirements for recovering the complete Oracle instance and database using the original Oracle server	1015
Recovering the complete Oracle instance and database using the original Oracle server	1016
Requirements for recovering the complete Oracle instance or database to a computer other than the original Oracle server	1017
Recovering the complete Oracle instance or database to a computer other than the original Oracle server	1017

Appendix I	Symantec Backup Exec Agent for Enterprise Vault	1019
	About the Agent for Enterprise Vault	1019
	Requirements for the Enterprise Vault Agent	1023
	About installing the Enterprise Vault Agent	1024
	About backup methods for Enterprise Vault backup jobs	1024
	Enterprise Vault backup options	1027
	About backing up Enterprise Vault components	1028
	About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases	1031
	About restoring Enterprise Vault	1031
	Enterprise Vault restore options	1033
	About restoring individual files and folders with the Enterprise Vault Agent	1036
	About automatic redirection of Enterprise Vault components under an Enterprise Vault server	1036
	About redirecting a restore for an Enterprise Vault component	1037
	Configuring Enterprise Vault to use the name of the new SQL Server that holds the Directory database	1039

Best practices for the Enterprise Vault Agent	1041
About the Backup Exec Migrator for Enterprise Vault	1041
Backup Exec Migrator for Enterprise Vault requirements	1042
How the Backup Exec Migrator works	1042
About using staged migrations with Backup Exec and the Backup Exec Migrator	1046
About Backup Exec Migrator events	1047
About Backup Exec Migrator logs	1047
About deleting files migrated by Backup Exec Migrator	1048
Configuring the Backup Exec Migrator	1049
About viewing migrated Enterprise Vault data	1057
About retrieving migrated Enterprise Vault data	1058
About the Partition Recovery Utility	1059

Appendix J

Symantec Backup Exec Agent for Lotus Domino	1063
About the Agent for Lotus Domino Server	1064
Lotus Domino Agent requirements	1064
About installing the Lotus Domino Agent on the Backup Exec server	1066
About the Lotus Domino Agent and the Domino Attachment and Object Service (DAOS)	1066
Best practices for restoring the missing .nlo files	1067
Viewing Lotus Domino databases	1068
About backing up Lotus Domino databases	1068
About selecting backup options for Lotus Domino databases	1070
Lotus Domino Agent backup options	1070
About automatic exclusion of Lotus Domino files during volume-level backups	1071
About supported Lotus Domino database configurations	1071
About Lotus Domino transaction logs	1072
About restoring Lotus Domino databases	1073
Restoring Lotus Domino databases	1074
Lotus Domino Agent restore options	1076
About redirecting restore jobs for Lotus Domino databases	1078
About redirecting the restore of DAOS NLO files	1079
How to prepare for disaster recovery on a Lotus Domino server	1079
Recovering a Lotus Domino server from a disaster	1079
About disaster recovery of a Lotus Domino server using archive logging	1081
Disabling the monitor change journal	1081

	Recovering the Lotus Domino server, databases, and transaction logs when archive logging is enabled	1082
	Re-enabling the monitor change journal	1083
	Recovering a Lotus Domino server that uses circular logging	1084
Appendix K	Symantec Backup Exec Agent for Microsoft Active Directory	1086
	About the Agent for Microsoft Active Directory	1086
	Requirements for the Agent for Microsoft Active Directory	1087
	About backing up Active Directory and ADAM/AD LDS	1088
	Editing options for Active Directory and ADAM/AD LDS backup jobs	1089
	Microsoft Active Directory backup job options	1091
	About restoring individual Active Directory and ADAM/AD LDS objects	1092
Appendix L	Symantec Backup Exec Central Admin Server Option	1095
	About the Central Admin Server Option	1096
	Requirements for installing CASO	1098
	How to choose the location for CASO storage and media data	1099
	About installing the Central Admin Server Option	1101
	Push-installing a managed Backup Exec server from the central administration server	1102
	Managed Backup Exec Server Configuration options	1106
	Installing a managed Backup Exec server across a firewall	1107
	Upgrading an existing CASO installation	1110
	Changing a Backup Exec server to a central administration server	1112
	Changing a Backup Exec server to a managed Backup Exec server	1113
	Deleting a managed Backup Exec server from a CASO environment	1114
	Renaming a central administration server	1116
	Renaming a managed Backup Exec server	1117
	How to reduce network traffic in CASO	1118
	CASO distributed, centralized, and replicated catalog locations	1118
	Changing the settings for a managed Backup Exec server	1120
	What happens when CASO communication thresholds are reached	1131

Enabling or disabling communications between the managed Backup Exec server and the central administration server	1132
Alerts and notifications in CASO	1132
Enabling managed Backup Exec servers to use any available network interface card	1133
About job delegation in CASO	1134
About copying jobs instead of delegating jobs in CASO	1134
About adding storage devices in a CASO environment	1135
How data lifecycle management (DLM) works in a CASO environment	1135
Obtaining media audit information for a managed Backup Exec server	1136
How to use Backup Exec server pools in CASO	1136
Selecting a Backup Exec server pool for backups	1137
Creating a Backup Exec server pool	1138
Adding managed Backup Exec servers to a Backup Exec server pool	1138
Deleting a Backup Exec server pool	1139
Removing a managed Backup Exec server from a Backup Exec server pool	1139
How centralized restore works in CASO	1140
How CASO restores data that resides on multiple storage devices	1141
About recovering failed jobs in CASO	1142
Pausing or resuming a managed Backup Exec server	1144
Stopping or starting Backup Exec services for a managed Backup Exec server	1144
Viewing managed Backup Exec server properties	1145
Viewing the settings for a central administration server	1148
Disaster recovery in CASO	1150
Troubleshooting CASO	1151
Running the Backup Exec Utility for CASO operations	1153
Uninstalling Backup Exec from the central administration server	1153
Uninstalling Backup Exec from a managed Backup Exec server	1154

Appendix M

Symantec Backup Exec Advanced Disk-based Backup Option	1155
About the Advanced Disk-based Backup Option	1155
How to use synthetic backups in place of recurring full backups	1156
Configuring off-host backup options for a backup job	1157

Setting default backup options for the Advanced Disk-based Backup	
Option	1159
About true image restore for synthetic backups	1161
How to use off-host backup to process remote computer backups on	
the Backup Exec server	1163
Troubleshooting off-host backups	1166

Appendix N	Symantec Backup Exec NDMP Option	1170
	Features of the NDMP Option	1170
	Requirements for using the NDMP Option	1171
	About installing the NDMP Option	1172
	Adding NDMP servers to Backup Exec	1172
	Sharing the tape drives on NDMP servers between multiple Backup	
	Exec servers	1174
	NDMP backup options for NDMP servers	1174
	NDMP server backup selections	1176
	How to use patterns to exclude files and directories from an NDMP	
	server backup	1177
	Supported configurations for duplicating data from NDMP	
	servers	1179
	About restoring and redirecting restore data for NDMP servers	1179
	NDMP server restore options	1180
	Setting the default backup options for the NDMP Option	1183
	Viewing the properties of an NDMP server	1184
	NDMP server properties	1184
	Viewing storage properties for an NDMP server	1185
	Storage properties for an NDMP server	1185

Appendix O	Symantec Backup Exec Agent for Linux	1187
	About the Agent for Linux	1188
	About open files and the Agent for Linux	1188
	Requirements for the Agent for Linux	1188
	About installing the Agent for Linux	1189
	Installing the Agent for Linux	1190
	About the Backup Exec operators group for the Agent for	
	Linux	1192
	About establishing trust for a remote Linux computer in the Backup	
	Exec list of servers	1193
	Establishing trust and adding a remote Linux computer to the	
	Backup Exec list of servers	1193
	Adding additional Backup Exec servers to which the Agent for Linux	
	can publish information	1194

About configuring the Agent for Linux	1195
About excluding files and directories from backup jobs for Linux computers	1196
Editing configuration options for Linux computers	1196
Configuration options for Linux computers	1197
About backing up a Linux computer by using the Agent for Linux	1204
Linux backup options	1205
About backing up Novell Open Enterprise Server (OES) components	1206
About restoring data to Linux computers	1207
About restoring Novell OES components	1208
Restore job options for Linux computers	1208
Edit the default backup job options for Linux computers	1209
Default backup job options for Linux computers	1209
Uninstalling the Agent for Linux	1211
Manually uninstalling the Agent for Linux	1212
Run-time scripts to remove when manually uninstalling the Agent for Linux	1213
Starting the Agent for Linux daemon	1214
Stopping the Agent for Linux daemon	1215
Troubleshooting the Agent for Linux	1215

Appendix P

Symantec Backup Exec Agent for Mac	1218
About the Agent for Mac	1219
Requirements for the Agent for Mac	1219
About the Backup Exec admin group on Macintosh systems	1219
Creating the Backup Exec admin group manually on Macintosh systems	1220
About installing the Agent for Mac	1221
Installing the Agent for Mac	1221
Uninstalling the Agent for Mac	1223
Manually uninstalling the Agent for Mac	1225
About configuring the Agent for Mac	1226
Editing configuration options for Macintosh computers	1227
Configuration options for Macintosh computers	1227
Starting the Agent for Mac	1232
Stopping the Agent for Mac	1233
About establishing trust for a remote Macintosh system in the Servers list	1233
Establishing trust and adding a remote Macintosh computer to the Backup Exec Servers list	1233

Adding additional Backup Exec servers to which the Agent for Mac can publish information	1234
About backing up data by using the Agent for Mac	1235
Editing the default backup options for Macintosh systems	1235
About excluding files and directories from backup jobs for Macintosh computers	1237
About restoring Macintosh systems	1238
Macintosh restore options	1238
Troubleshooting the Agent for Mac	1238

Appendix Q

Symantec Backup Exec Remote Media Agent for Linux	1240
About the Remote Media Agent for Linux	1241
How the Remote Media Agent for Linux works	1242
Requirements for the Remote Media Agent for Linux	1242
About open files and the Remote Media Agent for Linux	1243
About installing the Remote Media Agent for Linux	1243
Installing the Remote Media Agent for Linux	1244
Uninstalling the Remote Media Agent for Linux	1246
Starting the Remote Media Agent for Linux daemon	1247
Stopping the Remote Media Agent for Linux daemon	1248
About establishing trust for a Remote Media Agent for Linux computer in the Backup Exec list of servers	1248
Establishing trust and adding a Remote Media Agent for Linux computer to the Backup Exec list of servers	1249
Adding additional Backup Exec servers to which Remote Media Agent for Linux can publish	1249
Finding simulated tape library files	1250
About the Backup Exec operators group for the Remote Media Agent for Linux	1251
Creating the Backup Exec operators group manually for the Remote Media Agent for Linux	1251
About adding a Linux server as a Remote Media Agent for Linux	1252
Adding a Linux server as a Remote Media Agent for Linux	1252
Remote Media Agent for Linux options	1253
Changing the port for communications between the Backup Exec server and the Remote Media Agent for Linux	1254
About creating storage device pools for devices attached to the Remote Media Agent for Linux	1255
Editing properties for the Remote Media Agent for Linux	1255
Remote Media Agent for Linux properties	1255

Deleting a Remote Media Agent for Linux from the Backup Exec list of servers	1256
Sharing a Remote Media Agent for Linux between multiple Backup Exec servers	1257
About backing up data by using the Remote Media Agent for Linux	1257
About restoring data by using the Remote Media Agent for Linux	1257
About the Tape Library Simulator Utility	1258
Creating a simulated tape library	1259
Simulated Tape Library options	1259
Viewing simulated tape libraries properties	1260
Simulated tape library properties	1261
Deleting a simulated tape library	1261
Managing simulated tape libraries from the command line	1262
Command line switches for the Tape Library Simulator Utility	1263
Troubleshooting the Remote Media Agent for Linux	1264

Appendix R

Symantec Backup Exec Storage Provisioning Option	1267
Features of the Storage Provisioning Option	1268
Requirements for the Storage Provisioning Option	1269
About installing the Storage Provisioning Option	1269
Viewing storage array components in Backup Exec	1270
Configuring a storage array by using the Configure Storage Wizard	1270
Viewing properties for a storage array and its physical disks	1272
Properties for a storage array and its physical disks	1272
About the Any Virtual Disk Storage device pool in the Storage Provisioning Option	1278
About virtual disks in the Storage Provisioning Option	1278
Editing default options for a virtual disk on a storage array	1280
Configuring a virtual disk on a storage array	1285
Viewing properties for unconfigured virtual disks on a storage array	1285
Properties for unconfigured virtual disks on storage arrays	1285
Editing properties of virtual disks on storage arrays	1292
Properties for virtual disks on storage arrays	1297
Adding a hot spare by using the Configure Storage Wizard	1307
Changing a hot spare by using the Configure Storage Wizard	1308
Detecting a new storage array	1309
Renaming a virtual disk or storage array	1309

Identifying the physical disks of a virtual disk by using the blink feature	1310
Troubleshooting the Storage Provisioning Option	1311
Appendix S Symantec Backup Exec Archiving Option	1312
About the Archiving Option	1313
Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option	1314
Requirements for the Exchange Mailbox Archiving Option	1316
Requirements for the File System Archiving Option	1319
About granting permissions on the Exchange Server for the Backup Exec service account in the Archiving Option	1320
How to calculate disk space requirements for the Exchange Mailbox Archiving Option	1326
How to calculate disk space requirements for the File System Archiving Option	1329
Installing the Backup Exec Archiving Option	1332
About using the command line to install the Backup Exec Exchange Mailbox Archiving Option	1333
About the Internet Information Services (IIS) 7.0 role services installed by Backup Exec	1333
About uninstalling or reinstalling the Archiving Option	1335
About installing Enterprise Vault on a Backup Exec server on which the Archiving Option is installed	1336
About Enterprise Vault services for the Archiving Option	1336
Repairing the Backup Exec Archiving Option	1337
How the Archiving Option works	1337
Types of data not included in archive jobs	1338
Best practices for the Archiving Option	1339
About creating an archive job	1340
Viewing the servers that have archiving jobs	1341
Setting archive job options	1341
Editing default settings for the Archiving Option	1346
Default settings for the Archiving Option	1346
About single instance storage of archived items	1349
Enabling single instance storage of archived items	1350
About synchronizing archive permissions and settings	1350
About vault stores in the Archiving Option	1351
Editing or viewing vault store properties	1352
Vault store properties	1352
About deleting an Archiving Option vault store	1353
Deleting a vault store	1354

About vault store partitions in the Archiving Option	1354
Editing vault store partition properties	1355
Vault store partition properties	1355
About archives in the Archiving Option	1356
Editing archive properties	1356
Deleting an archive	1357
Archive properties	1357
About Archiving Option operation entries in the audit log	1358
About Exchange mailbox groups in archive jobs	1358
Managing Exchange mailbox groups	1359
Manage mailbox groups options	1360
Mailbox group options for mailbox group details	1361
About deleting archived data from its original location	1363
About archive settings in the Archiving Option	1363
Manage Archive Settings options	1363
Archive Settings Details	1364
About retention categories for archived items	1366
About managing index locations in the Backup Exec Archiving	
Option	1370
Viewing index locations	1371
Adding a new index location	1371
Deleting an index location	1372
Opening or closing an index location	1372
Manage Index Locations options	1373
About restoring items from the archives	1374
About searching for data in the archives	1375
About deleting data from the archives	1375
Deleting data from the archives	1376
Delete from Archive options	1376
About backing up and restoring the Archiving Option components from	
a remote Backup Exec server	1377
Preventing the deletion of expired archived items from an archive	1377
About backing up Archiving Option components	1377
Archive Option backup options	1380
About consistency checks for Archiving Option databases	1381
About disabling backup mode for Archiving Option	
components	1381
About restoring an Archiving Option component	1382
About running the Backup Exec Utility to complete redirected	
restores of Archiving Option components	1383
Running the Backup Exec Utility to complete a redirected restore	
of the Archiving Option databases	1384

Running the Backup Exec Utility before redirecting the restore of an Archiving Option vault store partition	1385
Running the Backup Exec Utility before redirecting the restore of Archiving Option index files	1385
Running the Backup Exec Utility to complete Archiving Option operations	1386
About moving Archiving Option components to a new location	1387
Troubleshooting archive jobs	1388
Viewing the Enterprise Vault event log for Archiving Option events	1388
Reports for the Archiving Option	1388
About Backup Exec Virtual Vault	1389
About the Backup Exec Outlook Add-In	1391
About installing the Backup Exec Outlook Add-In to end users' computers	1391
About configuring Outlook in cached Exchange mode	1396
Best practices for Virtual Vault	1396
About the Vault Cache	1397
About Vault Cache synchronization	1398
About preemptive caching	1399
About the temporary cache location on the Backup Exec server	1400
About enabling and disabling Virtual Vault	1401

Appendix T	Accessibility and Backup Exec	1403
	About accessibility and Backup Exec	1403
	About keyboard shortcuts in Backup Exec	1404
	Home tab keyboard shortcuts	1405
	Backup and Restore tab keyboard shortcuts	1405
	Job Monitor tab keyboard shortcuts	1413
	Storage tab keyboard shortcuts	1416
	Reports tab keyboard shortcuts	1424
	General keyboard navigation within the Backup Exec user interface	1425
	Keyboard navigation within dialog boxes in Backup Exec	1425
	List box navigation in Backup Exec	1426
	Tabbed dialog box navigation in Backup Exec	1426
	About setting accessibility options	1427

Glossary	1428
----------------	------

Index	1435
-------------	------

Introducing Backup Exec 2014

This chapter includes the following topics:

- [About Backup Exec](#)
- [How Backup Exec works](#)
- [What's new in Backup Exec 2014](#)
- [What's new in Backup Exec 2014 agents and options](#)

About Backup Exec

Symantec Backup Exec 2014 is a high-performance data management solution for Windows® servers networks. With its client/server design, Backup Exec provides fast, reliable backup and restore capabilities for servers, applications, and workstations across the network.

Backup Exec is available in several configurations that can accommodate networks of all sizes. In addition, Backup Exec's family of agents and options offer solutions for scaling your Backup Exec environment and extending platform and feature support.

You can find more information about Backup Exec's host of solutions at the following URL:

www.BackupExec.com

See [“What's new in Backup Exec 2014 ”](#) on page 39.

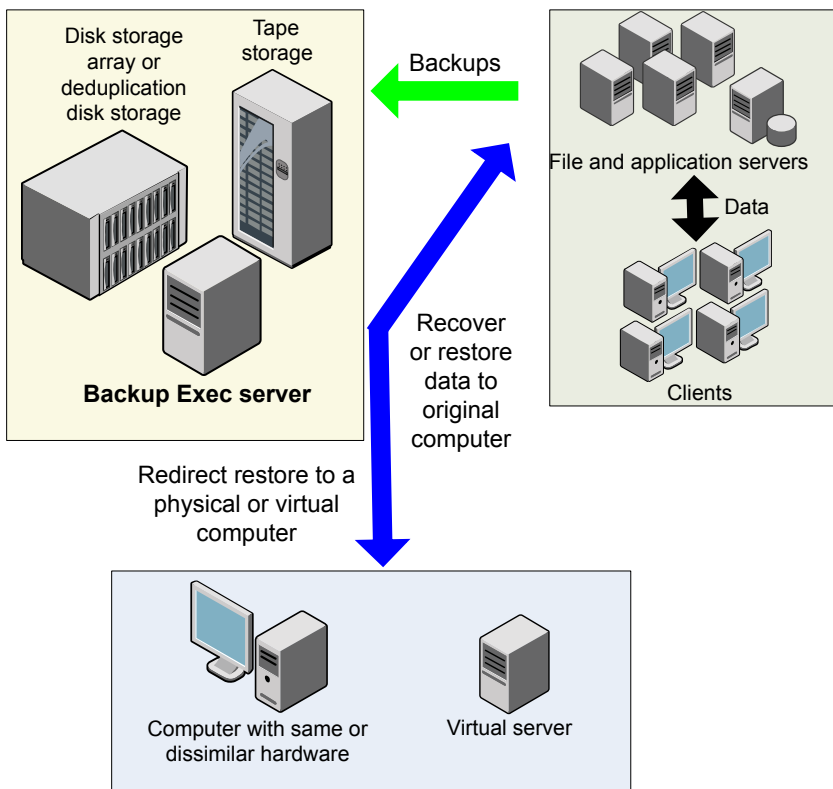
See [“What's new in Backup Exec 2014 agents and options”](#) on page 48.

See [“How Backup Exec works”](#) on page 38.

How Backup Exec works

You use the Backup Exec Administration Console to interact with Backup Exec for tasks like submitting backups, restoring data, configuring storage, and monitoring jobs. You can run the Administration Console from the Backup Exec server, which is a Windows server on which Backup Exec is installed, or from a remote computer. After backups, restores, or other operations are created, the Backup Exec server processes the jobs or delegates the jobs for processing in multi-Backup Exec server environments.

Figure 1-1 Backup and restore functionality for the entire network



Backup Exec includes the following features:

- Convenient backup scheduling
Backup Exec administrators can set up scheduled backups for Windows and Linux computers across the network. Backup Exec's flexible calendar-based administration lets you easily schedule backups for processing during off-peak hours.

- **Complete system recovery**
Backup Exec's Simplified Disaster Recovery takes all of the guesswork out of recovering an entire system. While configuring a backup, you get a clear indication that you have selected the data that is necessary to perform a Simplified Disaster Recovery-enabled backup. After you back up a computer's critical system components, use the **Create Simplified Disaster Recovery Disk Wizard** to create a Simplified Disaster Recovery disk image. You can then use the disk image to perform a disaster recovery of the computer.
- **Comprehensive monitoring and intuitive mechanisms for everyday tasks**
The **Job Monitor** provides a single location to monitor and manage all of your backup, restore, installation, and storage operation jobs. The **Home** tab lets you view statistics for your entire Backup Exec environment. From the **Servers** view, you can monitor the backup status for all of the computers on your network. Interactive alerts display the situations that require your attention.
Backup results can be viewed from a backup's job history. The job history contains statistics, errors, and other information pertaining to the backup. Backup Exec's catalog is a database of all backed-up data, and Backup Exec uses the catalog to track restore selections.
Wizards guide you through most Backup Exec operations, including the creation of backup and restore jobs, and the configuration of storage.
- **Automated data lifecycle management for disk-based storage**
Backup Exec uses data lifecycle management to automatically delete expired backup sets on disk storage, disk cartridge media, deduplication storage, storage arrays, and virtual disks. You specify how long to keep backup data when you create a backup job that is sent to a disk-based storage device. When the amount of time to keep the backup data expires, the data lifecycle management feature deletes the backup sets and reclaims the disk space for use by new backup sets.
See ["How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage"](#) on page 335.
- **Simplified device and media management**
Backup Exec uses the Advanced Device and Media Management (ADAMM) feature to manage data retention on tapes. ADAMM expires the backup sets that are stored on media according to a set of rules that you apply to the media.
See ["What's new in Backup Exec 2014"](#) on page 39.
See ["What's new in Backup Exec 2014 agents and options"](#) on page 48.

What's new in Backup Exec 2014

This release of Backup Exec includes the following new features and capabilities:

Table 1-1 What's new in Backup Exec 2014

Item	Description
Enhanced conversion to a virtual machine	<p>The following enhancements were made to the conversion to a virtual machine feature:</p> <ul style="list-style-type: none"> ■ Support for incremental backups. ■ Standalone conversion, which is a conversion that runs without a backup. ■ Support for disks that are larger than 2 TB (for VMware or for Hyper-V 2012 or later). ■ Support for exclusion of non-critical data from the conversion. <p>See "How conversion of physical computers to virtual machines works in Backup Exec" on page 479.</p>
Push-installation of the Agent for Windows to virtual machines	<p>The Add Server wizard lets you push-install the Agent for Windows to the virtual machines that you select. You can choose to install the Agent for Windows on all of the guest virtual machines that need it or on specific machines only.</p> <p>See "Push-installing the Agent for Windows to VMware virtual machines" on page 837.</p> <p>See "Push-installing the Agent for Windows to Hyper-V virtual machines" on page 872.</p>
Enhanced V-Ray Edition	<p>The V-Ray Edition is designed to back up and restore virtual machines and virtualized applications in VMware and Hyper-V environments. It includes the Deduplication Option, the Agent for Applications and Databases, and the Agent for VMware and Hyper-V. It is licensed according to the number of cores on the host. The Backup Exec V-Ray Tier A license is available for hosts that have 6 or fewer cores in a socket. The Backup Exec V-Ray Tier B license is available for hosts that have 8 or more cores in a socket.</p>

Table 1-1 What's new in Backup Exec 2014 (*continued*)

Item	Description
Updated job scheduling	<p>Backup Exec has an updated scheduling feature which offers improved flexibility and more control over your job scheduling. You can now select any specific days of the month to form a pattern that reoccurs every month. You can exclude individual dates for jobs. You can also restrict hourly recurrences to a designated block of time, such as business hours.</p> <p>See “How job scheduling works in Backup Exec” on page 215.</p>
Backup source ordering	<p>When you create a selection list for a job, you can now assign the order in which backup sources are backed up. Adjusting the order can be helpful if you need a specific source to be backed up before others, for example.</p> <p>See “Changing the order in which backup sources are backed up” on page 175.</p>
Editing multiple backup selections	<p>You can now edit the detailed backup selections for multiple jobs at the same time. Previously, you could only edit the selections that the backups had in common.</p> <p>See “Editing backup definitions” on page 206.</p>
Backing up multiple servers or applications	<p>You can back up multiple servers or applications at once using Backup Exec. You can back them up as part of one backup definition or you can back them up individually in separate backup definitions.</p> <p>See “Backing up data” on page 158.</p>

Table 1-1 What's new in Backup Exec 2014 (*continued*)

Item	Description
Mixed backup methods	<p>You can now create backup definitions that use both the incremental and differential backup methods in the same backup definition. Additionally, you can select agent or data-specific backup methods. For example, you can use a specific backup method for Microsoft SQL data that backs up and truncates the transaction log. You can mix full, incremental, and differential backup methods within a single job template to back up multiple types of data using data-specific backup methods.</p> <p>See "Backup methods in Backup Exec" on page 188.</p>
Delayed catalog operations	<p>You can configure a delayed catalog operation for any backups that are enabled for Granular Recovery Technology (GRT). GRT-enabled backup jobs require more time to catalog because of the granular information that they contain. You can now run the catalog operation as a separate job or you can schedule the catalog operation to avoid conflicts with your backup resources. You can also choose to not run a delayed catalog operation.</p> <p>See "Configuring delayed catalog operations to improve backup performance" on page 536.</p>

Table 1-1 What's new in Backup Exec 2014 (*continued*)

Item	Description
Improved credentials management	<p>You can now select a logon account and test it for use with a server or other backup source. Testing a logon account can help you determine if you have the appropriate credentials for a backup job before you attempt to run the job. The Logon Account Wizard now includes more guidance to help you select and configure the right kind of logon account for your backup job.</p> <p>Backup Exec also automatically checks your credentials in the background while you work to ensure that you can access your backup sources. If Backup Exec detects that you cannot access a backup source because of an issue with your credentials, it sends you an alert.</p> <p>See “Testing credentials for backup sources” on page 211.</p>
Restore Wizard enhancements	<p>You can start restore jobs by launching the Restore Wizard directly from the following:</p> <ul style="list-style-type: none">■ From a job in the Jobs views.■ From a backup set in the Backup Sets view.■ From media in the Media view. <p>See “Methods for restoring data in Backup Exec” on page 232.</p>

Table 1-1 What's new in Backup Exec 2014 *(continued)*

Item	Description
Support for user-mode drivers	<p>Symantec kernel-mode tape drivers and tapeinst.exe are no longer installed if Backup Exec 2014 is installed on Windows Server 2012 and later. The Backup Exec 2014 hardware compatibility list (HCL) does not list any drivers that require kernel-mode drivers. If you use older or unsupported tape devices, test the devices before you use Windows Server 2012 and later to determine the functionality in case kernel-mode drivers are required.</p> <p>Symantec user-mode drivers now access all supported tape devices. User-mode drivers have accessed most tape devices for multiple previous versions of Backup Exec.</p> <p>User-mode drivers are more stable, and are patched without restarting the server. Hardware devices that run user-mode drivers are easier to support and maintain.</p> <p>See "Installing Symantec tape device drivers" on page 360.</p>

Table 1-1 What's new in Backup Exec 2014 (*continued*)

Item	Description
New information available on the Storage tab	<p>Two new columns appear on the Storage view, and provide the following information:</p> <ul style="list-style-type: none"> ■ Active and Scheduled jobs Displays the number of storage operation jobs, backup jobs, and the restore jobs that run on this device. Click the text in this column to see details of the jobs that are running or that are scheduled to run. ■ Parent device Identifies the parent device if there is a relationship between devices. Examples include a relationship between a tape library and a tape drive, or between a storage array and a virtual disk. You must select the List view to see this column. <p>Additionally, storage capacity monitoring is now available for disk cartridge media.</p> <p>See “Storage tab overview in Backup Exec” on page 425.</p>
Support for scheduling some storage operation jobs	<p>You can schedule some storage operations as recurring jobs. You can also specify a schedule and a recipient for notification when these jobs run.</p> <p>Storage operations that you can schedule include the following operations:</p> <ul style="list-style-type: none"> ■ Scan ■ Inventory ■ Eject ■ Clean ■ Import media ■ Export media ■ Export expired media ■ Unlock a robotic library's front portal ■ Erase <p>See “About storage operation jobs” on page 424.</p>

Table 1-1 What's new in Backup Exec 2014 *(continued)*

Item	Description
Automatic deletion of expired backup sets and reclamation of disk space on disk-based storage, including disk cartridge media	<p>Backup Exec now uses a disk reclamation process called data lifecycle management (DLM) instead of Advanced Device and Media Management (ADAMM) to manage the data retention of backup sets on disk-based storage. DLM automatically deletes expired backup sets on disk-based storage and reclaims the disk space for use by new backup sets. DLM deletes backup sets after the amount of time expires that you specified when you created a backup job.</p> <p>Disk-based storage includes disk storage, disk cartridge media, deduplication storage, storage arrays, and virtual disks.</p> <p>Previous versions of Backup Exec let the backup sets expire but did not delete them from the disk, which can cause your disk storage to become full and jobs to fail. Backup Exec now proactively frees disk space by using data lifecycle management.</p> <p>See “How data lifecycle management (DLM) deletes expired backup sets on disk-based storage” on page 335.</p>

Table 1-1 What's new in Backup Exec 2014 (*continued*)

Item	Description
Support for Windows Server 2012/8/Server 2012 R2/8.1	<p>You can install a Backup Exec server on Windows Server 2012/Server 2012 R2 computers. Additionally, you can use the Agent for Windows to back up and restore data from Windows Server 2012/Server 2012 R2 computers.</p> <p>Note: You cannot install a Backup Exec server or the Agent for Windows on ReFS volumes or any volumes that have been enabled for deduplication in Windows.</p> <p>You can use the Agent for Windows to back up and restore data from Windows 8/8.1 computers.</p> <p>Note: Backup Exec does not support restoring the WindowsApps folder for Windows Server 2012 R2/8.1 computers on which Desktop Experience is enabled.</p> <p>Note: Backup Exec does not support restoring WIMBoot based systems.</p>
Support for Windows Server 2012 in Simplified Disaster Recovery	<p>Simplified Disaster Recovery (SDR) supports Windows Server 2012.</p> <p>See “About Simplified Disaster Recovery” on page 736.</p> <p>For specific operating system requirements for Simplified Disaster Recovery, see the Backup Exec Software Compatibility List at the following URL: http://entsupport.symantec.com/umi/v-269-1</p>
Job Monitor tab	<p>The Job Monitor provides a single location to monitor and manage all of your backup, restore, installation, and storage operation jobs.</p> <p>See “About the Job Monitor” on page 255.</p>

Table 1-1 What's new in Backup Exec 2014 *(continued)*

Item	Description
Details pane on the Backup and Restore tab	<p>The new Details pane is located on the lower portion of the Backup and Restore tab, under the list of servers. The Details pane provides additional details for the server that is selected in the list of servers. The Details pane lists the resources for the selected server and the selection status, backup status, and logon account information for each resource.</p> <p>When a Hyper-V host or VMware host is selected from the list of servers, the Details pane includes additional information and functionality. It provides details about the last 7 days of backup jobs, the date of the last backup, and the date of the next scheduled backup. In addition, it includes the ability to restore data and to filter the list of guest virtual machines.</p> <p>Note that the Details pane for Hyper-V and VMware provides a combination of current information and historical information. In addition, it includes the resources that may not appear in the list of servers, such as:</p> <ul style="list-style-type: none"> ■ Virtual machines that do not have the Agent for Windows installed. ■ Non-Windows virtual machines, such as Linux or Mac. ■ Virtual machines that do not have a GRT-enabled backup.

What's new in Backup Exec 2014 agents and options

This release of Backup Exec includes the following new features and capabilities in the agents and options:

Table 1-2 What's new in Backup Exec 2014 agents and options

Item	Description
Deduplication Option	<p>The following new features are available:</p> <ul style="list-style-type: none"> ■ Improved data integrity and recoverability ■ Support for 64-TB storage See “Requirements for the Deduplication Option” on page 798. ■ Increased queue processing speed ■ Improved cache loading performance ■ Improved garbage collection performance ■ Improved deduplication performance, that includes the following items: <ul style="list-style-type: none"> ■ Improved rehydration ■ Compression of segments that were previously uncompressed or were compressed inefficiently. ■ Ability to bypass the Windows localhost IPC bottleneck ■ Improved client-side directory cache management ■ Enhanced stream handler for backups of EMC and NetApp file systems <p>See “About the Deduplication Option” on page 795.</p>
Agent for VMware	<p>The Agent for VMware now includes support for VMware ESX/ESXi/vSphere/vCenter 5.1 and 5.5.</p> <p>In addition, the VSS Copy backup method is now supported. A full VSS backup is performed by default. However, the option to enable the VSS Copy backup method is now available in the Virtual Machines backup options.</p> <p>See “About the Agent for VMware” on page 830.</p>

Table 1-2 What's new in Backup Exec 2014 agents and options (*continued*)

Item	Description
Agent for Hyper-V	<p>The Agent for Hyper-V lets you back up and restore Hyper-V virtual machines that are hosted on Microsoft Windows 2012 and 2012 R2.</p> <p>See “Using the Agent for Hyper-V with Windows Server 2012/2012 R2” on page 865.</p>
Agent for Microsoft SQL Server	<p>The Agent for Microsoft SQL Server includes the following new features:</p> <ul style="list-style-type: none"> ■ Differential and Incremental backup methods for use with conversion to virtual jobs. ■ Automatic backup method that lets you back up an entire SQL instance even though some databases do not support log backups. ■ Attributes that provide the status of a SQL database in the Backup Selections dialog. ■ Full backups are automatically run on SQL databases if you run an incremental backup and one of the following occurs: <ul style="list-style-type: none"> ■ A new SQL database is added to the backup selections. ■ A full backup was not previously run on the database. ■ A product other than Backup Exec ran the latest full backup of the database. ■ A full copy was the only backup method used to back up the database.

Table 1-2 What's new in Backup Exec 2014 agents and options (*continued*)

Item	Description
Agent for Microsoft Exchange Server	<p>The Agent for Microsoft Exchange Server now supports the following:</p> <ul style="list-style-type: none"> ■ Granular Recovery Technology (GRT) of Microsoft Exchange Server 2013 ■ Microsoft Exchange Server 2013 Cumulative Update 3 ■ Microsoft Exchange Server 2013 Service Pack 1 <p>The Agent for Microsoft Exchange Server includes the new feature to grant minimal permissions for a user account to perform database backup and restores and Granular Recovery Technology operations on Exchange 2010/2013 servers.</p>
Agent for Microsoft SharePoint	<p>The Agent for Microsoft SharePoint now supports the following:</p> <ul style="list-style-type: none"> ■ Granular Recovery Technology (GRT) of Microsoft SharePoint 2013 ■ Claims-based authentication for Microsoft SharePoint 2010 and Microsoft SharePoint 2013 <p>The Agent for Microsoft SharePoint includes the new log backup method that backs up the data that is contained in the transaction log.</p>
Agent for Enterprise Vault	<p>The Enterprise Vault Agent now supports Enterprise Vault versions 10.0.3, 10.0.4 and 11.</p>
Agent for Linux	<p>The Agent for Linux now supports the following:</p> <ul style="list-style-type: none"> ■ RHEL 5.9, 6.3, and 6.4. AOFO is supported on these RHEL versions including RHEL 5.8 and 6.2 ■ SLES 11 SP2 with AOFO and OES 11 SP1 ■ Ubuntu 12.04 and 12.10 <p>See “About the Agent for Linux” on page 1188.</p>

Table 1-2 What's new in Backup Exec 2014 agents and options (*continued*)

Item	Description
Agent for Lotus Domino	<p>The Agent for Lotus Domino now supports the following:</p> <ul style="list-style-type: none"> ■ IBM Notes and Domino 9.0 Social Edition ■ The Copy on Restore option <p>Using this option, you can restore a Domino database with a different name. Thus, eliminating the need of deleting the current database. Backup Exec assigns the copied database a new database ID and a replica ID to avoid conflicts with the database that is online. Users do not experience any disruption because production database is always available. You can use this copied database to extract individual items such as emails.</p>

Installation

This chapter includes the following topics:

- [About the Backup Exec installation process and licensing options](#)
- [Backup Exec preinstallation checklist](#)
- [Running the Environment Check before installing or upgrading Backup Exec](#)
- [Microsoft SQL Server 2008 Express Edition components installed with Backup Exec](#)
- [Backup Exec's standard features](#)
- [System requirements for Backup Exec](#)
- [Installing Backup Exec by using the Installation Wizard](#)
- [Installing additional Backup Exec options to the local Backup Exec server](#)
- [Push-installing Backup Exec to remote computers](#)
- [Methods for installing the Agent for Windows](#)
- [Installing the Remote Administrator](#)
- [Installing Backup Exec using the command line \(silent mode\)](#)
- [Installing a trial version of Backup Exec agents and options](#)
- [About the installation log](#)
- [Viewing the Installation Summary Report](#)
- [Repairing Backup Exec](#)
- [Updating Backup Exec with LiveUpdate](#)
- [Viewing installed updates](#)

- [Viewing license information](#)
- [Backup Exec maintenance contract information](#)
- [About upgrading from previous versions of Backup Exec to Backup Exec 2014](#)
- [Post-installation tasks](#)
- [Uninstalling Backup Exec](#)
- [Uninstalling Backup Exec using the command line](#)
- [Uninstalling Backup Exec options from the local Backup Exec server](#)

About the Backup Exec installation process and licensing options

The Backup Exec installation wizard guides you through the process of installing Backup Exec and its agents and options. Using the installation wizard, you can install Backup Exec and its agents and options on a local computer or you can push-install them to a remote computer. The computer on which Backup Exec is installed is called the Backup Exec server. Additionally, you can install the Remote Administrator, which lets you administrate the Backup Exec server from a remote Windows server or workstation.

Note: You cannot install Backup Exec or the Agent for Windows on a volume that has been enabled for deduplication in Windows or on an ReFS volume.

When you install Backup Exec, you have the following licensing options:

Table 2-1 Licensing options

Item	Description
Enter serial numbers manually	<p>You can type the serial numbers that are listed on your license certificate. Serial numbers contain one letter and 10 numbers, such as M0123456789. After you add your serial numbers, Backup Exec polls the Symantec Web service to verify the serial numbers. An Internet connection is required to enable Backup Exec to verify the serial numbers.</p> <p>Note: Serial numbers from previous versions of Backup Exec cannot activate the most recent version of Backup Exec.</p> <p>The following are the two types of Backup Exec serial numbers:</p> <ul style="list-style-type: none">■ Product activation serial numbers activate Backup Exec products and features.■ Maintenance contract serial numbers activate the maintenance for your Backup Exec products. Maintenance contract serial numbers cannot activate Backup Exec products and features. <p>See "Backup Exec maintenance contract information" on page 124.</p>

Table 2-1 Licensing options (continued)

Item	Description
Import licenses from the Symantec License File	<p>You can import your Symantec License Files (.slf) from a network share or from a local drive.</p> <p>You may receive Symantec License Files in an email with your license certificate or you may need to go to the Symantec Licensing Portal website to obtain them. From the Symantec Licensing Portal website, you receive one file with all of the serial numbers that you registered. From the email, you receive multiple files with one serial number in each file.</p> <p>After installation, the .slf files can be found in the following locations:</p> <ul style="list-style-type: none"> ■ In Windows 2003: C:\Program Files\Common Files\Symantec Shared\Licenses ■ In Windows 2008 and later: C:\Programdata\Symantec Shared\Licenses
Install a 60-day trial version	<p>A 60-day trial version is available for Backup Exec and many of its agents and options. If you do not enter any serial numbers or Symantec license files during the installation process, a trial version is installed. If you enter a serial number or license file for Backup Exec, trial versions of many of the agents and options are available.</p>

After Backup Exec is installed, you can install additional agents and options from the Backup Exec user interface. You can install additional agents and options on the Backup Exec server or push-install Backup Exec or the Agent for Windows to remote computers.

Installation from a command line is also available. Command line installation is called silent mode installation. The silent mode installation uses the Setup.exe program on the Backup Exec installation media.

Backup Exec may install the following additional products during the installation process:

- Symantec LiveUpdate

- Microsoft XML Core Services (MSXML) 6.0
- Microsoft Report Viewer 2010 SP1
- Microsoft.NET 4
- Microsoft Visual C++ 2005 Service Pack 1 Redistributable Package MFC Security Update
- Microsoft Visual C++ 2008 Service Pack 1 Redistributable Package MFC Security Update
- Microsoft Visual C++ 2010 Service Pack 1 Redistributable Package MFC Security Update
- Microsoft Visual C++ 2012 Redistributable Package
- Microsoft SQL Server 2008 R2 Express with SP2

See [“Installing Backup Exec by using the Installation Wizard”](#) on page 66.

See [“Installing additional Backup Exec options to the local Backup Exec server”](#) on page 76.

See [“Push-installing Backup Exec to remote computers”](#) on page 78.

See [“Push-installing the Agent for Windows to remote computers”](#) on page 86.

See [“Installing Backup Exec using the command line \(silent mode\)”](#) on page 108.

See [“Backup Exec preinstallation checklist”](#) on page 57.

Backup Exec preinstallation checklist

Before you install Backup Exec, you should do the following:

- Run the Backup Exec Environment Check on the computer on which you want to install Backup Exec. The Environment Check analyzes the computer to make sure that the installation process can complete. If Backup Exec finds any configuration issues that can be fixed during the installation, or that may prevent the installation, warnings appear. Although the Environment Check runs automatically during installation, you may want to run it manually before you install Backup Exec or before you back up data with Backup Exec.

See [“Running the Environment Check before installing or upgrading Backup Exec”](#) on page 58.

- Check the Backup Exec Hardware Compatibility List to confirm that your storage device hardware is supported in this version of Backup Exec.

You can find a list of compatible types of storage in the Backup Exec hardware compatibility list:

<http://entsupport.symantec.com/umi/v-269-2>

- Install the storage device hardware (controller, drives, robotic libraries) on the Backup Exec server. Refer to the documentation that is included with your storage device hardware for installation instructions. Use the appropriate Windows hardware setup functions to configure your controller and storage devices. Refer to your Microsoft Windows documentation for more information. You can find a list of compatible types of storage in the Backup Exec hardware compatibility list:
<http://entsupport.symantec.com/umi/v-269-2>
- Check your Windows security settings to make sure that they work properly with the Backup Exec service account.
See [“Changing the credentials for a service account”](#) on page 621.
- If the drive on which you want to install Backup Exec is encrypted or compressed, and you want to use a default SQL Express database, verify that an unencrypted and uncompressed drive is available for SQL Express installation
- Check the computer name of the computer on which you want to install Backup Exec. It should only use standard ANSI characters. You may receive errors if you install Backup Exec on a computer with a name that uses non-standard characters.
- If you want to install Backup Exec to a non-English version of Windows, download the Microsoft SQL Server 2008 R2 Express with SP2 setup file for the language that you want to install from the Microsoft website before you install Backup Exec if all of the following are true:
 - You want to use a local Backup Exec SQL Express instance.
 - You have non-English SQL Server instances on the computer on which you want to install Backup Exec.

Running the Environment Check before installing or upgrading Backup Exec

The Symantec Backup Exec Environment Check is a utility that runs on a computer automatically during installation and that reports the following:

- If the computer meets the minimum requirements for installation, such as the operating system, disk and physical memory, and sufficient logon account privileges.
See [“System requirements for Backup Exec”](#) on page 63.
- If the third-party software that uses Backup Exec ports is configured correctly.
- If required components are installed, and if they are compatible with Backup Exec.

- If previous versions of Backup Exec and Backup Exec options are installed.
- If storage device hardware and associated drivers are properly installed and recognized by the Windows operating system.

One of the following results is reported for each item:

Table 2-2 Environment Check results

Result	Description
Passed	There are no incompatibilities to prevent the Backup Exec installation. For hardware, this result indicates that the hardware configuration is recognized by Backup Exec.
Warning	An incompatibility with Backup Exec exists. Some of the issues may be resolved during the Backup Exec installation. A warning does not prevent Backup Exec from installing. However, if the issues are not resolved during installation, jobs may fail.
Failed	An incompatibility with Backup Exec exists, and it will cause the installation to fail. Action is required before you can successfully install Backup Exec.

Although the Environment Check runs automatically during installation, you may want to run it manually before installing Backup Exec or before backing up data with Backup Exec.

To check your environment before installing

- 1 From the installation media browser (Browser.exe), click **Preinstallation**, and then click **Backup Exec**.
- 2 Click **Next**.
- 3 Do any of the following:

To check the configuration of the local computer

To check the configuration of a remote computer

Check **Local Environment Check**.

Check **Remote Environment Check**.
- 4 Click **Next**.
- 5 If you checked **Remote Environment Check** in step 3, do one of the following, and then click **Next**:

To select the name of a computer from a list Do the following:

- Click **Add Server From List**.
- Select the computer from the list, and then click **Next**.

To add the name of a computer manually Do the following:

- Click **Add Server Manually**.
- In the **Domain** field, type the name of the domain.
- In the **Computer Name** field, type the name of the computer.
- Click **OK**.
- Type the user name and password for this computer.
- Click **OK**.

To remove the name of a computer from the list of computers on which the Environment Check runs Do the following:

- Select the computer from the list.
- Click **Remove**.

- 6 Review the results of the Environment Check, and then to save the results, check **Save Results To**.

To change the location where the Environment Check results are saved, click **Change Path** to browse to a new location.

- 7 Click **Finish**.

Microsoft SQL Server 2008 Express Edition components installed with Backup Exec

The Backup Exec installation program installs Microsoft SQL Server 2008 R2 Express with SP2 components that are required to run Backup Exec.

Backup Exec prompts you to do one of the following:

- Install the required Microsoft SQL Server Express components with Backup Exec and create a default Backup Exec instance. This is the recommended action.

Note: SQL Express can be installed on a server that runs other instances of SQL Express or full versions of SQL.

- Select a Microsoft SQL Server 2005 (SP4) or SQL Server 2008/2012 instance that already exists on the network on which you want to run Backup Exec.

If you choose to install Backup Exec into an existing SQL Server 2005 or 2008 instance, make sure that one of the following is installed before you continue with the installation:

- SQL Server 2005 Service Pack 4 or later
- SQL Server 2008 Service Pack 3 or later

Caution: Backup Exec may not function properly if you install it into an existing SQL Server instance that uses case-sensitive collation. Symantec recommends that you avoid installing Backup Exec to a SQL Server instance that uses case-sensitive collation.

When Backup Exec is installed into an existing instance, the automated master database restore feature is not available. To recover the Master database, you must replace it with the Master database copy that Backup Exec automatically creates and updates when the Master database is backed up.

When Backup Exec applies updates, the SQL instance is stopped, which may cause other databases in the same instance to lose connectivity. If a remote SQL instance is used, ensure that Backup Exec has good network connectivity with the instance to avoid errors. A default local instance of SQL Express is recommended.

You cannot install multiple Backup Exec Databases on the same SQL Server instance.

Note: If you are installing a managed Backup Exec server, it is recommended that you select a local Microsoft SQL Server 2005 (SP4) instance or later on which to install the Backup Exec Database for this managed server. Do not select the same SQL Server instance that is used by the central administration server.

See [“System requirements for Backup Exec”](#) on page 63.

Backup Exec's standard features

The following features are included with Backup Exec at no additional cost. When you install Backup Exec, you can select any of these features that you want to use.

Table 2-3 Backup Exec's standard features

Feature	Description
Tape Device Drivers	<p>Installs the Symantec tape device drivers for all supported tape devices that are attached to the server. If there are no tape devices attached to your Backup Exec server, uncheck this option.</p> <p>Note: You do not need to install Symantec tape device drivers if Backup Exec runs on Windows Server 2012 and later. Symantec kernel-mode drivers and tapeinst.exe are no longer installed if Backup Exec runs on Windows Server 2012 or later.</p> <p>See "Installing Symantec tape device drivers" on page 360.</p>
Online Documentation	Installs the Backup Exec Administrator's Guide in a pdf file format.
Enable Robotic Library Support	Enables support for tape libraries, or optical robotic libraries and library storage systems. Backup Exec includes support for one drive in every robotic library. Each additional drive in a library requires a Library Expansion Option license.
Copy Server Configurations	Enables you to copy configuration settings and logon information between Backup Exec servers. This option is recommended for environments that contain multiple Backup Exec servers. This option is required for the Central Admin Server Option (CASO).
Managed Backup Exec server	Installs the managed Backup Exec server component of the Central Admin Server Option (CASO). You can install managed Backup Exec servers after you install a central administration server.
Virtual Tape Library Support	<p>Provides support for every single-drive Virtual Tape Library (VTL). You must purchase the Virtual Tape Library Unlimited Drive Option to support additional drives in each VTL.</p> <p>If you select this option, the Enable Robotic Library Support option is selected automatically. You cannot uncheck Enable Robotic Library Support unless you uncheck Virtual Tape Library Support.</p>

All other options and agents require the purchase of additional licenses. Installing a trial version enables many options that must be purchased separately and are not included as part of Backup Exec.

If you have a licensed version of Backup Exec, you can use a trial version of most options and agents for a specified period of time.

See “[Installing a trial version of Backup Exec agents and options](#)” on page 118.

System requirements for Backup Exec

The following are the minimum system requirements to run this version of Backup Exec:

Table 2-4 Minimum system requirements

Item	Requirements
Operating system	<p>You can find a list of compatible operating systems, platforms, and applications at the following URL:</p> <p>http://entsupport.symantec.com/umi/V-269-1</p> <p>You cannot install a Backup Exec server on a computer that runs the Windows Server Core installation option of Windows Server 2008/2012. You can only install the Backup Exec Agent for Windows on Server Core computers.</p> <p>On the computers that run Windows 2003, ensure that those computers have Windows Imaging Component installed before you install Backup Exec 2014. Windows Imaging Component is required to install Backup Exec 2014.</p> <p>See the Microsoft website for more information on downloading and installing Windows Imaging Component.</p> <p>You cannot install SQL Express or SQL Server on a Windows Server 2008 computer that is configured in a Read Only Domain Controller (RODC) role. The Read Only Domain Controller role does not let you use the local accounts that are required for SQL Express and SQL Server. When you install Backup Exec on an RODC computer you must select a remote SQL instance for the Backup Exec Database.</p> <p>For Windows Server 2012 and later computers, you cannot install Backup Exec or the Agent for Windows on a volume that is enabled for deduplication in Windows or an ReFS volume.</p>
Additional application support	You can use Backup Exec with Microsoft System Center Operation Manager (SCOM) 2007.
Internet browser	Internet Explorer 7.0 or later.
Processor	Intel Pentium, Xeon, AMD, or compatible.
Screen resolution	800 x 600

Table 2-4 Minimum system requirements (*continued*)

Item	Requirements
SQL Server or SQL Express	SQL Server 2008 R2 SP2.
Memory	<p>Required: 1GB RAM above the operating system's requirements for the exclusive use by Backup Exec.</p> <p>Recommended: 2 GB RAM (or more for better performance)</p> <p>Note: RAM requirements may vary depending on the operations performed, the options installed, and the specific computer configuration.</p> <p>For the Central Admin Server Option: 1 GB RAM is required; 2 GB RAM is recommended.</p> <p>Symantec Recovery Disk: 1 GB minimum (dedicated) for the multi-lingual version.</p> <p>Virtual Memory Recommendations: 20 MB above the Windows recommended size for total paging file size (total for all disk volumes). Refer to your Microsoft Windows documentation for instructions on how to view or set the paging file size.</p>
User interface language	<p>The Backup Exec user interface displays in the format that is configured in the Region and Language settings in the Control Panel. You can change the Backup Exec display language so that you can view user interface items in a different language.</p> <p>If Backup Exec does not support a language, the user interface displays in English. The user interface also displays in English if the menu and dialog boxes option is set to a language other than the language you want to display the user interface. If you do not use one of the supported languages listed in Table 2-5, you must install the English language pack in Windows.</p>
Installation disk space	<p>1.26 GB (Typical installation)</p> <p>1.91 GB (Includes all options)</p> <p>Note: Disk space requirements may vary depending on the operations performed, the options installed, and the specific system configuration. The Backup Exec Database and catalogs require additional space. An additional 525 MB is required for SQL Express. Any disk storage that you use also requires additional space.</p>

Table 2-4 Minimum system requirements (*continued*)

Item	Requirements
Other Hardware	<p>The following hardware is recommended:</p> <ul style="list-style-type: none"> ■ Network interface card or a virtual network adapter card. ■ CD/DVD drive. ■ A mouse.
Storage Hardware	<p>You can use storage media drives, robotic libraries, removable storage devices, and non-removable hard drives.</p> <p>You can find a list of compatible types of storage at the following URL: http://entsupport.symantec.com/umi/V-269-2</p> <p>Support is available for the first drive in each robotic library when you purchase Backup Exec. To enable support for each additional robotic library drive, you must purchase the Backup Exec Library Expansion Option.</p>

Table 2-5 Supported languages for the Backup Exec user interface

Language	Language code
Chinese (Simplified)	ZH
Chinese (Traditional)	CH
English	EN
French	FR
German	DE
Italian	IT
Japanese	JP
Korean	KO
Spanish	ES
Russian	RU
Portuguese (Brazilian)	PT

See “Installing Backup Exec by using the Installation Wizard” on page 66.

Installing Backup Exec by using the Installation Wizard

The Backup Exec installation program provides two methods of installation: typical and custom. A typical installation is a simpler installation method than a custom installation and is designed for small or uncomplicated environments. For example, if you use a local Backup Exec server and a few Backup Exec agents or options, then a typical installation may be best for you. A custom installation is designed for large or complex environments. You can also use the custom installation method if you prefer to set all of your options. For example, if you use a remote Backup Exec server or use the Enterprise Server Option, you should perform a custom installation.

With a typical installation, Backup Exec makes the following decisions for you, based on common installation scenarios:

- Backup Exec is installed to a local Backup Exec server.
- SQL Express is installed with the default instance.
- Agents and options are installed if you enter the licenses for them. If you do not enter any licenses, a trial version of Backup Exec is installed. A trial version enables many options that must be purchased separately and are not included as part of Backup Exec.

Note: The installation program prevents you from selecting a license for an agent or option that is not compatible with the typical installation method, such as the Enterprise Server Option or the Archiving Option.

- LiveUpdate runs automatically.

Note: Before you install, make sure that your licenses for the Backup Exec version and any agents or options that you want to install are available.

Licenses are required to install Backup Exec and its agents and options. However, you can install a trial version of Backup Exec without a license.

Note: If you are installing the Archiving Option, ensure that the server does not have any unformatted hard drives. The Archiving Option cannot be installed on an unformatted drive. If your server has an unformatted drive, you should either format the drive or remove it before you attempt to install the Archiving Option.

Choose either the typical installation or the custom installation:

[How to install a typical installation](#)

[How to install a custom installation](#)

How to install a typical installation

Follow these steps to install a typical installation of Backup Exec.

To install a typical installation of Backup Exec

- 1 From the installation media browser, click **Install Products**, and then select **Backup Exec**.

If the required version of Microsoft .NET Framework is not already installed on this computer, Backup Exec installs it. The installation of the Microsoft.NET Framework may take some time.
- 2 On the **Welcome** panel, read the license agreement, and then click **I accept the terms of the license agreement**.
- 3 Click **Next**.
- 4 On the **Installation Type** panel, click **Typical installation**, and then click **Next**.
The Environment Check runs automatically.
- 5 Review the results of the Environment Check.
- 6 Do one of the following:
 - If the Environment Check does not reveal any issues that may prevent a successful installation of Backup Exec, click **Next**.
 - If the Environment Check reveals any issues that may prevent a successful installation of Backup Exec, click **Cancel** to exit the wizard. Correct the issues before you attempt to install Backup Exec again.
- 7 Select one of the following methods to add licenses:

To enter serial numbers manually

Note: An Internet connection is required to validate the serial numbers. If you do not have an Internet connection, import the licenses from the Symantec License File included with your license certificate or install a trial version.

Do the following in the order listed:

- In the **Enter a Serial number** field, type the appropriate serial number from your license certificate.
- Click **Add to List**.
- Repeat for each serial number.

To import licenses from the Symantec License File

Do the following in the order listed:

- Click **Import Symantec License File**.
- Browse to the location of your license files, and then select the appropriate file.

To install a trial version

Do not type a serial number or import a license file. Go to step [8](#).

8 Click **Next**.

You may be prompted to enter contact information for maintenance contract serial numbers.

9 If you entered product activation serial numbers, then on the **Review Licenses** panel, do one of the following:

To install a licensed version of Backup Exec

Do the following in the order listed:

- In the **Select a Backup Exec edition license to install on the computer** field, select the Backup Exec edition to install.
- Check the check boxes for the agents or options you want to install.
- Click the drop-down menu, and then select the number of licenses that you want to install.

To install a trial version

In the **Select a Backup Exec edition license to install on the computer** field, select **Trial**.

10 Click **Next**.

- 11 On the **Service Account** panel, provide a user name, password, and domain for an Administrator account that the Backup Exec system services can use.

You cannot install Backup Exec with an account that has a blank password on Windows Server 2003/2008/2012 computers unless Windows is configured to allow it. If you try to do so, the following error message appears when Backup Exec services are created:

The account name and/or password supplied is not valid. Re-enter the login information and try again.

You can, however, configure Windows to allow for blank passwords. For more information, see your Windows documentation.

- 12 If you want to change the directory where Backup Exec files are installed, click **Change**, and then select a new location.

If you change the directory to a new location, ensure that you select a secure location where you can store sensitive data such as passwords.

- 13 Click **Next**.

- 14 If the **SQL Express Setup** panel appears, perform the following steps to identify the location of the SQL Express setup file:

- Click **Browse**.
- Navigate to the location where you downloaded the SQL Express 2008 R2 SP2 setup file.
- Click **OK**.
- Click **Next**.

- 15 On the **Remote Computers** panel, do one of the following:

To install the Agent for Windows on one remote computer

Do the following in the order listed:

- Click **Add**.
- Select **Add a Single Computer**.
- Type the fully qualified name of the remote computer or click **Browse Remote Computers** to locate the remote computer.
- Under **Remote computer credentials**, type the credentials that Backup Exec can use to connect to the remote servers.
You must use Administrator credentials.
- Click **Next**.
- In the **Destination Folder** field, enter the path where you want to install the files.
- Click **Next**.
- After all of the computers in the list are validated and the list is complete, click **Next**.

To install the Agent for Windows on multiple computers using the same settings

Do the following in the order listed:

- Click **Add**.
- Select **Add Multiple Computers with the Same Settings**.
- Type the fully qualified name of the remote computer or click **Browse** to locate the remote computer.
- Click **Add to List**.
Type the fully qualified name and then click **Add to List** for all of the remote computers for which you want to push-install the Agent for Windows.
- Under **Remote computer credentials**, type the credentials that Backup Exec can use to connect to the remote servers.
You must use Administrator credentials.
- Click **Next**.
- In the **Destination Folder** field, enter the path where you want to install the files.
- Click **Next**.
- After all of the computers in the list are validated and the list is complete, click **Next**.

To proceed without push-installing the Agent for Windows

Click **Next**.

- 16 Review the Backup Exec installation summary, and then click **Install**.
- 17 If you installed the Agent for Windows on remote computers, on the **Remote Installation** dialog box, click **Next**.
- 18 When the installation is complete, you can run LiveUpdate, choose to restart the system, view the readme, or remove the Backup Exec shortcut from the desktop.
- 19 Click **Next**, and then click **Finish** to exit the wizard.

If you chose to restart the system, the computer will restart automatically.

The installation process creates an installation log named Bkupinst2014.htm in one of the following directories on the computer where Backup Exec is installed.

- For Windows 2003: %allusersprofile%\Application Data\Symantec\Backup Exec\Logs

- For Windows 2008 and later: %programdata%\Symantec\Backup Exec\Logs
See [“About the installation log”](#) on page 119.

How to install a custom installation

Follow these steps to install a custom installation of Backup Exec.

To install a custom installation of Backup Exec

- 1 From the installation media browser, click **Install Products**, and then select **Backup Exec**.

If the required version of Microsoft .NET Framework is not already installed on this computer, Backup Exec installs it. The installation of the Microsoft.NET Framework may take some time.
- 2 On the **Welcome** panel, read the license agreement, and then click **I accept the terms of the license agreement**.
- 3 Click **Next**.
- 4 On the **Installation Type** panel, click **Custom installation**, and then click **Next**.
- 5 On the **Menu** panel, check **Local Installation**, and then select **Install Backup Exec software and options**.
- 6 Click **Next**.

The Backup Exec Environment Check runs automatically after you click **Next**.
- 7 Review the results of the Environment Check.
- 8 Do one of the following:
 - If the Environment Check does not reveal any issues that may prevent a successful installation of Backup Exec, click **Next**.
 - If the Environment Check reveals any issues that may prevent a successful installation of Backup Exec, click **Cancel** to exit the wizard. Correct the issues before you attempt to install Backup Exec again.
- 9 Select one of the following methods to enter licenses:

To enter serial numbers manually	<p>Do the following in the order listed:</p> <ul style="list-style-type: none"> ■ In the Enter a Serial number field, type the appropriate serial number from your license certificate. ■ Click Add to List. ■ Repeat for each license for each option or agent that you want to install. ■ Click Next to validate the serial numbers.
<p>Note: An Internet connection is required to validate the serial numbers. If you do not have an Internet connection, import the licenses from the Symantec License File or install a trial version.</p>	
To import licenses from the Symantec License file	<p>Do the following in the order listed:</p> <ul style="list-style-type: none"> ■ Click Import Symantec License File. ■ Browse to the location of your license files, and then select the appropriate file.
To install a trial version	<p>Do not type a serial number or import a license file. Go to step 10.</p>

10 Click **Next**.

You may be prompted to enter contact information for maintenance contract serial numbers.

11 If you entered product activation serial numbers, on the **Review Licenses** panel, do one of the following:

To install a licensed version of Backup Exec	<p>Do the following in the order listed:</p> <ul style="list-style-type: none"> ■ In the Select a Backup Exec edition license to install on the computer field, select the Backup Exec edition to install. ■ Check the check boxes for the agents or options you want to install. ■ Click the drop-down menu, and then select the number of licenses that you want to install.
To install a trial version	<p>In the Select a Backup Exec edition license to install on the computer field, select Trial.</p>

12 Click **Next**.

- 13 On the **Configure Options** panel, select any additional options that you want to install.

For example, you can select additional standard features, or you can select the agents or the options that are available for a trial installation.

- 14 Click **Next**.

If you selected the File System Archiving Option or the Microsoft Exchange Mailbox Archiving Option, the Archiving Option Environment Check runs in the background. The Archiving Option Environment Check verifies that the computer meets the minimum requirements for installing and configuring the options. If the computer does not meet the minimum requirements, you must uncheck the archiving options or fix the errors before you can continue with the installation.

- 15 If you want to install Backup Exec for any additional languages, select the language, and then click **Next**.

- 16 On the **Destination** panel, do the following:

- Review the disk space requirements for the items that you selected to install.
- If you want to change the directory where the Backup Exec files are installed, click **Change**, and then select a new directory or create a new folder.
Symantec recommends that you do not select a mount point as the destination directory because if you delete the mount point, Backup Exec is uninstalled.

- 17 Click **Next**.

- 18 Provide a user name, password, and domain for an Administrator account that the Backup Exec system services can use, and then click **Next**.

You cannot install Backup Exec with an account that has a blank password on Windows Server 2003/2008/2012 computers unless Windows is configured to allow it. If you try to do so, the following error message appears when Backup Exec services are created:

The account name and/or password supplied is not valid. Re-enter the login information and try again.

You can, however, configure Windows to allow for blank passwords. For more information, see your Windows documentation.

- 19 On the **Choose SQL Server** panel, do one of the following to select a location to store the Backup Exec Database.

Note: The **Choose SQL Server** panel does not appear for upgrades. You cannot change the database location during the upgrade process. If you want to change the database location after the upgrade, use BE Utility.

To create a local Backup Exec SQL Express instance

Do the following in the order listed:

- Click **Create a local Backup Exec SQL Express instance to store the Backup Exec database.**
- To change the location of the Backup Exec SQL Express instance, click **Browse**.
- Select the location, and then click **OK**.

To use an existing SQL Server 2005/2008/2012 instance

Do the following in the order listed:

- Click **Use an existing instance of SQL Server 2005 SP4 or later, SQL Server 2008 SP3 or later, or SQL Server 2012.**
- Select the instance.

When Backup Exec is installed into an existing instance, the automated Master database restore feature is not available. To recover the Master database, replace it with the Master database copy that Backup Exec automatically creates and updates when the Master database is backed up.

See [“Microsoft SQL Server 2008 Express Edition components installed with Backup Exec”](#) on page 60.

20 Click Next.

Backup Exec attempts to connect to the instance.

21 If the SQL Express Setup panel appears, perform the following steps to identify the location of the SQL Express setup file:

- Click **Browse**.
- Navigate to the location where you downloaded the SQL Express 2008 R2 SP2 setup file.
- Click **OK**.
- Click **Next**.

- 22 If you are prompted, select how the **Symantec Device Driver Installer** should install device drivers for the tape storage devices that are connected to the server, and then click **Next**.

Symantec recommends that you select **Use Symantec device drivers for all tape devices**.

- 23 If you are prompted, enter information or choose settings for the additional options that you want to install, and then click **Next** after each selection.
- 24 Review the Backup Exec installation summary, and then click **Install**.

The installation process takes several minutes to complete. During the process, the progress bar may not move for several minutes.

- 25 When the installation is complete, you can run LiveUpdate, choose to restart the system, view the readme, or remove the Backup Exec shortcut from the desktop.

- 26 Click **Next**, and then click **Finish** to exit the wizard.

If you chose to restart the system, the computer will restart automatically.

Installing additional Backup Exec options to the local Backup Exec server

You can install agents and options when you install Backup Exec. However, if you have already installed Backup Exec and want to install additional options, review the documentation for those options to ensure that your system meets all minimum requirements. The Backup Exec services may be stopped while the additional options are installed. If any active jobs are running, you are prompted to stop them, or to wait for the jobs to finish.

If you have a licensed version of Backup Exec, you can use a trial version of most options and agents for a specified period of time.

See [“Installing a trial version of Backup Exec agents and options”](#) on page 118.

Note: If the Central Admin Server Option is installed, and you want to install additional options on a managed Backup Exec server, you can pause the managed Backup Exec server. When a managed Backup Exec server is paused, the administration server does not delegate jobs to it. When the installation is complete, un-pause, or resume, the managed Backup Exec server.

See [“Pausing or resuming a managed Backup Exec server”](#) on page 1144.

To install additional Backup Exec options to the local Backup Exec server

- 1 Click the Backup Exec button, select **Installation and Licensing**, and then select **Install Options and Licenses on this Backup Exec Server**.

You may be prompted to insert the installation media.

- 2 Do one of the following:

To enter serial numbers manually

Note: An Internet connection is required to validate the serial numbers. If you do not have an Internet connection, import the licenses from the Symantec License File or install a trial version.

Do the following in the order listed:

- In the **Enter a Serial number** field, type the appropriate serial number from your license certificate.
- Click **Add to List**.
- Repeat for each serial number.

To import licenses from the Symantec license file

Do the following in the order listed:

- Click **Import Symantec License File**.
- Browse to the location of your license files, and then select the appropriate file.

To install a trial version

Do not type a serial number or import a license file. Go to step [5](#).

- 3 Click **Next**.

You may be prompted to enter contact information for maintenance contract serial numbers.

- 4 If you entered product activation serial numbers, on the **Review Licenses** panel, do one of the following:

To install a licensed version of Backup Exec

Do the following in the order listed:

- In the **Select a Backup Exec edition license to install on the computer** field, select the Backup Exec edition to install.
- Check the check boxes for the agents or options you want to install.
- Click the drop-down menu, and then select the number of licenses that you want to install.

To install a trial version

In the **Select a Backup Exec edition license to install on the computer** field, select **Trial**.

- 5 Click **Next**.
- 6 Do the following:

■ Check the check boxes for the additional options that you want to install.

■ Uncheck the check boxes for the options that you want to remove.
- 7 Click **Next**.
- 8 If you are prompted, enter information or choose settings for the additional options that you want to install. Click **Next** after each selection.
- 9 Review the Backup Exec installation summary, and then click **Install**.

The Backup Exec services are stopped while the additional options are installed. If any active jobs are running, you are prompted to stop them, or to wait for the jobs to finish.

When the installation is complete, the services are restarted.
- 10 Click **Finish**.

Push-installing Backup Exec to remote computers

If you install Backup Exec through Terminal Services and the installation media is on a shared drive (network share) you must use a UNC path. Installation by mapped drives is not supported.

You can set up multiple server installations. Backup Exec processes up to five remote computer installations concurrently.

Before you install Backup Exec to remote computers, you should review the special considerations.

Table 2-6 Special considerations for installing Backup Exec to remote computers

Item	Consideration
32-bit and 64-bit computers	If you try to push-install Backup Exec from a 32-bit computer to a 64-bit computer, you may be prompted to insert the 64-bit installation media.

Table 2-6 Special considerations for installing Backup Exec to remote computers *(continued)*

Item	Consideration
Windows Server 2003 SP2	<p>To push-install Backup Exec to a Windows Server 2003 computer, you must enable File and Printer Sharing on the Windows Firewall Exceptions list for the following ports:</p> <ul style="list-style-type: none"> ■ 135 (RPC) ■ 445 (TCP) ■ 103X (mostly 1037) ■ 441 (RPC) <p>For more information about the Windows Firewall Exceptions list, refer to your Microsoft Windows documentation.</p> <p>During the installation process, Backup Exec sets the Remote Launch and remote access security permissions for the Administrator's group.</p> <p>You should enable the "Allow remote administration exception" group policy for the computer to which you push the installation.</p>
Windows Server 2008 or later	<p>To push-install Backup Exec to a computer that runs Windows Server 2008 or later, you must enable the following items on the destination computer's Windows Firewall Exceptions list</p> <ul style="list-style-type: none"> ■ File and Printer Sharing ■ Windows Management Instrumentation (WMI) <p>For more information, refer to your Microsoft Windows documentation.</p> <p>You cannot install Backup Exec on a volume that has been enabled for deduplication in Windows or on an ReFS volume.</p>
Symantec Endpoint Protection (SEP) 11.0 or later	<p>To push-install Backup Exec to a computer that runs Symantec Endpoint Protection (SEP) version 11.0 or later, you must configure SEP to share files and printers. The file and printer sharing feature is turned off by default.</p>

Note: You can also use Microsoft's Add or Remove Programs utility to install Backup Exec to a remote computer. See your Microsoft documentation for more information.

The installation process creates an installation log named Bkupinst2014.htm in one of the following directories on the computer where Backup Exec is installed.

- For Windows 2003: %allusersprofile%\Application Data\Symantec\Backup Exec\Logs
- For Windows 2008 and later: %programdata%\Symantec\Backup Exec\Logs

To push-install Backup Exec to remote computers

1 Do one of the following:

To push-install Backup Exec to remote computers from the installation media

Do the following steps in the order listed:

- From the installation media browser, click **Install Products**, and then click **Backup Exec**.
- On the **Welcome** panel, click **Next**.
- Select **I accept the terms of the license agreement**, and then click **Next**.
- Select **Custom installation**.
- Uncheck **Local Installation**, and then check **Remote Installation**.
- Click **Next**.

To push-install Backup Exec to remote computers from the Backup Exec server

Click the Backup Exec button, select **Installation and Licensing**, and then select **Install Agents and Backup Exec Servers on Other Servers**.

2 On the **Remote Computers** panel, do one of the following:

To install Backup Exec on one remote computer

Do the following in the order listed:

- Click **Add**.
- Select **Add a Single Computer**.
- Select **Symantec Backup Exec**, and then click **Next**.
- Type the fully qualified name of the remote computer or click **Browse Remote Computers** to locate the remote computer.

To install Backup Exec on multiple computers using the same settings

Do the following in the order listed:

- Click **Add**.
- Select **Add Multiple Computers with the Same Settings**.
- Select **Symantec Backup Exec**, and then click **Next**.
- Type the fully qualified name of the remote computer or click **Browse** to locate the remote computer.
- Click **Add to List**.
Type the fully qualified name and then click **Add to List** for all of the remote computers for which you want to push install the options.

- 3 Under **Remote computer credentials**, type the credentials that Backup Exec can use to connect to the remote servers.

You must use Administrator credentials.

- 4 Click **Next**.

- 5 Select one of the following methods to enter licenses:

To enter serial numbers from your license certificate

Note: An Internet connection is required to validate the serial numbers. If you do not have an Internet connection, import the licenses from the Symantec License File or install a trial version.

Do the following in the order listed:

- In the **Enter a Serial number** field, type the appropriate serial number from your license certificate.
- Click **Add to List**.
- Repeat for each license for each option or agent that you want to install.
- Click **Next** to validate the serial numbers.

To import licenses from a Symantec License File

Do the following in the order listed:

- Click **Import Symantec License File**.
- Browse to the location of your license files, and then select the appropriate file.

To install a trial version Do not type a serial number or import a license file. Go to step 8.

6 Click Next.

You may be prompted to enter contact information for maintenance contract serial numbers.

7 If you entered product activation serial numbers, on the Review Licenses panel, do one of the following:

To install a licensed version of Backup Exec

Do the following in the order listed:

- In the **Select a Backup Exec edition license to install on the computer** field, select the Backup Exec edition to install.
- Check the check boxes for the agents or options you want to install.
- Click the drop-down menu, and then select the number of licenses that you want to install.

To install a trial version

In the **Select a Backup Exec edition license to install on the computer** field, select **Trial**.

8 Click Next.

9 On the Configure Options panel, select any additional options that you want to install.

For example, you can select additional standard features, or you can select the agents or the options that are available for a trial installation.

10 In the Destination Folder field, enter the location where you want to install Backup Exec.

11 Click Next.

12 Complete the service account credentials options as follows:

User Name	<p>Type the user name for an Administrator account that the Backup Exec services can use.</p> <p>If the remote computer is in a domain, use a domain administrator's account or an equivalent account that is part of the domain administrator's group.</p> <p>If the remote computer is in a workgroup, use an administrator's account or an equivalent account that is part of the administrator's group on the computer.</p>
Password	<p>Type the password for an administrator account that the Backup Exec services can use.</p>
Domain	<p>If the computer is in a domain, select the domain in which the computer is located.</p> <p>If the computer is in a workgroup, select the computer name.</p>

13 Click **Next**.

14 Do one of the following to select a location on which to store the Backup Exec Database, and then click **Next**.

To create a local Backup Exec SQL Express instance	<p>Do the following in the order listed:</p> <ul style="list-style-type: none">■ Click Create a local Backup Exec SQL Express instance to store the Backup Exec database.■ To change the location of the database, type the new location in the Destination Folder field.
--	--

To use an existing SQL Server 2005 or SQL Server 2008 instance

Do the following in the order listed:

- Click **Use an existing instance of SQL Server 2005 SP4 or later or SQL Server 2008/2012.**
- Select the instance.

When Backup Exec is installed into an existing instance, the automated master database restore feature is not available. To recover the Master database, you must replace it with the Master database copy that Backup Exec automatically creates and updates when the Master database is backed up.

See [“Microsoft SQL Server 2008 Express Edition components installed with Backup Exec”](#) on page 60.

Backup Exec attempts to connect to the instance.

This step is skipped during upgrades.

- 15 Click **Next**.
- 16 Make a selection for tape device drivers, and then click **Next**.

Note: You do not need to install Symantec tape device drivers if Backup Exec runs on Windows Server 2012 and later. Symantec kernel-mode drivers and tapeinst.exe are no longer installed if Backup Exec runs on Windows Server 2012 and later.

See [“Installing Symantec tape device drivers ”](#) on page 360.

- 17 Click **Next**.
- 18 If you are prompted, enter information or choose settings for additional options that are being installed, and then click **Next** or **OK** after each selection.
- 19 After Backup Exec validates the remote computers, you can change the list in any of the following ways:

To manually add one remote computer

Click **Add**, and then click **Add a Single Computer**.

To manually add multiple remote computers

Click **Add**, and then click **Add Multiple Computers with the Same Settings**.

To add multiple remote computers by importing an existing list of computers

Click **Import and Export**, and then select one of the following options:

- Select **Import from File** to enable Backup Exec to add the names of the remote computers from a selected list.
- Select **Import Servers Published to this Backup Exec server** to enable Backup Exec to add the names of all the remote computers that are set up to publish to this Backup Exec server.

You must enter remote computer logon credentials for the list of remote computers.

To change the product that you selected to install or to change other properties you selected for this installation

Select the remote computer that you want to change, and then click **Edit**.

To delete a remote computer from the list

Select the remote computer that you want to delete, and then click **Delete**.

To save this list of remote computers and the associated remote computer logon credentials

Verify that **Save the server list for future remote install sessions** is checked.

This option enables the names and the credentials of all of the remote computers to be added automatically the next time you install Backup Exec or options to these remote computers.

To save the list of remote computers to an XML file

Click **Import and Export**, and then click **Export to File**.

You can select the location to save the Push_Export.xml file. This option is useful if you want to use the same list for multiple Backup Exec servers. When you import the list, you must re-enter the remote computer logon credentials.

To fix the errors that were located during the validation

Right-click the name of the computer, and then click **Fix Error**.

To enable Backup Exec to attempt to re-validate an invalid remote computer

Right-click the name of the computer, and then click **Retry Validation**.

20 After all of the computers in the list are validated and the list is complete, click **Next**.

- 21 Review the Backup Exec installation summary, and then click **Install**.

See [“About the installation log”](#) on page 119.

- 22 Click **Next**, and then click **Finish** to exit the wizard.

If you did not restart the remote computer, you may need to do it now in order for the configuration to take effect.

Methods for installing the Agent for Windows

You can install the Agent for Windows by using the following methods, depending on your environment:

- Push-install the Agent for Windows to one or more remote computers from the Backup Exec server.
See [“Push-installing the Agent for Windows to remote computers”](#) on page 86.
- Add the remote computer to the list of servers and install the Agent for Windows on the remote computer.
See [“Adding servers that you want to back up to the list of servers on the Backup and Restore tab”](#) on page 153.
- Use a Microsoft Active Directory network to centrally manage the installation of the Agent for Windows to computers in the network.
See [“Installing the Agent for Windows in an Active Directory network”](#) on page 94.
- Use a command prompt.
See [“Using a command prompt to install the Agent for Windows on a remote computer”](#) on page 98.
- Use command script files.
See [“Using a command script to install the Agent for Windows”](#) on page 101.

Push-installing the Agent for Windows to remote computers

You can push-install the Agent for Windows to remote computers from a Backup Exec server. Push installations save time by eliminating the need for local access at the target computer for the installation to be successful. You can push-install the Agent for Windows to an unlimited number of remote computers. Backup Exec can process up to five active push-installations at a time.

Review the following special considerations before you install the Agent for Windows on remote computers.

Table 2-7 Special considerations for installing the Agent for Windows

Item	Consideration
ForceGuest configuration	<p>You cannot push-install the Agent for Windows when the remote computer is in the ForceGuest configuration and it is not in a domain. ForceGuest is an operating system configuration that limits incoming users to Guest-level access. Instead, use the installation media or the network to install the Agent for Windows on the Windows computer. You can also turn off ForceGuest. Refer to your Microsoft Windows documentation for more information.</p> <p>See “Installing Backup Exec using the command line (silent mode)” on page 108.</p>
Server core option of Windows Server 2008/2012	<p>Backup Exec installs a command-line version of the Agent for Windows on the computers that run the Server Core installation option of Windows Server 2008/2012. The Backup Exec Agent Utility command-line applet is installed with the Agent for Windows. This applet lets you monitor Backup Exec operations on the remote computer.</p> <p>See “Backup Exec Agent Utility Command Line Applet switches” on page 790.</p>
Windows deduplication and ReFS volumes	<p>You cannot install the Agent for Windows on a volume that is enabled for deduplication in Windows or on an ReFS volume.</p>

Table 2-7 Special considerations for installing the Agent for Windows
(continued)

Item	Consideration
Windows Server 2003 SP2	<p>To push-install Backup Exec options to a Windows Server 2003 SP2 computer, you must enable File and Printer Sharing on the Windows Firewall Exceptions list for the following ports:</p> <ul style="list-style-type: none"> ■ 135 (RPC) ■ 445 (TCP) ■ 103X (mostly 1037) ■ 441 (RPC) <p>For more information about the Windows Firewall Exceptions list, refer to your Microsoft Windows documentation.</p> <p>During the installation process, Backup Exec sets the Remote Launch and remote access security permissions for the Administrator's group.</p> <p>You should enable the "Allow remote administration exception" group policy for the computer to which you push the installation.</p>
Windows Vista/ Server 2008/7/ 8/Server 2012/Server 2012 R2	<p>To push-install Backup Exec options to a computer that runs Windows Vista/Server 2008/7/8/Server 2012/Server 2012 R2, you must enable certain items on the destination computer's Windows Firewall Exceptions list. You must enable the following items:</p> <ul style="list-style-type: none"> ■ File and Printer Sharing ■ Windows Management Instrumentation (WMI) <p>For more information, refer to your Microsoft Windows documentation.</p> <p>To push-install to a computer that runs Windows Vista/Server 2008/7/8/Server 2012, the destination computer must be part of a domain.</p> <p>For more information, refer to the Microsoft knowledge base.</p>

Table 2-7 Special considerations for installing the Agent for Windows
(continued)

Item	Consideration
Symantec Endpoint Protection 11.0 or later	To push-install options to a computer that runs Symantec Endpoint Protection (SEP) version 11.0 or later, you must configure SEP to share files and printers. File and printer sharing is turned off by default.
Trust the Backup Exec server and remote computer	When you connect to a remote computer from the Backup Exec server, you must establish a trust between the Backup Exec server and the remote computer to ensure secure communication. To establish the trust, you must add the remote computer to the list of servers on the Backup and Restore tab. See “About the list of servers on the Backup and Restore tab” on page 152.

The installation process creates an installation log named Bkupinst2014.htm on the computer where Backup Exec is installed, and also creates an installation log named RAWSinst2014.htm on the remote computer.

See [“About the installation log”](#) on page 119.

If there are problems installing the Agent for Windows using this method, you can try to manually install the Agent for Windows.

See [“Using a command prompt to install the Agent for Windows on a remote computer”](#) on page 98.

To push-install the Agent for Windows to remote computers

- 1 Do one of the following:

To push-install the Agent for Windows to remote computers from the installation media

Do the following steps in the order listed:

- From the installation media browser, click **Installation**, and then click **Backup Exec**.
- On the **Welcome** panel, select **I accept the terms of the license agreement**, and then click **Next**.
- Click **Custom installation**.
- Uncheck **Local Installation**, and then check **Remote Installation**.
- Click **Next**.

To push-install the Agent for Windows to remote computers from the Backup Exec server

Click the Backup Exec button, select **Installation and Licensing**, and then select **Install Agents and Backup Exec Servers on Other Servers**.

2 Do one of the following

To install the Agent for Windows on one remote computer

Do the following steps in the order listed:

- On the **Remote Computers** panel, click **Add**.
- Select **Add a Single Computer**.
- Select **Agent for Windows**, and then click **Next**.
- Type the fully qualified name of the remote computer or click **Browse Remote Computers** to locate the remote computer.

To install the Agent for Windows on multiple computers using the same settings

Do the following steps in the order listed:

- On the **Remote Computers** panel, click **Add**.
- Select **Add Multiple Computers with the Same Settings**.
- Select **Agent for Windows**, and then click **Next**.
- Type the fully qualified name of the remote computer or click **Browse** to locate the remote computer.
- Click **Add to List**.
 Type the fully qualified name and then click **Add to List** for all of the remote computers for which you want to push install the options.

- 3 Under **Remote computer credentials**, type the credentials that Backup Exec can use to connect to the remote servers.

You must use Administrator credentials.

- 4 Click **Next**.
- 5 In the **Destination Folder** field, enter the path where you want to install the files.
- 6 Click **Next**.
- 7 After Backup Exec validates the remote computers, you can change the list in any of the following ways:

To manually add one remote computer

Click **Add**, and then click **Add a Single Computer**.

To manually add multiple remote computers

Click **Add**, and then click **Add Multiple Computers with the Same Settings**.

To add multiple remote computers by importing an existing list of computers

Click **Import and Export**, and then select one of the following options:

- Select **Import from File** to enable Backup Exec to add the names of the remote computers from a selected list.
- Select **Import Servers Published to this Backup Exec server** to enable Backup Exec to add the names of all the remote computers that are set up to publish to this Backup Exec server.

You must enter remote computer logon credentials for the list of remote computers.

To change the product that you selected to install or to change other properties you selected for this installation

Select the remote computer that you want to change, and then click **Edit**.

To delete a remote computer from the list

Select the remote computer that you want to delete, and then click **Delete**.

To save this list of remote computers and the associated remote computer logon credentials

Verify that **Save the server list for future remote install sessions** is checked.

This option enables the names of all of the remote computers and their credentials to be added automatically the next time you want to install Backup Exec or options to these remote computers

To save the list of remote computers to an XML file

Click **Import and Export**, and then click **Export to File**.

You can select the location to save the XML file. This option is useful if you want to use the same list for multiple Backup Exec servers. When you import the list, you must re-enter the remote computer logon credentials.

To fix the errors that were located during the validation

Right-click the name of the computer, and then click **Fix Errors**.

To enable Backup Exec to attempt to re-validate an invalid remote computer

Right-click the name of the computer, and then click **Retry Validation**.

- 8 After all of the computers in the list are validated and the list is complete, click **Next**.

- 9 Review the Backup Exec installation summary, and then click **Install**.

See [“About the installation log”](#) on page 119.

- 10 Click **Next**, and then click **Finish** to exit the wizard.

If you did not restart the remote computer, you may need to do it now in order for the configuration to take effect.

Installing updates to the Agent for Windows on remote computers

When a Backup Exec server is updated with patches, an alert is generated to warn you that the Agent for Windows on remote computers must be updated with the same patches. Additionally, in the properties for the remote computer, the property **Do the updates installed on this server match the updates installed on the backup server** indicates whether the remote computer is up to date with the Backup Exec server. From the Backup Exec console, you can update the remote computers immediately, at a scheduled time, or on a recurring schedule. You can also update a group of remote computers together.

See [“Updating Backup Exec with LiveUpdate”](#) on page 121.

To install updates for the Agent for Windows

- 1 On the **Backup and Restore** tab, right-click the remote computer or the group that needs to be updated.
- 2 Select **Update**.

- 3 On the **Install Updates** dialog box, select the option for when you want to install the updates.

Recurrence	Select this option to create a recurring schedule for the job.
Recurrence Pattern	If you choose to make the job recur on a schedule, configure the frequency with which the job recurs. You can select to run the job in hourly, daily, weekly, monthly, or yearly increments.
Starting on	Enter the date on which you want the schedule to take effect.
Calendar	View all the scheduled jobs on the calendar to check for scheduling conflicts.
Keep the job scheduled for x hours before it is rescheduled	Specify the maximum amount of time after the scheduled start time at which you want Backup Exec to consider the job to be missed and reschedule it.
Cancel the job if it is still running x hours after its scheduled start time	Specify the amount of time after the job's scheduled start time at which you want Backup Exec to cancel the job if it is still running.
Run now with no recurring schedule	Select this option to run the job immediately without scheduling any more instances of it for the future.
Run on	Select a specific date on which to run the job without scheduling any more instances of it for the future.
Restart the computer automatically after installing the updates to the Symantec Backup Exec Agent for Windows when a restart is required	Select this option to enable Backup Exec to automatically restart the remote computer, if required.

- 4 Click **OK**.

Installing the Agent for Windows in an Active Directory network

You can centrally manage the installation of the Backup Exec Agent for Windows to computers in an Active Directory network. You configure the installation once, and then use a Group Policy Object to assign that installation to computers in an

Organizational Unit. The options are installed automatically whenever a computer in the Organizational Unit is started.

Note: Review your organization’s deployment plans before you implement a rollout of the Backup Exec Agent for Windows to client computers. You should also review your Group Policy Desktop Management and Active Directory documentation.

Table 2-8 How to install the Agent for Windows in an Active Directory Network

Action	Description
Create a transform for the Agent for Windows. See the section called “Creating a transform” on page 96.	<p>A transform contains the changes that you want to make to the Windows Installer package for the Agent for Windows when a computer starts, such as the installation path. You must create separate transforms for 32-bit computers and 64-bit computers.</p> <p>Requirements to create a transform are as follows:</p> <ul style="list-style-type: none">■ The computer on which you want to create the transform must have Microsoft Windows 2003 or later.■ The computers on which you want to install the Agent for Windows must be running MSXML 6.0.■ The computers on which you want to install the Agent for Windows must be running Microsoft Visual C++ Runtime 8.0/9.0/10.0/11.0. <p>You can configure a Group Policy object to deploy all of the Microsoft Visual C++ Runtime components or install them manually on each Agent for Windows computer. For more information about configuring a Group Policy object, see your Microsoft Windows documentation.</p> <ul style="list-style-type: none">■ Only assignment to computers is supported. Assignment to users is not supported.

Table 2-8 How to install the Agent for Windows in an Active Directory Network
(continued)

Action	Description
<p>Create a distribution point (share) that contains the source file of the Agent for Windows that you want to install.</p> <p>See the section called “Creating a software distribution point (share)” on page 97.</p>	<p>You must copy the transform that you create, and the Backup Exec RAWS32 or RAWSX64 directory, to the distribution point.</p>
<p>Configure a Group Policy Object to assign the transform and the RAWS32 or RAWSX64 directory in the distribution point to computers in an Active Directory Organizational Unit.</p> <p>See the section called “Configuring a Group Policy Object” on page 97.</p>	<p>The software is installed automatically when the computers in the Organizational Unit are started.</p>

Creating a transform

To create a transform

- 1 Do one of the following:
 - From the Backup Exec installation media browser, click **Install Products**, and then click **Agent for Windows**.
 - From a Backup Exec server on which Backup Exec is installed, go to \Program Files\Symantec\Backup Exec\Agents\RAWS32 or RAWSX64 and double-click **Setup.exe**.
- 2 On the **Welcome** panel, click **Next**.
- 3 On the **Install Type** panel, click **Create a Transform to use Active Directory to install the Agent for Windows**, and then click **Next**.
- 4 On the **Install Option** panel, in the **Destination Folder** area, enter the path where you want to install the files.
- 5 Click **Next**.
- 6 Enter a file name and a path where the transform will be created, and then click **Next**.

Use a meaningful file name for the transform. For example, the name can include the names of the options in the transform and the platform you plan to apply the transform to, such as AgentDefaultPathNoPublishing.
- 7 To create the transform, click **Install**.
- 8 After the transform is created, set up a distribution point for the source files.

Creating a software distribution point (share)

To install the Agent for Windows in an Active Directory network, you must create a software distribution point after you create a transform.

Table 2-9 How to create a software distribution point (share)

Step	Description
Step 1	Create a shared folder, and then set permissions so that the client computers that will run the installation have access to the shared folder.
Step 2	Copy the following directories from the Backup Exec server to the shared folder: <ul style="list-style-type: none"> ■ RAWS32 or RAWSX64 ■ MSXML ■ VCRedist ■ DotNetFx By default, these folders are located in \Program Files\Symantec\Backup Exec\Agents.
Step 3	Copy the transform from the path where it was created to the RAWS32 or RAWSX64 directory on the shared folder.
Step 4	Configure a Group Policy Object to deploy the source files.

Configuring a Group Policy Object

To install the Agent for Windows in an Active Directory network, you must configure a Group Policy Object after you create a software distribution point and create a transform.

To configure a Group Policy Object to deploy the software

- 1 From the Active Directory snap-in that manages users and groups, click **Properties**, and create a new Group Policy Object or edit an existing one.
Refer to your Microsoft Windows documentation for information on creating a Group Policy Object.
- 2 Under **Computer Configuration**, expand **Software Settings**.
- 3 Right-click **Software Installation**, click **New**, and then click **Package**.

- 4 On the **File Open** dialog box, browse to the software distribution point by using the Universal Naming Convention (UNC) name, for example, \\server name\share name, select the package file, and then click **Open**.
- 5 Select the package file **Symantec Backup Exec Agent for Windows.msi**, and then click **Open**.
- 6 When you are prompted, apply the **Advanced Option**.
- 7 After Active Directory checks the MSI package, on the **General Properties** tab, make sure that the correct versions of the options are being installed.
- 8 On the **Deployment** tab, set up the configuration for your environment.

Make sure the option **Make this 32-bit x86 application available to WIN64 machines** is not selected.

If you want the Agent for Windows to be uninstalled if the computer is removed from the Organization Unit, select the option **Uninstall this application when it falls out of the scope of management**.
- 9 On the **Modifications** tab, click **Add**, browse to the share, and select the transform that you created.
- 10 Select **Open**, and make any other changes that are necessary, and then click **OK**.
- 11 Close all of the dialog boxes.

When a computer in the Organizational Unit that you specified is started, the transform is processed and the options that you specified are installed.
- 12 View the installation log that is created on the destination computers to verify the installation of the Agent for Windows.

Using a command prompt to install the Agent for Windows on a remote computer

You can install the Agent for Windows by using a command prompt.

The installation process creates an installation log named RAWSin2014.htm.

See [“About the installation log”](#) on page 119.

To use a command prompt to install the Agent for Windows on a remote computer

- 1 At a remote computer, map a drive letter to the Agents directory. By default, the Agents directory is located at the following path:
`\Program Files\Symantec\Backup Exec\Agents`
 or you can copy the following folders to the same local directory:

To install to a 32-bit computer:	RAWS32, MSXML, VCRedist, and DotNetFx folders
To install to a 64-bit computer:	RAWSX64, MSXML, VCRedist, and DotNetFx folders
- 2 Open a command prompt and type the drive letter that you mapped in step 1 and the following path:

To install to a 32-bit computer:	<code>\RAWS32</code>
To install to a 64-bit computer:	<code>\RAWSX64</code>

3 Do one of the following:

To install the Agent for Windows to a 32-bit computer without publishing enabled Run the following command:

```
setup.exe /RANT32: /S: /DISADVRT:
```

To install the Agent for Windows to a 32-bit computer with publishing enabled Run the following command:

```
setup.exe /RANT32: /S: /ADVRT:
Backup Exec server name 1 Backup
Exec server name 2
```

To install the Agent for Windows to a 32-bit computer and restart the computer automatically Run the following command:

```
setup.exe /RANT32: /S: /BOOT:
```

To install the Agent for Windows to a 64-bit computer without publishing enabled Run the following command:

```
setup.exe /RAWSX64: /S: /DISADVRT:
```

To install the Agent for Windows to a 64-bit computer with publishing enabled Run the following command:

```
setup.exe /RAWSX64: /S: /ADVRT:
Backup Exec server name 1 Backup
Exec server name 2
```

To install the Agent for Windows to a 64-bit computer and restart the computer automatically Run the following command:

```
setup.exe /RAWSX64 /S: /BOOT:
```

The Agent for Windows is installed on the remote computer in the following directory:

If you installed the Agent for Windows to a 32-bit computer: \Program Files\Symantec\Backup Exec\RAWS32

If you installed the Agent for Windows to a 64-bit computer: \Program Files\Symantec\Backup Exec\RAWSx64

Using a command prompt to uninstall the Agent for Windows from a remote computer

You can uninstall the Agent for Windows by using a command prompt.

To use a command prompt to uninstall the Agent for Windows from a remote computer

- 1 At the remote computer, map a drive letter to the Agent for Windows directory using the following path:

To uninstall the Agent for Windows from a 32-bit computer: `\Program Files\Symantec\Backup Exec\Agents\RAWS32`

To uninstall the Agent for Windows from a 64-bit computer: `\Program Files\Symantec\Backup Exec\Agents\RAWSX64`

- 2 Open a command prompt, and then type the drive letter that you mapped in step 1.
- 3 Run the following command:

To uninstall the Agent for Windows from a 32-bit computer: `setup.exe /RANT32: /S: /U:`

The `/S:` parameter is used to run the operation in silent mode, without the benefit of a user interface. The `/U:` parameter specifies an uninstall operation.

To uninstall the Agent for Windows from a 64-bit computer: `setup.exe /RAWSX64: /S: /U:`

See [“Using a command prompt to install the Agent for Windows on a remote computer”](#) on page 98.

Using a command script to install the Agent for Windows

You can use command script files to install the Agent for Windows. The command script files are included in the Backup Exec installation directory.

The installation process creates an installation log named `RAWSinst2014.htm`.

See [“About the installation log”](#) on page 119.

To use a command script to install the Agent for Windows

- 1 Map a drive letter to the Agents directory on a Backup Exec server. By default, the Agents directory is located at the following path:

\Program Files\Symantec\Backup Exec\Agents

- 2 Do one of the following:

To install the Agent for Windows on a 32-bit computer

In the RAW32 directory, double-click **setupaa**.

To install the Agent for Windows on a 64-bit computer

In the RAW64 directory, double-click **setupaa64**.

Using a command script to uninstall the Agent for Windows

A command script file is available to uninstall the Agent for Windows.

To use a command script to uninstall the Agent for Windows

- 1 Map a drive letter to the Agents directory on a Backup Exec server. By default, the Agents directory is located at the following path

\Program Files\Symantec\Backup Exec\Agents

- 2 Do one of the following:

To uninstall the Agent for Windows from a 32-bit computer

In the RAW32 directory, double-click **Uninstallaafo**.

To uninstall the Agent for Windows from a 64-bit computer

In the RAW64 directory, double-click **Uninstallaafo64**.

- 3 Restart the remote computer.

See [“Using a command script to install the Agent for Windows”](#) on page 101.

Installing the Remote Administrator

The Remote Administrator lets you administer the Backup Exec server from a remote Windows server or workstation. To support the Remote Administrator, the Backup Exec system services must be running on the Backup Exec server that you want to administer.

Note: The Remote Administrator can be used to administer the servers that have Backup Exec 2014 installed. It cannot be used to administer the servers that have Backup Exec 2012 or earlier installed. However, you can use side-by-side installations of the Remote Administrator for Backup Exec 2014 and Backup Exec 2012.

To install the Remote Administrator

- 1 From the installation media browser, click **Install Products**.
- 2 Click **Backup Exec**.
- 3 On the **Welcome** panel, select **I accept the terms of the license agreement**, and then click **Next**.
- 4 On the **Installation Type** panel, select **Custom installation**, and then click **Next**.
- 5 Check **Local Installation**, and then click **Install Remote Administration Console only**.
- 6 Click **Next**.
- 7 On the **Destination** panel, do the following:
 - Review the disk space requirements for the installation.
 - To change the location where the files are installed, click **Change** to select another directory for the installation.
- 8 Click **Next**.
- 9 Review the Backup Exec installation summary, and then click **Install**.
- 10 Click **Next**, and then click **Finish** to exit the wizard.

See [“Running the Remote Administrator”](#) on page 105.

Installing the Remote Administrator using the command line

You can use silent mode installation to install the Remote Administrator. Options for the Remote Administrator are specified with the use of additional command switches.

Note: The Remote Administrator can be used to administer the servers that have Backup Exec 2014 installed. It cannot be used to administer the servers that have Backup Exec 2012 or earlier installed. However, you can use side-by-side installations of the Remote Administrator for Backup Exec 2014 and Backup Exec 2012.

To install the Remote Administrator using the command line

- 1 Open a Windows command prompt.
- 2 Change to the drive containing the Backup Exec installation media.
- 3 Change directories to one of the following:

For 32-bit computers \be\winnt\install\be32

For 64-bit computers \be\winnt\install\bex64

- 4 Type `setup /RA:` and the appropriate switches. For example:

```
setup /RA: /S:
```

The command line switches used for silent mode installation of the Remote Administrator are described in the following table.

Remember the following general rules for using these switches:

- Substitute values appropriate for your environment for values in *italics*; for example, substitute your password for *password*.
- Enclose the value in quotation marks if it contains spaces, such as "Program Files\Symantec\Backup Exec".

Table 2-10 Command line switches for Remote Administrator silent mode installation

Switch	Additional Switches	Description
/RA:		Installs Remote Administrator using the options that are specified with the additional switches.
	/DEST:" <i>drive:path</i> "	Specifies the local path where Remote Administrator will be installed. Otherwise, the default path <code>Program Files\Symantec\Backup Exec</code> is used.
	/DOCS:	Installs the online documentation.

Table 2-10 Command line switches for Remote Administrator silent mode installation (*continued*)

Switch	Additional Switches	Description
	/NOINSTALL:	Lets you select all install options without actually installing the Backup Exec software. This option can be used with the /CPF: switch.
	/CPF:"path\filename.cpf"	Creates a file containing all of the installation parameters provided. Note that the file is not encrypted, which exposes parameters such as the password.
-?		Provides help on all command-line operations, usage, and special switches.

See [“Installing Backup Exec using the command line \(silent mode\)”](#) on page 108.

Running the Remote Administrator

The Remote Administrator lets you administer the Backup Exec server from a remote Windows server or workstation. To support the Remote Administrator, the Backup Exec server requires that the Backup Exec system services must be running.

You may be prompted for a user name and password to browse some network shares even if you are logged into the Remote Administrator computer under an account that is valid for those shares. Provide a domain-qualified user name and password when prompted (for example, domain1\howard).

For workgroup accounts, when logging in between different workgroups, you can provide only a user ID when prompted, and leave the workgroup line blank.

Note: The Remote Administrator can be used to administer the servers that have Backup Exec 2014 installed. It cannot be used to administer the servers that have Backup Exec 2012 or earlier installed. However, you can use side-by-side installations of the Remote Administrator for Backup Exec 2014 and Backup Exec 2012.

See [“Installing the Remote Administrator”](#) on page 102.

To run the Remote Administrator

- 1 Click **Start**.
- 2 Point to Programs, and then click **Symantec Backup Exec**.

If you are connecting to a Remote Administration Console from a Backup Exec server, click the Backup Exec button, and then select **Connect to Backup Exec Server**.

3 Select the appropriate options.

Manage services

Select this option to access the Backup Exec Services Manager to stop and start services or to set the logon credentials that are used to run the services.

Server name

Enter the name of the Backup Exec server. You can select the name from the list or type the name of the server if you are running the Remote Administrator from a Backup Exec server.

Each server in the domain that has Backup Exec installed automatically appears in the list box.

User name

Enter an administrator user name for the server to which you want to connect. Enter the user name using the Domain\Username format.

You cannot log on to the Remote Administration Console with a user name that has a blank password on Windows Server 2003/2008 and Vista computers. You must configure Windows to allow blank passwords. Otherwise, the error message "Logon failure: user account restriction" appears. For more information, see your Windows documentation.

Password

Enter the password for the user.

Domain

Enter the domain to which the user belongs. Enter the name of the domain that was used in the **User name** option.

The status of the local services appears at the bottom of this dialog box. If you try to connect to a server and the connection fails, this dialog box displays the services status for the server you attempted to connect to.

4 Click **OK**.

Installing Backup Exec using the command line (silent mode)

Installing Backup Exec using the command line is referred to as silent mode installation. This method of installation uses the setup.exe program on the Backup Exec installation media, a series of command switches, and the /S: switch.

Requirements for Command Line Installation include the following:

- Backup Exec installation media.
- Administrator privileges on the computer where you want to install, configure, or uninstall Backup Exec.

The installation process creates an installation log named Bkupinst2014.htm on the computer where Backup Exec is installed.

See [“About the installation log”](#) on page 119.

To install Backup Exec using the command line (silent mode)

- 1 Open a Windows command prompt.
- 2 Change to the drive containing the Backup Exec installation media.
- 3 Change directories to one of the following:

For 32-bit computers \be\winnt\install\be32

For 64-bit computers \be\winnt\install\bex64

- 4 Type `setup /TS:` and the appropriate switches. For example:

```
setup /TS: /USER:<user> /DOM:domain /PASS:password
/SLF:C:\path\slf.slf,C:\path\slf2.slf /S:
```

See [“Command line switches for silent mode installation of Backup Exec”](#) on page 109.

If you use the command line switches without the /S: switch, the Backup Exec installation program launches with the command line parameters as defaults for the installation options. For example, if /S: had been left in the above example, the Backup Exec installation program launches with the user name, domain, password, and license appearing on the installation dialog boxes.

- 5 Press **Enter**.

Command line switches for silent mode installation of Backup Exec

The command line switches used for silent mode installation of Backup Exec are described in the following table.

The following are general rules for using these switches:

- Substitute values appropriate for your environment for the values that are shown in italics. For example, substitute a user name for *user*.
- Enclose the value in quotation marks if it contains spaces, such as "Operations Weekly Backup".

See [“Installing Backup Exec using the command line \(silent mode\)”](#) on page 108.

Table 2-11 Command line switches for silent mode installation of Backup Exec

Switch	Additional Switches	Description
/TS:		Installs Backup Exec using the options that are specified with the additional switches. The /USER: <i>user</i> , /DOM: <i>dm</i> , /PASS: <i>pw</i> switches are required.
	/USER: <i>user</i> /DOM: <i>dm</i> /PASS: <i>pw</i>	Required. Specifies an existing user, domain, and password for the Backup Exec system service account. Silent mode installation does not create a user. Note: When using /PASS:, if a quote is needed as part of the password, specify it as \". For example, if the password is pass\"word, type it as /PASS:pass\"word. If the characters \" are used as part of the password, you must precede each character with a \. For example, if the password is pass\"word, type it as /PASS:pass\\\"word.
	/DEST: <i>drive</i> : <i>path</i>	Specifies the local path where Backup Exec is installed. Otherwise, the default path <code>Program Files\Symantec\Backup Exec</code> is used.
	/DOCS:	Installs the online documentation.

Table 2-11 Command line switches for silent mode installation of Backup Exec
(continued)

Switch	Additional Switches	Description
	<i>/BELANG:language</i>	<p>Installs the Backup Exec language resource files.</p> <p>Specify one or more of the following switches with the Backup Exec language switch to indicate which language files you want to install:</p> <ul style="list-style-type: none"> ■ EN installs English. ■ DE installs German. ■ ES installs Spanish. ■ FR installs French. ■ IT installs Italian. ■ PT installs Portuguese. ■ RU installs Russian. ■ JP installs Japanese. ■ KO installs Korean. ■ ZH installs Simplified Chinese. ■ CH installs Traditional Chinese. <p>The following example shows how the /BELANG switch can be used to install English, German, and Spanish:</p> <p><i>/BELANG:"EN DE ES"</i></p>
	<i>/NOINSTALL:</i>	<p>Lets you select all install options without installing the Backup Exec software.</p> <p>This option can be used with the /CPF: switch.</p>
	<i>/CPF:path\filename.cpf</i>	<p>Creates a file containing all of the installation parameters provided. Note that the file is not encrypted, which exposes parameters.</p>

Table 2-11

Command line switches for silent mode installation of Backup Exec
(continued)

Switch	Additional Switches	Description
	/SLF: <i>slf file</i>	<p>Specifies one or more licenses to use for installing Backup Exec and additional options. Licenses are not required to install the Remote Administrator. You may specify up to 99 licenses. If none are specified, then a trial copy of Backup Exec is installed.</p> <p>The following examples show how the /SLF switch can be used:</p> <p>/SLF:C:\path\slf1.slf</p> <p>/SLF:C:\path\slf1.slf,C:\path\ slf2.slf, C:\path\ slf3.slf</p> <p>Note: If you install a license for an option or agent, you must also type a switch that specifies the option or agent. The switches that specify an option or agent are included in this table.</p>
	/DISABLETELEMETRY:	Disables the option to send Backup Exec usage data to Symantec over the web.
	/BOOT:	Automatically initiates a restart of the computer during a silent install or uninstall.

Table 2-11 Command line switches for silent mode installation of Backup Exec
(continued)

Switch	Additional Switches	Description
	/TD:NEW, ALL, or NONE	<p>Note: None of the additional /TD switches are supported for Windows Server 2012 or later.</p> <p>/TD:NEW installs tape drivers only for the drives that do not have drivers loaded.</p> <p>/TD:ALL installs tape drivers for all drives.</p> <p>/TD:NONE does not install tape device drivers.</p> <p>Note: To install the Symantec tape drivers in Windows 2008 R2, the Windows driver signing policy must be set to Ignore. However, for Windows 2003/2008 the driver installation fails when the signing policy is set to Ignore. You can install the drivers using Symantec Device Driver Installer (tapeinst.exe) instead. See your Microsoft Windows documentation for more information about the signing policy.</p>
	/DBSERVER: <i>server\instance</i>	Installs the Backup Exec Database to the specified SQL server.
	/BACKUPDATA:	Specifies if a copy of your current Backup Exec Database should be stored during an upgrade of Backup Exec.

Table 2-11

Command line switches for silent mode installation of Backup Exec
(continued)

Switch	Additional Switches	Description
	/BACKUPDIR:	Specifies the location to store the Backup Exec Database during an upgrade of Backup Exec. The folder you select must be empty. A copy of your current Backup Exec Database is placed in the location that you specify and is used if the upgrade fails. Note: This switch is ignored if the corresponding /BACKUPDATA: switch is not specified.
	/DBINSTPATH: <i>SQL Express destination folder</i>	Installs the default instance of SQL Express in the specified folder.
	/NOUPDATE:	Skips the installation of Symantec LiveUpdate.
	/ADVRT: <i>Backup Exec server name</i>	Installs the Agent for Windows and enables publishing.
	/DISADVRT:	Installs the Agent for Windows without publishing it.
	/SQLXSETUP: <i>SQL Express Install Package</i>	Specifies the location of the language-specific install package for Microsoft SQL Server 2008 R2 Express Edition with SP2.
	/LOADER:	Installs the Library Expansion Option. This option is not supported for Windows Server 2012 or later.
	/QUICKSTART:	Installs the QuickStart Edition. You must enter a QuickStart license to install this edition.
	/BEWS:	Installs the standard Backup Exec Edition. You must enter a Backup Exec license to install this edition.
	/NFR:	Installs the Not For Resale Edition. You must enter a Not For Resale license to install this edition.

Table 2-11

Command line switches for silent mode installation of Backup Exec
(continued)

Switch	Additional Switches	Description
	/TRIAL:	Installs the Trial Edition.
	/SMALLBUSINESS:	Installs the Small Business Edition. You must enter a Small Business Edition license to install this edition.
	/ESSENTIAL:	Installs the Essential Protection Edition. You must enter an Essential Protection license to install this edition.
	/TOTAL:	Installs the Total Protection Edition. You must enter a Total Protection license to install this edition.
	/VRAY:	Installs the V-Ray Edition. You must enter a V-Ray license to install this edition.
	/CAPACITY:	Installs the Capacity Edition. You must enter a Capacity license to install this edition.
	/APPLICATIONS:	Installs the Agent for Applications and Databases.
	/VIRT:	Installs the Agent for VMware and Hyper-V.
	/ENTSERVER:	<p>Installs the Enterprise Server Option.</p> <p>You must use one or both of the following switches with the Enterprise Server Option switch to indicate which options you want to install.</p> <ul style="list-style-type: none"> ■ /CASO: Installs the Central Admin Server Option. ■ /ADBO: Installs the Advanced Disk-based Backup Option.

Table 2-11 Command line switches for silent mode installation of Backup Exec
(continued)

Switch	Additional Switches	Description
	/ADBO:	Installs the Advanced Disk-based Backup Option. You must use /ENTSERVER: with this switch.
	/CASO:	Installs the Central Admin Server Option. You must use /ENTSERVER: with this switch.
	/MMS:CAS <i>server name</i>	Creates a managed Backup Exec server for use with the Central Admin Server Option.
	/CASOPVLLOCAL: <1 or 0>	/CASOPVLLOCAL:1 indicates that device and media data is stored locally on the managed server. Use this switch with /MMS:. /CASOPVLLOCAL:0 indicates that device and media data is stored on the administration server. Use this switch with /MMS:.
	/ACCESSCATALOGSANDRESTORE:	Enables unrestricted access to catalogs and backup sets for restore. This switch is used with the /MMS:<CAS server name> switch and it replaces the /SSO:<primary server name> switch.
	/NTA:	Installs the Agent for Windows.
	/MAC:	Installs the Agent for Mac.
	/NDMP:	Installs the NDMP Option.
	/RALS:	Installs the Agent for Linux.
	/STORPROV:	Installs the Storage Provisioning Option.
	/FIXEDSPO:	Installs the basic Storage Provisioning Option.
	/DEDUPE:	Installs the Deduplication Option.

Table 2-11

Command line switches for silent mode installation of Backup Exec
(continued)

Switch	Additional Switches	Description
	/EXCHARCH:	<p>Installs the Exchange Mailbox Archiving Option.</p> <p>This switch is used with the /APPLICATIONS: switch.</p>
	/NTFS:	Installs the File System Archiving Option.
	/VTL:	Installs the VTL Unlimited Drive Option.
	/RMAL:	Installs the Remote Media Agent for Linux.
	/COPYCONFIG:	Installs the Copy Server Configuration option.
	/BEAODNSALIAS:	Specifies a fully-qualified DNS alias for the Exchange Mailbox Archiving Option.
	/BEAOINDEXPATH:	Specifies the location of the index files for the Exchange Mailbox Archiving Option. The default location is %ProgramFiles%\Symantec\Backup Exec\ArchiveIndex.
	/BEAOCACHEPATH:	Specifies the location of the temporary cache for the Exchange Mailbox Archiving Option. The default location is %ProgramFiles%\Symantec\Backup Exec\AOCache.
	/BEAOCACHESIZE:	Specifies the maximum size of the temporary cache location for the Exchange Mailbox Archiving Option. The default size is 10 GB.
-?		Provides help on all command-line operations, usage, and special switches.

Creating and using installation parameter files

If you use the command line switches without the /S: switch, the Backup Exec installation program launches with the command line parameters as defaults for the installation options. For example, suppose you type:

```
SETUP /TS: /USER:user /DOM:domain /PASS:password /SLF:"C:\path
name\slf1.slf"
```

The Backup Exec installation program is launched. The screens that let you enter the logon credentials and the license will appear with the information you provided on the command line.

You can also use the /CPF: command to create a parameter file that contains all of the command line options you provided. This parameter file can then be used to provide the options for installing either Backup Exec or the Remote Administrator. Note that the file is not encrypted, which exposes parameters such as the password.

To create installation parameter files

- 1 Open a Windows command prompt.
- 2 Change to the drive containing the Backup Exec installation media.
- 3 Change directories to one of the following:

For 32-bit computers `\be\winnt\install\be32`

For 64-bit computers `\be\winnt\install\bex64`

- 4 Type `setup /TS:` and the appropriate switches, including /CPF: and the full path name of the parameter file. For example, type:

```
setup /TS: /USER:user /DOM:domain /PASS:password /SLF:"C:\path
name\slf1.slf" /CPF:"A:\file name" /S:
```

Backup Exec will be installed on your server and a parameter file containing the user name, domain, password, and license will be saved to a removable device. You can use this parameter file to install to another computer.

To use installation parameter files

- 1 Open a Windows command prompt.
- 2 Change to the drive containing the Backup Exec installation media.
- 3 Change directories to one of the following:

- For 32-bit computers `BE\WINNT\INSTALL\BE32`

- For 64-bit computers `BE\WINNT\INSTALL\BE64`

- 4 Type: `SETUP /PARAMS:"A:\file name" /S:`
- 5 If you want to overwrite a parameter, specify the new parameter. For example, to change the password, type: `SETUP /PARAMS:"A:\file name" /PASS:new password /S:`

Installing a trial version of Backup Exec agents and options

You can install a trial version of most Backup Exec agents and options at any time after the core product is licensed. Each agent and each option has its own independent trial period. When a trial period is about to expire, Backup Exec warns you with an alert.

You can view a list of agents and options that are available for a trial period. You can also view the amount of time that is left in each individual trial period.

See [“Viewing license information”](#) on page 123.

To install a trial version of Backup Exec agents and options

- 1 Click the Backup Exec button, select **Installation and Licensing**, and then select **Install Options and Licenses on this Backup Exec Server**.
- 2 Click **Next**.
Do not enter serial numbers or a Symantec License file.
- 3 Click **Next**.
- 4 Select the agents or options that you want to evaluate.
- 5 Click **Next**.
- 6 If you are prompted, enter information or choose settings for the additional options that you want to install. Click **Next** after each selection.
- 7 Read the Backup Exec installation review, and then click **Install**.

The Backup Exec services are stopped while the additional options are installed. If any active jobs are in progress, you are prompted to stop them, or to wait for the jobs to finish.

When the installation is complete, the services restart.

- 8 Click **Finish**.

About the installation log

Backup Exec creates an installation log file, named Bkupinst2014.htm, when you install Backup Exec and when you install patches. This log file can help you troubleshoot installation problems. The log file provides links to tech notes for the most common errors. If you install the Agent for Windows, a log file called RAWSin2014.htm is also created.

In addition, the text in the log file uses the following colors so you can identify warnings and errors:

Table 2-12 Installation log colors

This color	Indicates
Black	Normal operations
Orange	Warning messages
Red	Error messages

The Bkupinst2014.htm file is located in the following locations:

- For Windows 2003: %allusersprofile%\Application Data\Symantec\Backup Exec\Logs
- For Windows 2008 and later: %ProgramData%\Symantec\Backup Exec\Logs

Note: The ProgramData folder is a hidden folder. If you do not see the ProgramData folder, refer to the Microsoft Windows documentation for instructions on how to display hidden folders.

Viewing the Installation Summary Report

Backup Exec creates an Installation Summary Report that includes the configuration settings that you selected during the installation process. The Installation Summary Report is updated with the product name and serial numbers when you install additional agents or options. It is also updated when you remove agents or options from Backup Exec.

The Installation Summary Report is stored in the following locations:

- For Windows 2003: %allusersprofile%\Application Data\Symantec\Backup Exec\Logs\InstallSummary\<computer name>.htm

- For Windows 2008 and later: %programdata%\Symantec\Backup Exec\Logs\InstallSummary\<computer name>.htm

The Installation Summary Report is available for viewing from the Backup Exec Administration Console or the **Home** tab at any time after the installation has completed.

To view the Installation Summary Report

- ◆ Do one of the following:

To view the Installation Summary Report from the Administration Console	Click the Backup Exec button, select Installation and Licensing , and then select Installation Summary Report .
---	---

To view the Installation Summary Report from the Home tab	On the Home tab, in the Installation and Upgrades group, click Installation Summary Report .
--	---

Repairing Backup Exec

If you have missing or corrupted Backup Exec files or registry keys on the local Backup Exec server, run the Repair option. The program stops all Backup Exec services, reinstalls corrupted files and registry keys, reinstalls tape devices (standalone drives and libraries), and restarts the services. The database is not reinstalled.

Any changes that are made to Backup Exec program files and registry keys are reset to the original settings.

To repair Backup Exec

- 1 Close the Backup Exec application.
- 2 From the Windows Control Panel, select the option to uninstall a program.
- 3 Select **Symantec Backup Exec 2014**, and then click **Change**.
- 4 Select **Local installation** and **Repair**, and then click **Next**.
Ensure that the **Remote installation** option is not selected.
- 5 If you are prompted to enter credentials for the Backup Exec service account, type the correct credentials, and then click **Next**.
- 6 Select **Install**.
You may be prompted to insert the installation media.
- 7 Click **Finish**.

Updating Backup Exec with LiveUpdate

Symantec LiveUpdate, which provides updates of Backup Exec, is installed automatically with Backup Exec. Backup Exec installs the latest version of LiveUpdate. If a previous version of LiveUpdate is detected on the computer, Backup Exec upgrades it.

LiveUpdate can be run manually or can be configured to run automatically. If you enable the automatic update option, you can configure LiveUpdate to poll the main Symantec web server on a scheduled interval. By default, LiveUpdate checks for updates every Sunday night at 10pm. If there is an update, LiveUpdate notifies you with an alert. The automatic update option only searches for Backup Exec updates. It does not show updates for other Symantec products that use LiveUpdate. Likewise, when LiveUpdate is scheduled to automatically update other Symantec products, it does not search for Backup Exec updates. In addition to scheduling LiveUpdate, you can also run it manually at any time. You can access LiveUpdate from Backup Exec, but you cannot access it from the Windows Start menu. If LiveUpdate installs any files, the Bkupinst2014.htm installation log file is updated with information about those files.

Note: During the installation and update process, the Backup Exec services are stopped and started one time during a LiveUpdate session, regardless of the number of updates that are being installed. All selected patches are installed in order.

When LiveUpdate installs updates on the Backup Exec server, it also determines if computers on which the Agent for Windows is installed have the latest updates. If you do not have the latest updates you receive an alert to install the updates.

You can use the LiveUpdate Administrator Utility with LiveUpdate. The LiveUpdate Administrator Utility allows an administrator to modify LiveUpdate so that network users can download program and virus definition updates from an internal server rather than going to the Symantec LiveUpdate server over the Internet.

Scheduling automatic Backup Exec updates using LiveUpdate

You can schedule LiveUpdate to check for updates as follows:

- Every day at a specific time
- Every week on a specific day of the week and at a specific time
- Every month on a specific day of the month and at a specific time

When you schedule automatic updates through Backup Exec, the settings apply only to updates for Backup Exec. Changes that you make to the LiveUpdate

schedule for Backup Exec do not affect the schedule for any other software applications that use LiveUpdate.

At the scheduled time, LiveUpdate automatically connects to the appropriate website, and then determines if your files need to be updated. Depending on the options that you select, Backup Exec either downloads and installs the files in the proper location or sends an alert to notify you that updates are available.

Backup Exec sends the following LiveUpdate alerts:

Table 2-13 LiveUpdate alerts

Backup Exec sends this alert	When
LiveUpdate Informational Alert	An update is installed successfully.
LiveUpdate Warning Alert	An update is installed successfully. However, you must restart the computer.
LiveUpdate Error Alert	An update fails to install.

To schedule automatic updates using LiveUpdate

- Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- In the left pane, select **LiveUpdate**.
- Select **Check for updates automatically according to a schedule**.
- Select the schedule to check for updates.
- Do one of the following:
 - Select **Download and install all available updates automatically** if you want the updates to be installed automatically.
 - Select **Send an alert when updates are available; do not download or install updates** if you do not want the updates to be installed automatically, but do want to receive notification that updates are available.
- Click **OK**.

Running LiveUpdate manually

You can either set a schedule for LiveUpdate or run LiveUpdate manually at any time to check for updates. You can configure LiveUpdate to run in either Interactive mode or Express mode. Interactive mode gives you the flexibility to choose which updates you want to install. Express mode automatically installs all of the Backup Exec updates. For information about how to change the LiveUpdate mode, see the LiveUpdate documentation.

Note: By default, LiveUpdate is configured for Interactive mode. If you change it to Express mode you must cancel the LiveUpdate session and restart it before the change takes place.

To run LiveUpdate manually

1 Click the Backup Exec button, select **Installation and Licensing**, and then select **LiveUpdate**.

2 Do one of the following:

If LiveUpdate is set for Express mode Click **Start**.

If LiveUpdate is set for Interactive mode Click **Next**.

See [“Viewing installed updates”](#) on page 123.

See [“Installing updates to the Agent for Windows on remote computers”](#) on page 93.

Viewing installed updates

You can view the service packs that are installed on a Backup Exec server. You must be logged on with administrator privileges.

If a service pack is installed before another service pack, that service pack no longer displays as installed since the later service pack contains the earlier service pack

A hot fix that is offered after the service pack is released is displayed with the previous service pack.

To view installed updates

- ◆ Click the Backup Exec button, select **Installation and Licensing**, and then select **Installed Updates**.

Viewing license information

You can view information about the Backup Exec options that are licensed and installed on a Backup Exec server. You can also view a list of agents and options that are available for a trial, as well as how much time is left in each individual trial period.

To view license information

- ◆ Click the Backup Exec button, select **Installation and Licensing**, and then select **License Information**.

Backup Exec maintenance contract information

After you purchase maintenance contracts for your Backup Exec products, Symantec automatically updates the Symantec Licensing Portal website with your maintenance contract information. Maintenance contract information includes the contract serial number and the contract expiration date.

To retrieve the contract expiration dates, you enter your maintenance contract serial numbers in the installation wizard. The installation wizard connects to the Symantec Web service, at which point you may be prompted to enter customer and technical contact information. The installation wizard then retrieves the maintenance contract information for each contract that you have purchased. Backup Exec then uses the contract expiration information to automatically set Backup Exec alerts that remind you to renew the maintenance contracts before they expire. Reminder alerts are set at 30-day, 60-day, and 90-day intervals, based on the expiration date of the maintenance contract. If you do not renew the maintenance contract, an alert is sent when the maintenance contract expires.

All Backup Exec products and maintenance contracts have Symantec serial numbers. The serial numbers appear on the printed certificate that you receive with your order. To determine the correct serial numbers to enter in the installation wizard, look for the Maintenance/Subscription columns on your certificate. Each maintenance contract specifies a start date and an end date; these dates appear in the Maintenance/Subscription columns.

To view Backup Exec maintenance contract information

- ◆ Click the Backup Exec button, select **Installation and Licensing**, and then select **Maintenance Contract Information**.

See [“Updating expired maintenance contracts”](#) on page 124.

Updating expired maintenance contracts

When your maintenance contracts expire, follow these steps to update them.

Note: You cannot use the Remote Administrator on a remote Windows server or workstation to update maintenance contracts.

Table 2-14 How to update expired maintenance contracts

Step	Additional information
Purchase new maintenance contracts.	Access the Symantec Licensing Portal website to purchase new contracts.

Table 2-14 How to update expired maintenance contracts (*continued*)

Step	Additional information
Wait for your new serial numbers to arrive.	The new serial numbers should arrive by email within two to five business days of the purchase date.
Launch the installation wizard from the Backup Exec Administration Console.	Use the option Install options and licenses on this Backup Exec server on the Installation and Licensing menu, which is accessed from the Backup Exec button.
Use the installation wizard to add the new serial numbers, and then remove the expired serial numbers.	After you select the expired serial numbers from the list, use the Remove option.

See [“Backup Exec maintenance contract information”](#) on page 124.

Managing maintenance contract customer numbers

Backup Exec provides a place where you can store all of your maintenance contract customer numbers. You need to provide these numbers when you call technical support.

To manage maintenance contract customer numbers

- Click the Backup Exec button, select **Installation and Licensing**, and then select **Maintenance Contract Customer Numbers**.
- Do one of the following:
 - To add a new customer number, click **New**, and then enter your customer number and any notes for this number.
 - To remove a customer number, select the number from the list, and then click **Delete**.
- Click **Close**.

About upgrading from previous versions of Backup Exec to Backup Exec 2014

You can use the Backup Exec installation media to upgrade from Backup Exec version 12.5 or later to the current version. You do not need to uninstall the previous version of Backup Exec before upgrading; the current version is installed over the

previous version. Separate installations of different versions of Backup Exec cannot exist on the same computer.

Most settings, all catalogs, and all data directories from previous versions of Backup Exec are kept, unless you choose to remove them. This version of Backup Exec can read and restore data from any previous version of Backup Exec or Backup Exec for NetWare, except where Symantec has made end-of-life decisions.

Note: Upgrading to the current version of Backup Exec from a version before 12.5 requires a multi-step process. For example, to upgrade from Backup Exec 11.x to Backup Exec 2014, you must first upgrade to a version such as Backup Exec 12.5 or Backup Exec 2010 before you can upgrade to Backup Exec 2014.

When you upgrade from previous versions, Backup Exec automatically converts your existing definitions, configurations, and jobs to the current version and some of your jobs may be combined or moved. After the migration completes, Backup Exec displays the Migration Report that you must review and acknowledge before Backup Exec continues the upgrade process. In this report, you can see how your jobs were migrated.

See [“Using the Migration Report to determine how existing jobs changed after an upgrade to Backup Exec 2014”](#) on page 129.

Backup Exec 2014 provides backward compatibility as follows:

- Backup Exec 2014 can communicate with Backup Exec 12.5 Remote Agent for Windows Systems and later.
- Backup Exec 2014 supports side-by-side installations of the Remote Administration Console for Backup Exec 2010 and later. For example, you can have a Remote Administration Console for Backup Exec 2012 and a Remote Administration Console for Backup Exec 2014.
- Backup Exec 2014 Central Admin Server Option server can communicate with Backup Exec 2010 R3 (with the most recent service packs) for the purpose of rolling upgrades.

A Remote Administration Console that uses a previous version of Backup Exec cannot be used with a Backup Exec server on which the current version is installed. For example, a Backup Exec 2010 Remote Administration Console cannot manage a Backup Exec 2014 Backup Exec server.

Before you upgrade Backup Exec, do the following:

- Delete the job histories and the catalogs that you no longer need to shorten the upgrade window.
- Run a database maintenance job.

- Verify that all available updates are installed for your current version of Backup Exec.
- Locate your license information and verify that your licenses are current. You must enter license information for Backup Exec 2014 when you upgrade.

You cannot change the configuration of your Backup Exec servers or the database location during the upgrade process. For example, you cannot change an administration server to a managed server. If you want to change the configuration of your Backup Exec servers, do it either before or after you upgrade to the current version. If you want to change the database location after the upgrade, use BEUtility.

Note: If you upgrade from a previous version of Backup Exec that uses a non-English version of Windows, you must download the SQL Server 2008 R2 Express with SP2 setup file for that language from the Microsoft website.

See [“Installing Backup Exec by using the Installation Wizard”](#) on page 66.

Upgrade checklist for Backup Exec 2010 and earlier

Before you upgrade from a previous version of Backup Exec to the current version, do the following:

- Ensure that your backups are up to date. Symantec recommends that you always run full backups before and after you upgrade the applications or operating systems on any backup sources.
- Review the overwrite protection periods of the media sets that are on disk storage. Extend the overwrite protection periods of the media sets that you want to keep longer.

Changes have been made to the way that Advanced Device and Media Management (ADAMM) handles the overwrite protection for the media sets that are on disk-based storage. Previous versions of Backup Exec let the media sets expire but did not delete them from the disk, which can cause your disk storage to become full and jobs to fail. Backup Exec 2014 now proactively frees disk space through a process that reclaims the disk space for use by new backup sets.

Backup Exec 2014 uses the disk reclamation process called data lifecycle management (DLM) instead of ADAMM to manage the data retention of backup sets on disk-based storage. DLM automatically deletes the expired backup sets on disk-based storage and reclaims the disk space for use by new backup sets. DLM deletes backup sets after the amount of time that you specify when you create a backup job expires.

For the jobs that you upgrade, DLM deletes the backup sets for which the overwrite protection period is expired.

If you want to review or change the overwrite protection periods of your expired backup sets after you upgrade, you can temporarily disable DLM, make any changes to the data retention of backup sets, and then re-enable DLM.

For more information, see the following URL:

<http://entsupport.symantec.com/umi/V-269-45>

- Disable your antivirus software.
- Check the Backup Exec Software Compatibility List (SCL) and the Hardware Compatibility List (HCL) to verify that the applications that you want to back up and your storage devices are supported with this version of Backup Exec. You can find the SCL and HCL at the following URL:
<http://entsupport.symantec.com/umi/V-269-1>
<http://entsupport.symantec.com/umi/V-269-2>
- Download all available upgrades and hot fixes for the version of Backup Exec that you want to install.
- Plan to perform the upgrade when system downtime won't affect users.
- Ensure that your serial numbers or Symantec License Files are available. You must enter new Backup Exec 2014 license information during the upgrade. Review the document *Backup Exec Licensing and Pricing Guide* for details on updated licensing and pricing at the following URL:
<http://entsupport.symantec.com/umi/V-269-50>
- Review the following topics in the *Backup Exec 2014 Administrator's Guide*:
 - *System requirements for Backup Exec*
 See “[System requirements for Backup Exec](#)” on page 63.
 - *How data lifecycle management (DLM) deletes expired backup sets on disk-based storage*
 See “[How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage](#)” on page 335.
 - *About upgrading from previous versions of Backup Exec*
 See “[About upgrading from previous versions of Backup Exec to Backup Exec 2014](#)” on page 125.
- Review the document *Best practices for installing Backup Exec* on the Backup Exec knowledge base.
- Review the document *What's Different in Backup Exec 2014* to learn about product-level differences at the following URL:
<http://www.symantec.com/docs/DOC5216>

Using the Migration Report to determine how existing jobs changed after an upgrade to Backup Exec 2014

When you upgrade Backup Exec 2014 from a previous version other than Backup Exec 2012, your existing definitions, configurations, and jobs are converted automatically to the current version. When the jobs are migrated, some of the jobs may be combined or moved. After the migration completes, Backup Exec displays the Migration Report for you to view and acknowledge. In this report, you can see how your jobs were migrated. The information that is included in the Migration Report cannot be recreated after the upgrade completes.

Symantec recommends that you review the Migration Report thoroughly to determine how your existing jobs have changed and how you may need to adjust your jobs manually. The Migration Report is available for viewing from the Backup Exec Administration Console or the **Home** tab at any time after the migration completes.

The Migration Report is stored in the following path:

\Program Files\Symantec\Backup Exec\Data\MigrationReportFiles\Data-Migration-Report.html.

Note: To view the Migration Report, JavaScript must be enabled in your web browser. If the server on which you installed Backup Exec does not have a browser with JavaScript enabled, you can copy the Migration Report to another server that has JavaScript enabled.

To view the Migration Report

◆ Do one of the following:

To view the Migration Report from the Administration Console	Click the Backup Exec button, select Installation and Licensing , and then select Migration Report .
To view the Migration Report from the Home tab	On the Home tab, in the Installation and Upgrades group, click Migration Report .

See [“Configuring the Home tab”](#) on page 139.

Post-installation tasks

For best results, do the following after installing Backup Exec:

- Create disk storage so that Backup Exec can automatically manage the lifecycle of your backup data.

See [“Configuring disk storage”](#) on page 318.

- Make sure that your storage devices are connected and configured properly.
- Decide what types of storage devices you want to use for your backup jobs. You can configure storage devices when you prepare your Backup Exec environment.
- Understand how Backup Exec provides overwrite protection for your tape media. See [“Media overwrite protection levels for tape media”](#) on page 394.
- Understand the default media set for tape media and its four-week overwrite protection period. See [“Default media sets”](#) on page 379.
- Understand Data Lifecycle Management. See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 335.
- Learn about creating new media sets with different retention periods. See [“Creating media sets for tapes”](#) on page 386.
- Decide which credentials you want your Backup Exec logon account to use when browsing and making backup selections. You can use an existing Backup Exec logon account, or create a new one. See [“Backup Exec logon accounts”](#) on page 610.

Uninstalling Backup Exec

Use Microsoft’s Add or Remove Programs option to remove Backup Exec from a computer. For additional information on Add or Remove Programs, refer to your Microsoft documentation.

Uninstalling Backup Exec also removes Symantec tape class drivers. If you reinstall Backup Exec and want to use Symantec tape class drivers, you must reinstall them.

To uninstall Backup Exec

- 1 Close Backup Exec.
- 2 From the Windows Control Panel, select the option to uninstall a program.
- 3 Select **Symantec Backup Exec 2014**, and then click **Uninstall**.
- 4 When you are prompted to confirm that you want to uninstall Backup Exec from your computer, click **Yes**.
- 5 Select whether you want to remove only the Backup Exec program files or Backup Exec and all of its associated files.

6 Click **Next**.

If the uninstall program fails, click **View Installation Log File** for additional information.

7 If you are prompted, restart the computer.

See [“Uninstalling Backup Exec options from the local Backup Exec server”](#) on page 131.

Uninstalling Backup Exec using the command line

If Backup Exec is already installed, you can use the setup.exe program to uninstall Backup Exec program files and Backup Exec data.

To uninstall Backup Exec using the command line

- 1 Open a Windows command prompt.
- 2 Change to the drive containing the Backup Exec installation media.
- 3 Change directories to one of the following:

For 32-bit computers `\be\winnt\install\be32`

For 64-bit computers `\be\winnt\install\bex64`

- 4 To remove the Backup Exec program files but keep all of the Backup Exec data, type:

```
SETUP /UNINSTALL:
```

To remove the Backup Exec program files and the Backup Exec data, type:

```
SETUP /REMOVEALL:
```

See [“Installing Backup Exec using the command line \(silent mode\)”](#) on page 108.

Uninstalling Backup Exec options from the local Backup Exec server

The Installation Wizard removes Backup Exec options from the local Backup Exec server. All corresponding files, registry keys, and configurations are removed.

Note: Symantec license files remain on the server after options are uninstalled. Do not delete the Symantec license files while Backup Exec is installed. Deleting the Symantec license files causes the trial version to go into effect.

To uninstall Backup Exec options from the local Backup Exec server

- 1 Click the Backup Exec button, select **Installation and Licensing**, and then select **Install Options and licenses on this Backup Exec Server**.
- 2 On the **Add Licenses** panel, in the **Activated Licenses** list, select the item that you want to remove, and then click **Remove Serial Numbers**.
- 3 Click **Next**.
- 4 On the **Review Licenses** panel, uncheck the check boxes for the licenses that want to remove.
- 5 Click **Next**.
- 6 On the **Configure Options** panel, verify that the option you want to remove is not checked, and then click **Next**.
- 7 Uncheck the check box for the language that you want to remove.
- 8 Click **Next**.
- 9 If you are prompted to enter credentials for the Backup Exec service account, type the correct credentials, and then click **Next**.
- 10 Read the installation summary, and then click **Install** to start the process.
- 11 When the Installation Wizard has completed, click **Finish**.

See [“Uninstalling Backup Exec”](#) on page 130.

Getting Started

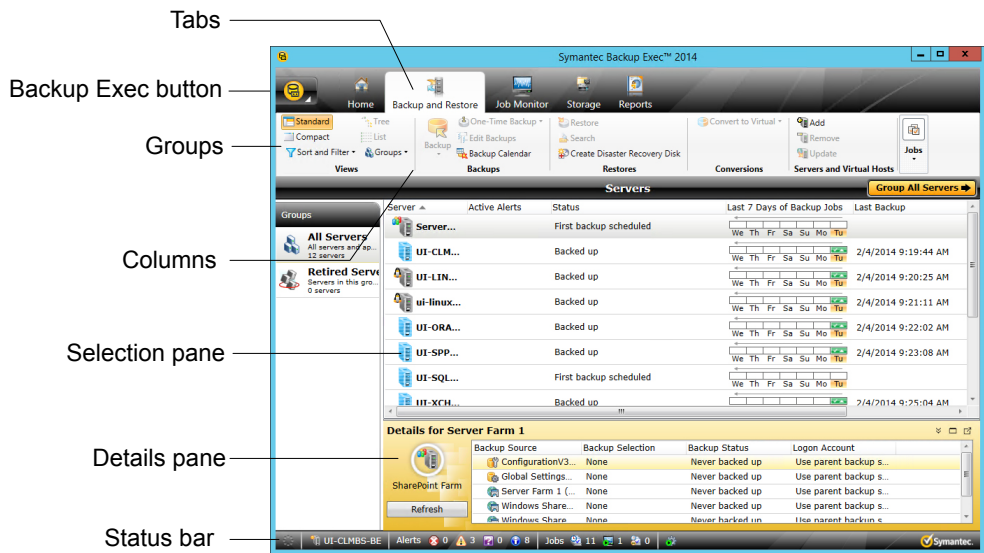
This chapter includes the following topics:

- [About the Backup Exec Administration Console](#)
- [Displaying the version information for Backup Exec](#)
- [How to sort, filter, and copy information on the Backup Exec Administration Console](#)
- [Customizing views on the Backup Exec Administration Console](#)
- [Configuring the Home tab](#)
- [Configuring the Symantec RSS Reader](#)
- [Creating a disaster preparation plan \(DPP\)](#)
- [Getting started with backups](#)

About the Backup Exec Administration Console

Use the Backup Exec Administration Console to run backups, restore data, monitor jobs, configure storage, and run reports.

Figure 3-1 Backup Exec Administration Console



The administration console screen includes the following components:

Table 3-1 Administration console components

Item	Description
Backup Exec button	The Backup Exec button displays on the upper left side of the administration console. To display the options in the Backup Exec button, click the Backup Exec button, select the menu name, and then select an option. You can launch Backup Exec operations by clicking options from a menu.

Table 3-1 Administration console components (*continued*)

Item	Description
Tabs	<p>Tabs at the top of the screen let you navigate Backup Exec.</p> <p>You can access the following views from the navigation bar:</p> <ul style="list-style-type: none"> ■ Home. Provides quick access to the Backup Exec information that you use frequently. Customize the Home view by adding or deleting items. ■ Backup and Restore. Create a backup or restore job. ■ Job Monitor. Monitor and manage backup, restore, installation, and storage operation jobs. ■ Storage. Configure storage, run storage operations, and manage media. ■ Reports. View, print, save, and schedule reports about the Backup Exec server, operations, and device and media usage, and create custom reports. View reports in Backup Exec in PDF or HTML format, and save and print reports in PDF, XML, HTML, Microsoft Excel (XLS), and comma-separated value (CSV) formats
Groups	<p>Groups display on the tabs in the administration console, and contain the commands that initiate actions such as creating a new backup job or configuring storage. The commands in the groups are dynamic, changing according to the selection. Some commands are unavailable until you select an item on the console screen or until you run a prerequisite task.</p>
Columns	<p>Customize columns by doing any of the following actions:</p> <ul style="list-style-type: none"> ■ Drag and drop columns to change their location. ■ Right-click a column heading to select the columns to display or to sort and filter the column content. ■ Click the column heading to change the order of the column. <p>For example, names of reports display in alphabetical order by default. To display report names in reverse alphabetical order, click the Name column heading on the Reports view.</p>
Selection pane	<p>Select items to work with, such as servers to back up or restore.</p>

Table 3-1 Administration console components (continued)

Item	Description
Details pane	<p>Additional details appear for the server that you select in the list of servers. The Details pane lists the resources for the selected server and the selection status, backup status, and logon account information for each resource.</p> <p>When you select a Hyper-V host or VMware host in the list of servers, the Details pane includes the following details:</p> <ul style="list-style-type: none">■ The last 7 days of backup jobs■ The date of the last backup■ The date of the next scheduled backup <p>You can also restore data and filter the list of guest virtual machines from this pane.</p>
Status bar	<p>The status bar appears on the bottom of the administration console and provides information about the Backup Exec server, jobs that are running or are scheduled, alerts, and services.</p>
Refresh	<p>Click F5 to refresh the user interface on the administration console.</p>

Displaying the version information for Backup Exec

You can display information about the version of Backup Exec that is installed.

To display the version information for Backup Exec

- 1 Click the Backup Exec button, select **Help and Documentation**, and then click **About Backup Exec**.
- 2 Click **OK**.

How to sort, filter, and copy information on the Backup Exec Administration Console

You can customize the information that displays on the **Backup and Restore** tab, the **Job Monitor** tab, and the **Storage** tab.

You can do any or all of the following actions:

- Choose a default configuration that Backup Exec provides, such as **Servers with Active Jobs** or **Failed Jobs**.
- Specify a sort order for the columns that appear in the views.

- Specify the values that you want to use to filter the information that Backup Exec displays.
- Specify the columns that you want to appear and the order in which they should appear.
- Create and save a configuration to use again.
- Copy list items to the Clipboard and then copy them to any application that supports copy-and-paste.

To sort or filter information on the Backup Exec Administration Console

1 Do one of the following:

- | | |
|---|---|
| To customize a view of the computers in the list of servers | On the Backup and Restore tab, in the Views group, click Sort and Filter . |
| To customize a view of the storage devices | On the Storage tab, in the Views Group, click List , and then click Sort and Filter . |

2 Do any of the following:

- | | |
|--|--|
| To select a default configuration, such as Servers with Failed Backups , or to select a configuration that you created and saved previously | Click Configurations and select a configuration. |
| To specify an ascending or descending sort order for the columns | Click Sort , choose the options as appropriate, and then click OK . |
| To specify one or more columns to filter for specific values | Click Filter , choose the options as appropriate, and then click OK . |
| To specify the columns that you want to display and the order in which they should appear | Click Columns , choose the options as appropriate, and then click OK . |
| To create and save a configuration | Click Save , choose the options as appropriate, and then click OK . |

To edit a configured view

- 1 Do one of the following:

To edit a configuration from the Backup and Restore tab	On the Backup and Restore tab, in the Views group, click Sort and Filter .
--	---

To edit a configuration from the Storage tab	On the Storage tab, in the Views group, click Sort and Filter .
---	--

- 2 Click **Configurations**.

- 3 Select the configuration that you want to edit, and then click the pencil icon.

To delete a configured view

- 1 Do one of the following:

To delete a configuration from the Backup and Restore tab	On the Backup and Restore tab, in the Views group, click Sort and Filter .
--	---

To delete a configuration from the Storage tab	On the Storage tab, in the Views group, click Sort and Filter .
---	--

- 2 Click **Configurations**.

- 3 Select the configuration that you want to delete, and then click the delete icon.

To copy information on the Backup Exec Administration Console

- 1 On any tab except the **Home** tab, right-click an item in the list view.

- 2 Click **Copy**.

- 3 Open any application that supports copy-and-paste, and then paste the information.

Customizing views on the Backup Exec Administration Console

You can customize how you view the information that displays on the **Backup and Restore** tab, the **Job Monitor** tab, and the **Storage** tab.

Table 3-2 Views on the Backup Exec Administration Console

View	Description
Standard	Displays the information in a view that provides descriptive text.
Compact	Displays the information in a view that conserves space.
Tree	Displays the items in a hierarchical view. This view is disabled for the list of servers on the Backup and Restore tab.
List	Displays the items in a list that you can sort by columns. This view is disabled for the list of servers on the Backup and Restore tab.

To customize views on the Backup Exec Administration Console

- ◆ On the **Backup and Restore** tab, the **Job Monitor** tab, or the **Storage** tab, in the **Views** group, click **Standard**, **Compact**, **Tree**, or **List**.

Note: The **Tree** and **List** view is disabled for the list of servers on the **Backup and Restore** tab.

Configuring the Home tab

You can configure the **Home** tab by selecting the items that you want to display. You can drag and drop items to move them to another location on the **Home** tab or you can maximize a single item. The **Home** tab items contain Backup Exec data and links to features that you use frequently.

To configure the Home tab

- 1 On the **Home** tab, in the **Layout** group, click the layout for the items that you want to display.
- 2 In the **System Health** and **Support** groups, select the check box for the items that you want to display.
- 3 Drag the items to a column and position in which you want them to display to further customize the **Home** tab.

You can configure the **Layout**, and hide or display items in the **System Health** and the **Support** groups.

Restoring the Home tab's default layout

You can quickly restore the **Home** tab to its default configuration at any time.

To restore the Home tab's default configuration

- ◆ On the **Home** tab, in the **Layout** group, click **Reset Home Tab**.

Layout group

You can select one of the following layout configurations to display the items on the **Home** tab.

Table 3-3 Home Tab Layout items

Item	Description
One Column	Displays the Home tab items in one column.
Two Columns	Displays the Home tab items in two columns.
Narrow/Wide	Displays the Home tab items in two columns with a narrow panel and a wide panel.
Three Columns	Displays the Home tab items in three columns.
Reset Home Tab	Restores the contents of the Home tab to the default configuration.

System Health group

The items in the **System Health** group provide overviews of alerts, backup jobs, backup size data, storage status, the Symantec ThreatCon Level, and Simplified Disaster Recovery. You can select the following items to display on the Backup Exec **Home** tab.

Table 3-4 System Health group items

Item	Description
Active Alerts	<p>Lets you view all alerts that have not received a response. You can filter the alerts to view specific types of alerts, the source of the alerts, and the amount of time that alerts occurred.</p> <p>You can display any or all of the following types of alerts:</p> <ul style="list-style-type: none">■ Error■ Warning■ Attention Required■ Information
Alert History	<p>Lets you view the property and response information for alerts.</p>
Backup Status	<p>Provides a summary view of the backup job status for the servers that are backed up or available for backup.</p>
Backup Size	<p>Provides a summary view of the amount of data that is backed up. You can customize the number of days for which you display information about the backup size. You can also select the type of backups that display.</p>
Storage Status	<p>Provides a summary view of the amount of space that is available on your storage. The storage information includes the total capacity that displays the amount space that is used for the different types of data.</p>
Symantec ThreatCon Level	<p>Provides an overall view of global Internet security level. Symantec's ThreatCon levels are based on a 1-4 rating system, with level 4 being the highest threat level.</p> <p>You must have Symantec Endpoint Protection 11.0 or later installed on the same computer as Backup Exec to view this item.</p>

Table 3-4 System Health group items (*continued*)

Item	Description
Simplified Disaster Recovery	<p>Provides a status of whether the Simplified Disaster Recovery disk image has been created.</p> <p>If the Simplified Disaster Recovery disk image has not been created, you can click the Create Disaster Recovery link to launch a wizard that guides you through the process.</p> <p>See “About Simplified Disaster Recovery” on page 736.</p>

Support group

The items in the **Support** group provide technical support, documentation, licensing and maintenance contracts, and the Symantec RSS Reader resources. You can select the following items to display on the Backup Exec **Home** tab.

Table 3-5 Support group items

Item	Description
Technical Support	<p>Provides the following support options to help you understand product features and functionality or troubleshoot issues:</p> <ul style="list-style-type: none">■ Backup Exec Tech Center■ Backup Exec Technical Support■ Best Practices■ Use MySupport to manage new or existing support cases■ Symantec Remote Assistance■ Register to receive notifications■ Get Backup Exec updates
Documentation	<p>Provides the following documentation options to help you understand product features and functionality or troubleshoot issues:</p> <ul style="list-style-type: none">■ View Readme■ View Administrator's Guide (PDF)■ View Administrator's Guide Addendum (PDF)

Table 3-5 Support group items (*continued*)

Item	Description
Licensing and Maintenance	Provides the following licensing and maintenance options to help you manage maintenance contracts and licenses: <ul style="list-style-type: none">■ View license information■ View maintenance contract information
Symantec RSS Reader	Lets you view and add Backup Exec and Symantec RSS feeds.
Installation and Upgrades	Provides the following reports: <ul style="list-style-type: none">■ Migration report■ Installation Summary report This item only displays when you upgrade Backup Exec to a later version. See “Using the Migration Report to determine how existing jobs changed after an upgrade to Backup Exec 2014” on page 129. See “Viewing the Installation Summary Report” on page 119.

Configuring the Symantec RSS Reader

You can customize the Symantec RSS Reader and select the default Backup Exec feeds that display in the reader. You can add additional Symantec and Backup Exec RSS feeds or remove RSS feeds.

The Symantec RSS Reader sorts articles by the date and the time. The reader displays the last entry of an article in the RSS feed; however, you can choose to view the full article.

The Symantec RSS Reader refreshes the RSS feeds every 15 minutes when the item is open in the **Home** tab. If the RSS feed is not open in the reader, the RSS feed does not refresh.

To view an article in the Symantec RSS Reader

- 1 On the **Home** tab, in the **Support** group, select the **Symantec RSS Reader** check box.
- 2 In the **Symantec RSS Reader**, click the arrow next to the RSS feed that contains the article.

- 3 Click the hyperlink for the article that you want to open.

The Symantec RSS Reader opens a new window that contains a portion of the article from the RSS feed.

- 4 Click **Go to full Article** to open Internet Explorer and view the entire contents of the article.

To customize the Symantec RSS feeds to the Symantec RSS Reader

- 1 On the **Home** tab, in the **Support** group, select the **Symantec RSS Reader** check box.
- 2 In the **Symantec RSS Reader**, click the pencil icon to add an RSS feed.
- 3 Enter any of the following:

URL	Indicates the location of the RSS feed that you want to add to the Symantec RSS Reader.
Name	Indicates the name of the RSS feed that you want to display in the Symantec RSS Reader.
Click here to see more Symantec RSS feeds	Shows a list of Symantec RSS feeds that you can add to the Symantec RSS Reader.

- 4 Click **OK**.

To remove an RSS feed from the Symantec RSS Reader

- 1 On the tab, in the **Support** group, select the **Symantec RSS Reader** check box.
- 2 Do one of the following:

To remove a default Backup Exec RSS feed	Clear the check box of the Backup Exec RSS feed.
To remove an RSS feed that you added to the Symantec RSS Reader	Click the red X next to the name of the RSS feed.

Creating a disaster preparation plan (DPP)

Disaster preparation planning is the implementation of strategies and procedures that minimize damage in the event a catastrophe destroys your data.

The following basic methods are available for disaster recovery:

- Automated recovery. Backup Exec's Simplified Disaster Recovery (SDR) option automates the disaster recovery process for Windows computers.
- Manual recovery. You can manually recover both local and remote Windows computers.

The purpose of a Disaster Preparation Plan (DPP) is to return to an operational status as quickly as possible. Backup Exec is a crucial component of the DPP. The DPP you put in place with your Backup Exec system should be customized to your network environment.

For more information about how to perform a manual disaster recovery, click the following URL:

<http://entsupport.symantec.com/umi/V-290-706>

While environments vary in different organizations, consider the following elements when creating a comprehensive DPP.

Table 3-6 Key elements of a DPP

Element	Description
Hardware protection	The hardware devices on your network (CPUs, drives, video) are susceptible to damage from many disaster situations. Uninterruptible power supplies (UPS), surge protectors, and security monitoring devices are the equipment most often used today to protect hardware. If you do not already have these items in place, you should consider installing them. The initial investment could be justified many times over in the event of a disaster.
The ability to maintain business operations during a disaster period	Make sure that proper precautions are taken by everyone to implement plans for network interruptions. For example, the phones in the sales department won't stop ringing because the server is down, so orders may have to be handwritten until the server is up again. Each department should work out strategies for such occurrences. If the proper precautions are taken, the server can be rebuilt quickly and operations can still continue.
A sound backup strategy.	A well-designed backup strategy that includes a strong media rotation scheme plays a key role in quickly restoring your file server.

Table 3-6 Key elements of a DPP (*continued*)

Element	Description
Off-site and duplicate stage backups.	It is imperative that you regularly move the backed-up data to an off-site storage facility. If you use disk as your storage medium, consider adding a stage to duplicate backups to other storage. This ensures that if something happens to your facility, all of your backups are not destroyed. Depending on the importance of your data, you may choose to use several off-site storage facilities.
Effective DPP management	A person or group of people should constantly supervise your organization's disaster preparation efforts. This person or group should install and maintain hardware protection devices, make sure all departments have a plan if the server goes down temporarily, and make sure that backups are made and rotated off-site regularly. Document your Disaster Preparation Plan for reference purposes.

See [“Getting started with backups”](#) on page 146.

Getting started with backups

After you install Backup Exec, you can run a backup job. The following table describes the recommended process for getting started.

Table 3-7 Getting started with backups

Step	Description
1. Configure storage	<p>You must have a configured storage device before you can run any backup jobs.</p> <p>If no storage devices are already configured, such as tape drives or robotic libraries, then you can use the Configure Storage wizard to set up storage devices in Backup Exec. The wizard guides you through configuring all of the storage that Backup Exec supports.</p> <p>To start the wizard, on the Storage tab, in the Configure group, click Configure Storage. The wizard guides you through the rest of the process.</p> <p>See “Using the Configure Storage wizard” on page 428.</p>

Table 3-7 Getting started with backups (*continued*)

Step	Description
2. Configure Backup Exec logon accounts	<p>You can use the default Backup Exec logon account, which is the system logon account for the Backup Exec server.</p> <p>Backup Exec uses the system logon account by default for most backups. The system logon account contains the credentials of the Backup Exec service account. If the service account does not have rights to access the data that you select for backup, you can use the Logon Account Wizard. Use this wizard to create additional logon accounts that do contain the necessary credentials for accessing that data.</p> <p>To start the Logon Account Wizard, click the Backup Exec button. Select Configuration and Settings > Logon Accounts > Logon Account Wizard.</p> <p>See “Backup Exec logon accounts” on page 610.</p>
3. Run a backup job	<p>You can back up any of the computers on which you installed the appropriate Backup Exec agent, such the Agent for Windows, or the Agent for Mac.</p> <p>If you didn't install a Backup Exec agent on a computer, you can do so now. On the Backup and Restore tab, in the Servers group, click Add, and then click the appropriate selection. The wizard guides you through the rest of the process.</p> <p>To start a backup job, on the Backup and Restore tab, select the computer that you want to back up and then in the Backups group, click Backup. Click a menu item, such as Back Up to Disk. You can click Edit to change any of the defaults, or you can click OK to accept all of the defaults and let Backup Exec schedule the job.</p> <p>See “Backing up data” on page 158.</p>

Table 3-7 Getting started with backups (*continued*)

Step	Description
4. Create the Simplified Disaster Recovery disk image	<p>By default, Backup Exec selects all the data on a computer for backup, including the critical system components that you need to perform a full system restore by using SDR. SDR-enabled backups are those backups for which all of the critical system components are selected for backup. You must have SDR-enabled backups to use Backup Exec to rebuild a computer and restore it to a functional state.</p> <p>See “How to ensure that backups are enabled for Simplified Disaster Recovery” on page 743.</p> <p>After you have run SDR-enabled backups for all of the computers that you want to protect, you should create a Simplified Disaster Recovery (SDR) disk image.</p> <p>On the Backup and Restore tab, in the Restores group, click Create Disaster Recovery Disk. The wizard guides you through the rest of the process.</p> <p>See “Creating a Simplified Disaster Recovery disk image” on page 750.</p>

Backups

This chapter includes the following topics:

- [How to prepare for your first backup with Backup Exec](#)
- [Required user rights for backup jobs](#)
- [About the list of servers on the Backup and Restore tab](#)
- [Adding servers that you want to back up to the list of servers on the Backup and Restore tab](#)
- [Removing servers from the list of servers on the Backup and Restore tab](#)
- [Creating a server group](#)
- [Hiding or viewing server groups on the Backup and Restore tab](#)
- [Editing a server group](#)
- [Retiring servers from Backup Exec](#)
- [Removing server groups from Backup Exec](#)
- [Backing up data](#)
- [About selecting data to back up](#)
- [Changing the order in which backup sources are backed up](#)
- [Excluding files from backups](#)
- [Including specific files with a backup job's backup selections](#)
- [About backing up critical system components](#)
- [Backup Exec Shadow Copy Components file system](#)
- [Backup methods in Backup Exec](#)

- Configuring backup methods for backup jobs
- How Backup Exec determines if a file has been backed up
- Configuring Backup Exec to automatically delete files after a backup
- Configuring network options for backup jobs
- Running the next scheduled backup job before its scheduled time
- Editing backup definitions
- Testing credentials for backup sources
- Replacing the credentials for a backup source
- Creating new credentials for a backup source
- Deleting retired or unused backup sources from the Credentials pane
- How job scheduling works in Backup Exec
- Including a specific date in the schedule for a backup job
- Preventing backup jobs from running on a specific date
- Viewing all scheduled backup jobs on a calendar
- Adding a stage to a backup definition
- Editing a stage
- Duplicating backup sets or a job history manually
- Running a test run job manually
- Verifying backed up data manually
- Copying data from a virtual tape library to a physical tape device using DirectCopy to tape

How to prepare for your first backup with Backup Exec

Before you back up data, you should develop a backup strategy that includes the backup method, frequency, and data retention methods that are appropriate for your organization. A backup strategy is the collection of procedures you implement as a solution for backing up your environment. You may have different strategies for different areas of the organization.

You may need to analyze your backup environment to determine the following:

- How much data needs to be backed up?
- How long will the backups take?
- How much storage is required?

For more information, go to the following URL:

<http://entsupport.symantec.com/umi/V-269-38>

Before you run a backup job, you should ensure that you have the proper user rights.

See [“Required user rights for backup jobs”](#) on page 151.

You must also configure storage before creating backup jobs. You can set up Backup Exec to use specific storage devices or logical groupings of devices, such as storage pools.

Specifically, you might want to perform the following tasks to help you manage storage hardware and media most effectively:

- Create disk-based storage so that Backup Exec can automatically manage backup data retention.
See [“Configuring disk storage”](#) on page 318.
- Set up storage device pools to load-balance jobs.
See [“Creating storage device pools”](#) on page 416.
- Create media sets to manage data retention for tape cartridge media.
See [“Default media sets”](#) on page 379.
- Configure deduplication disk storage to optimize storage and network bandwidth.
See [“About the Deduplication Option”](#) on page 795.
- Configure vault stores for the Archiving Option.
See [“About vault stores in the Archiving Option”](#) on page 1351.

Required user rights for backup jobs

To perform any backup operations, the following Windows user rights are required for the service account and any Backup Exec logon accounts:

- Act as part of the operating system.
- Create a token object.
- Back up files and directories.
- Restore files and directories.

- Manage auditing and security log.
- Logon as a batch job (only for Windows Vista and later).

For more information about user rights in Windows operating systems, see your Microsoft documentation.

See [“Changing the credentials for a service account”](#) on page 621.

See [“Backup Exec logon accounts”](#) on page 610.

About the list of servers on the Backup and Restore tab

You can view a list of servers on the **Backup and Restore** tab. The servers that display in the list include any servers that Backup Exec discovered during an upgrade, any servers that you manually add to Backup Exec, and any servers that Backup Exec discovers during a catalog operation. Servers must be added to the list so that you can select them for backup jobs.

You can also monitor server activity and job status from the list of servers. By default, Backup Exec displays a server's alerts, backup status, and a calendar of the last seven days of backup jobs. It also displays the date and time of the previous and upcoming scheduled backups. You can customize the columns on this list to display additional information.

You can select to view any of the following details about each server in the list:

- Server
- Server type
- Version of the server
- Backup Exec version
- Data source types
- Backup selections
- Active alerts
- Status
- Last seven days of backup jobs
- Last backup
- Next backup
- Percent complete

- Elapsed time
- Byte count
- Average job rate
- Description

Windows servers must have the Agent for Windows installed on them before you can add them to the list of servers. When you add Windows servers to Backup Exec, you have the option to install the Agent for Windows to them remotely.

If you no longer want to monitor or back up a server with Backup Exec, you can remove it from the list of servers.

See [“Adding servers that you want to back up to the list of servers on the Backup and Restore tab”](#) on page 153.

See [“Removing servers from the list of servers on the Backup and Restore tab”](#) on page 154.

Adding servers that you want to back up to the list of servers on the Backup and Restore tab

Before you can create a backup definition, you must add the servers that you want to protect to the list of servers on the **Backup and Restore** tab. You can add servers during the push-installation process or you can complete the following procedure to add servers at any time.

To add servers to the list of servers

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Complete the steps to add a server or servers to the list of servers.

Note: If Backup Exec discovered servers using the **Discover Data to Back Up** option, they display on the **Browse** dialog box under the heading **Servers without an Agent for Windows installed**.

See [“Adding discovered servers to the list of servers in Backup Exec”](#) on page 573.

See [“About the list of servers on the Backup and Restore tab”](#) on page 152.

Removing servers from the list of servers on the Backup and Restore tab

If you no longer want to monitor or back up a server with Backup Exec, you can remove it from the list of servers on the Backup and Restore tab. You can no longer back up or restore data from servers after you remove them from the list of servers.

Note: If you remove a server from the list and it has scheduled jobs pending, the jobs are deleted. The jobs do not run as scheduled. Do not remove a server from the list of servers if you still want to back up that server.

You should not use this procedure to delete a managed Backup Exec server from a CASO environment. You should remove managed Backup Exec servers using the **Storage** tab.

See [“Removing a managed Backup Exec server from a Backup Exec server pool”](#) on page 1139.

To remove servers from the list of servers

- 1 On the **Backup and Restore** tab, right-click the server that you want to remove from the list of servers.
- 2 Click **Remove**.
- 3 Click **Yes** to confirm that you want to remove the server from the list of servers.

See [“About the list of servers on the Backup and Restore tab”](#) on page 152.

Creating a server group

Server groups are a way to organize and view server information in the list of servers. You can create server groups based on any criteria. You may want to group servers with a specific type of data or servers that reside in a specific location. Then, when you view server groups, only the server group that you select displays in the list of servers on the **Backup and Restore** tab. Viewing server groups lets you quickly monitor the status of all of the servers in the group at a glance. You can also back up an entire server group.

See [“Backing up data”](#) on page 158.

Backup Exec comes with two preconfigured server groups. The All Servers server group contains all of the servers in the list of servers. The Retired Servers server group is intended for any servers that you no longer actively monitor with Backup Exec. Servers no longer appear in the All Servers server group after you move them to the Retired Servers server group.

To create a server group

- 1 On the **Backup and Restore** tab, in the **Views** group, click **Groups**.
- 2 Click **Add**.
- 3 In the **Group name** field, type a name for the server group that you want to create. You may want to name the server group to indicate the type of servers in the group or the location at which the servers reside, for example.
- 4 In the **Description** field, type a description for the server group.
- 5 (Optional) Do any of the following to filter the servers in your environment so that you can find the servers that you want to add to the server group:

To filter servers by type	In the Server type field, select the type of server for which you want to search.
To filter servers by data type	In the Data type field, select the type of data that the server for which you want to search contains.
To filter servers by name	In the Name contains field, type all or part of the server name.

- 6 In the **Servers** group box, select the servers that you want to add to the server group, and then click **OK**.

See [“Hiding or viewing server groups on the Backup and Restore tab”](#) on page 155.

See [“Editing a server group”](#) on page 156.

See [“Retiring servers from Backup Exec”](#) on page 157.

See [“Removing server groups from Backup Exec”](#) on page 158.

Hiding or viewing server groups on the Backup and Restore tab

You view server groups on the **Groups** pane. The **Groups** pane is enabled by default when you install Backup Exec. If you do not use server groups, you can hide the **Groups** pane. Double-click a server group on the **Groups** pane to view more detailed information about the server group's jobs, job history, and any active alerts.

To hide or view server groups on the Backup and Restore tab

- 1 On the **Backup and Restore** tab, in the **Views** group, click **Groups**.
- 2 Select **Show Server Groups**.

The **Show Server Groups** option lets you hide or show the **Groups** pane to the left of the list of servers.

See [“Creating a server group”](#) on page 154.

Editing a server group

Existing server groups can be edited in the following ways:

- Add servers to the server group.
- Remove servers from the server group.
- Change the server group's name.
- Change the server group's description.

To edit a server group

- 1 On the **Backup and Restore** tab, in the **Groups** pane, right-click the group that you want to edit.
- 2 Select **Edit**.
- 3 Do any of the following, as necessary:

To add servers to the server group

In the **Servers** group box, select the servers that you want to add to the server group.

You can use the following fields to filter the list of servers so that you can find the server that you want to add:

- **Server type**
- **Data type**
- **Name contains**

To remove servers from the server group

In the **Servers** group box, deselect the servers that you want to remove from the server group.

You can use the following fields to filter the list of servers so that you can find the server that you want to remove:

- **Server type**
- **Data type**
- **Name contains**

To change the server group's name In the **Group name** field, type the new name for the server group.

To change the server group's description In the **Description** field, type the new description for the server group.

4 Click **OK**.

See [“Creating a server group”](#) on page 154.

Retiring servers from Backup Exec

You can retire servers from Backup Exec by moving them to the Retired Servers server group. The Retired Servers server group is intended for any servers that you no longer actively monitor with Backup Exec. You can still view any information about the retired servers on the **Backup and Restore** tab by clicking on the Retired Servers server group. However, the retired servers no longer appear in the All Servers server group with the servers that you regularly back up and monitor. It may be useful to retire servers if you use Backup Exec to monitor a large number of servers.

You cannot create new backup jobs for any servers that are in the Retired Servers server group. However, any scheduled backup jobs still run on retired servers. You can restore data from retired servers.

If you retire a server and then you decide that you want to move it back to the All Servers server group, you can click on it and drag it from the Retired Servers server group to the All Servers server group.

To retire servers from server groups

- 1 On the **Backup and Restore** tab, in the **Groups** pane, right-click **Retired Servers**.
- 2 Click **Edit**.

- 3 (Optional) Do any of the following to filter the servers in your environment so that you can find the servers that you want to retire:

To filter servers by type	In the Server type field, select the type of server for which you want to search.
To filter servers by data type	In the Data type field, select the type of data that the server for which you want to search contains.
To filter servers by name	In the Name contains field, type all or part of the server name.

- 4 In the **Servers** group box, select the servers that you want to retire and then click **OK**.

See [“Creating a server group”](#) on page 154.

Removing server groups from Backup Exec

If you no longer want to use a server group, you can remove it from Backup Exec. Removing a server group does not affect the servers in the group. It only removes the association between the servers and your ability to see them in the group. You can still back up and monitor servers after you remove the server group to which they belong. You cannot remove either the All Servers server group or the Retired Servers server group.

To remove server groups

- 1 On the **Backup and Restore** tab, in the **Groups** pane, right-click the server group that you want to delete.
- 2 Click **Remove**.
- 3 Confirm that you want to remove the server group.

See [“Creating a server group”](#) on page 154.

Backing up data

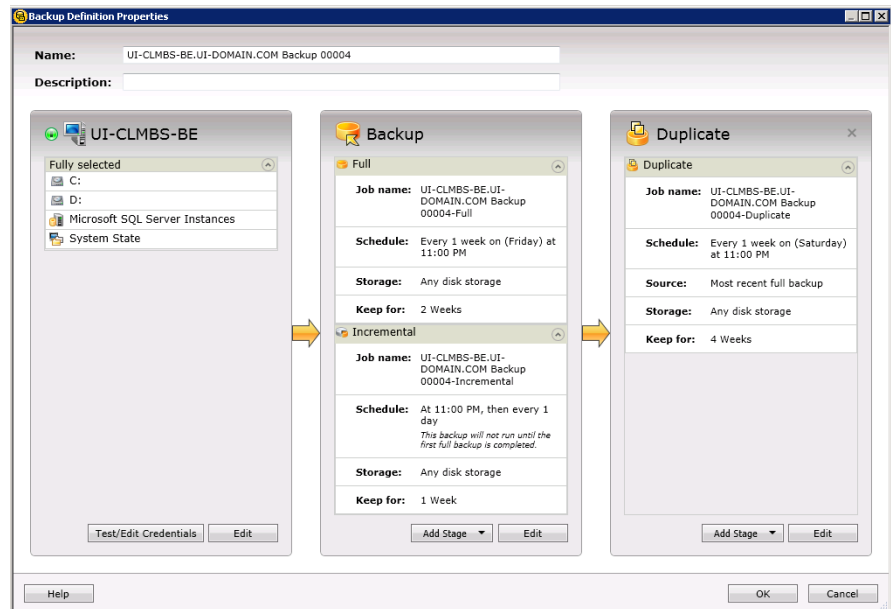
When you want to back up data, you create a container that is called the backup definition.

Backup definitions contain backup selections, job templates, and stages.

Table 4-1 Backup definition contents

Item	Description
Backup selections	Backup selections include any servers, volumes, or data that you have selected to back up.
Job templates	<p>Job templates are the collection of settings that Backup Exec uses to create jobs. Backup job settings can include scheduling options, storage device options, or backup methods for selected types of data, for example. When you want to run a backup, Backup Exec combines the job template with the backup selections to create a backup job that runs according to the options that you specified.</p> <p>Backup definitions always contain one job template that uses the full backup method, but recurring jobs can also contain incremental, differential, or data-specific job templates.</p> <p>See “Backup methods in Backup Exec” on page 188.</p>
Stages	Stages are optional tasks that you can run with backup jobs. Backup definitions can contain one or more stages. You can create stages that duplicate your backup data, create virtual machines with your backup data, or archive your backup data.

Figure 4-1 Backup Definition (with backup selections, full and incremental backup job templates, and a duplicate stage)



Backup Exec offers many choices for creating backup jobs to protect your data, including the following:

- Create a recurring backup definition to back up the full or partial contents of a single server or multiple servers
If you select to back up multiple servers or applications, you can create separate backup definitions for each server or application. Alternatively, you can create one backup definition that includes all servers and applications. If you create separate backup definitions, it is easier to identify problems when backup jobs fail. Also, if an issue arises with one server that causes a backup job to fail, the other backup jobs can still complete successfully. If you create one backup definition that includes all of your servers and applications, it takes less work to monitor the job status. You can select a default method for backing up multiple servers in the **Backups** dialog box in the Backup Exec Settings.
See [“Configuring the default setting for backing up multiple servers or applications”](#) on page 566.

Note: You cannot back up multiple servers with a synthetic backup job or a conversion to virtual job.

- Create a backup definition to run only once
A one-time backup only runs once without any recurring instances. You may want to use a one-time backup to create a baseline for a server before you upgrade it or install new software. After Backup Exec finishes running a one-time backup, it deletes the job rather than saving it with your recurring jobs. If you want to view information about a one-time backup after the job is complete, you can still view its job history.

Warning: Data lifecycle management deletes all expired backup sets that are created by a one-time backup job. DLM does not keep the last backup set after the retention date expires if the backup set is from a one-time backup. To prevent the backup sets from being automatically deleted, you can manually retain specific backup sets or you can change the expiration date of the backup set.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 335.

- Create a new backup definition using an existing backup definition's settings
If you want to create a backup definition that is similar to an existing backup definition, you can apply the existing definition's settings to a new definition. Any backup methods, job settings, and stages are copied into a new backup definition for the server or servers that you selected to back up. All that you have to do is select the backup selections. You can override any of the job settings, if necessary.
- Create a server group out of similar computers and back up the entire group at once
Server groups are a way to organize and view server information in the list of servers. You can create server groups based on any criteria. You may want to group servers with a specific type of data or servers that reside in a specific location.

You can also edit existing backup definitions to modify their schedules, backup selections, or other settings.

To protect remote computers, you must install the Agent for Windows on the remote computer. The Agent for Windows is a system service that runs on Windows servers and workstations. It provides efficient backup processing by locally performing the tasks that, in typical backup technologies, require extensive network interaction.

See [“Methods for installing the Agent for Windows”](#) on page 86.

To back up data

- 1 Select one of the following methods to back up data:

To create recurring backup jobs

Complete the following steps:

- On the **Backup and Restore** tab, select the server, servers, or server group that you want to back up.
- Right-click the server, servers, or server group.
- On the **Backup** menu, select the backup option that you want to use.

To create one-time backups

Complete the following steps:

- On the **Backup and Restore** tab, select the server, servers, or server group that you want to back up.
- In the **Backups** group, click **One-Time Backup**.
- Select the backup option that you want to use.

To create a new backup definition using the settings from an existing backup definition

Complete the following steps:

- On the **Backup and Restore** tab, right-click the server or servers that you want to back up.
You cannot reuse a backup definition's settings to create a backup job for a server group.
- On the **Backup** menu, select **Create a New Backup Using the Settings from an Existing Backup**.
- On the **Backup Job Selection** dialog box, select the backup definition that contains the settings that you want to copy.
- Click **OK**.

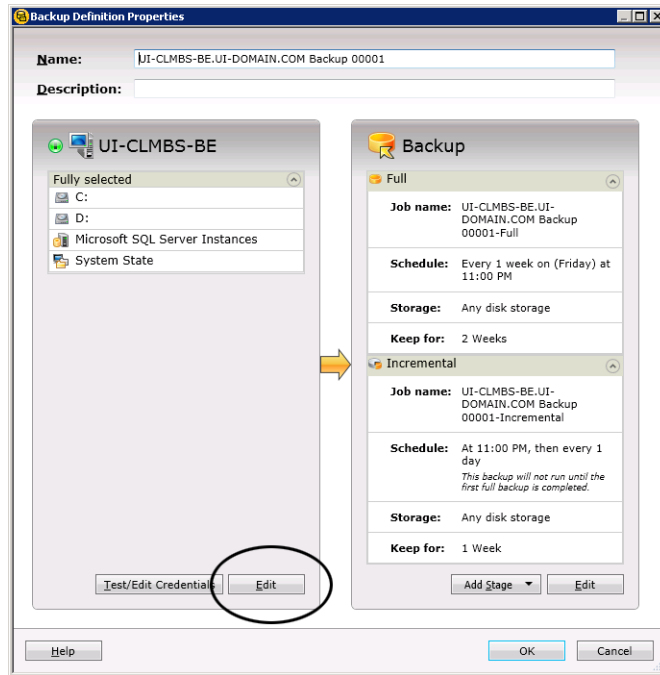
To create a backup definition for a server group

Complete the following steps:

- On the **Backup and Restore** tab, in the **Groups** pane, right-click the server group that you want to back up.
- On the **Backup** menu, select the backup option that you want to use.

- 2 In the **Name** field, type a name for the backup definition.
- 3 In the **Description** field, type a description for the backup definition.

- 4 In the **Selections** box, click **Edit**.



- 5 To add additional servers to the backup definition:
- Click **Add**.
 - Select the server or servers that you want to add to the backup definition.
 - Click **OK**.
- 6 Select the data that you want to back up.

Servers are selected in their entirety by default. If you do not want to back up the entire server, double-click the server name to see all of the contents. Select the check boxes for each item that you want to back up.

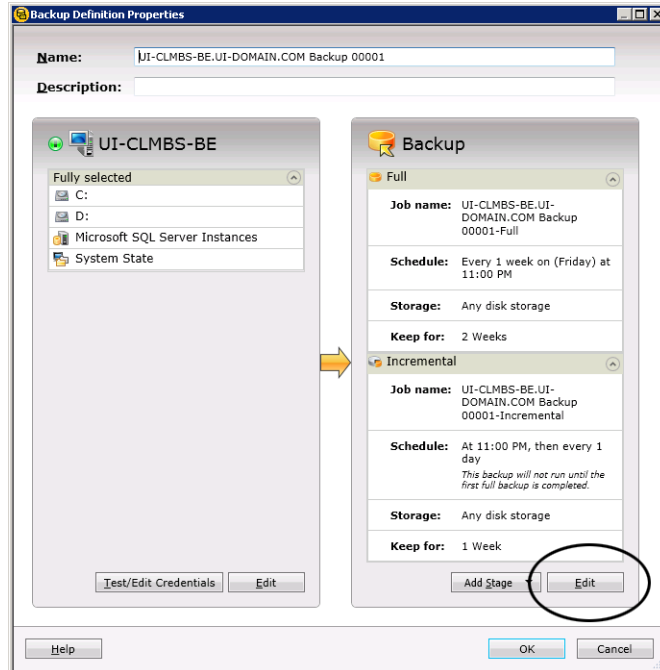
Note: Deselecting a server's critical system components creates backup sets that cannot be used for some restore scenarios.

See [“About backing up critical system components”](#) on page 183.

- 7 Click **OK**.

- 8 In the **Backup** box, click **Edit**.

Note: If you copied the backup settings from an existing backup definition or if you do not want to change any of the existing or default settings, you can skip to step 13



- 9 In the left pane, click **Schedule**, and then select when you want the backup job or jobs to run.
- 10 In the left pane, click **Storage**, and then select the storage device that you want to use for the backup job or jobs.

- 11 In the left pane, select any additional options that apply to the backup job or jobs.

The remaining options in the left pane are optional. The options vary depending on what you selected to back up.

For example, you might want to set up notifications so that users can receive email or texts regarding this backup definition.

Network

Select this option to specify the network interface that Backup Exec uses to access remote computers.

See [“Configuring network options for backup jobs”](#) on page 204.

Notification

Select this option to configure Backup Exec to notify specified recipients when the backup job is completed.

Each backup job and stage can be configured with different notification recipients. Backup Exec can notify people by email or text message.

See [“Notification options for jobs”](#) on page 305.

Test Run

Select this option to configure a test job that automatically tests storage capacity, credentials, and media integrity.

The test job can help you determine if there are any problems that might keep the backup job from completing successfully.

See [“Configuring automatic test run jobs for backup jobs”](#) on page 533.

Verify

Select this option to create a job that automatically verifies whether all of the data was successfully backed up when the job is completed.

A verify job can also help you determine whether the media you use is defective.

See [“Configuring automatic verify operations for backup jobs”](#) on page 534.

Delayed Catalog

Select this option to configure the delayed catalog operation for any GRT-enabled jobs. You can choose to run a delayed catalog operation, schedule the catalog operation for another time, or choose to not run a delayed catalog operation.

See [“Configuring delayed catalog operations to improve backup performance”](#) on page 536.

Advanced Open File

Select this option to configure the snapshot settings that Backup Exec uses to process the backup job. Snapshot technology lets Backup Exec capture any files that are open when a backup job runs.

You can also enable checkpoint restart, which lets you resume interrupted backup jobs.

See [“Configuring Advanced Open File options for backup jobs”](#) on page 540.

See [“Configuring checkpoint restart”](#) on page 543.

Advanced Disk-based Backup

Select this option to configure off-host backup processing for the backup job.

See [“Configuring off-host backup options for a backup job”](#) on page 1157.

Security

Select this option to configure Backup Exec to work with Symantec Endpoint Protection. You can configure the backup job to run automatically when the Symantec ThreatCon reaches a specified level. Integrating your backup jobs with Symantec Endpoint Protection helps make sure that your vital data is safely backed up as soon as global threats are detected.

See [“Configuring Backup Exec to work with Symantec Endpoint Protection”](#) on page 545.

Pre/Post Commands

Select this option to configure any commands that you want to run either before the backup job begins or after the backup job is completed.

See [“Configuring pre/post commands for backup jobs”](#) on page 547.

Files and Folders

Select this option to configure how Backup Exec processes file system attributes such as junction points and symbolic links.

See [“Configuring file and folder options for backup jobs”](#) on page 552.

Enterprise Vault	Select this option to configure Enterprise Vault options. See “Enterprise Vault backup options” on page 1027.
Linux and Macintosh	Select this option to configure options for any Linux or Macintosh computers that are included in the backup job. See “Linux backup options” on page 1205. See “Default backup job options for Macintosh systems” on page 1235.
Lotus Domino	Select this option to configure options for any Lotus Domino data that is included in the backup job. See “Lotus Domino Agent backup options” on page 1070.
Microsoft Active Directory	Select this option to configure options for any Microsoft Active Directory data that is included in the backup job. See “Microsoft Active Directory backup job options” on page 1091.
Microsoft Exchange	Select this option to configure options for any Microsoft Exchange data that is included in the backup job. See “Backing up Exchange data” on page 950.
Virtual Machines	Select this option to configure options for any virtual machines that are included in the backup job. See “Backing up Microsoft Hyper-V virtual machines” on page 873. See “Backing up VMware data” on page 838.
Microsoft SharePoint	Select this option to configure options for any Microsoft SharePoint data that is included in the backup job. See “Backing up Microsoft SharePoint data” on page 974.
Microsoft SQL	Select this option to configure options for any Microsoft SQL data that is included in the backup job. See “Backing up SQL databases and transaction logs” on page 899.
NDMP	Select this option to configure options for any NDMP data that is included in the backup job. See “NDMP backup options for NDMP servers” on page 1174.

Oracle	Select this option to configure options for any Oracle data that is included in the backup job. See “Oracle backup options” on page 1011.
Archive Option	Select this option to configure how Backup Exec works with any archiving stages in the backup definition. See “Archive Option backup options” on page 1380.
Exclusions	Select this option to exclude specific files or specific types of files from backups. See “Excluding files from backups” on page 178.

- 12 When you are finished selecting all applicable options for the backup definition, click **OK**.
- 13 (Optional) You can add stages to this backup definition to duplicate your data, create virtual machines with your backup data, or archive your backup data.

Note: You cannot add a stage to one-time backups.

See [“Adding a stage to a backup definition”](#) on page 219.

- 14 Click **OK** to save the backup definition settings.

The backup jobs and stages run at the time that you selected.

See [“Creating a server group”](#) on page 154.

See [“Editing backup definitions”](#) on page 206.

See [“About selecting data to back up”](#) on page 169.

See [“Changing the order in which backup sources are backed up”](#) on page 175.

How Backup Exec backs up and restores volumes that are enabled for deduplication in Windows

Windows Server 2012 introduced native file system deduplication. When a Windows volume deduplicates data, the deduplicated data is considered to be optimized. Data in its original, non-deduplicated format is considered to be non-optimized.

Backup Exec supports the backup of volumes that have Windows deduplication enabled. If you back up optimized data, Backup Exec backs it up in its original, non-optimized format. Ensure that you have enough space to back up the non-optimized data on the backup target before you run the backup job. The optimized files are not affected on the backup source itself.

When you restore the data that you backed up, Backup Exec restores the files as non-optimized. Ensure that you have enough disk space to restore the non-optimized data. You should free space on the volume on which you want to restore the files. Run a Windows garbage collector to optimize the space on the volume before you run the restore.

How Backup Exec backs up and restores Microsoft Virtual Hard Disk (vhd) files

Microsoft Windows 2008 R2 gave users the ability to create native Virtual Hard Disk (vhd) files. A vhd file is a virtual hard disk contained in a single file. For more information about vhd files, see your Microsoft Windows documentation.

Backup Exec gives you the ability to back up and restore native vhd files. If a native vhd file is not mounted, you can back up and restore it with the volume on which it resides.

If a native vhd file is mounted to a drive letter or to an empty folder path, the file is skipped during backup jobs. You cannot include a mounted vhd as part of your backup selections. To back up the data in a mounted vhd file, select its mount point in the backup selections.

See [“Backing up data”](#) on page 158.

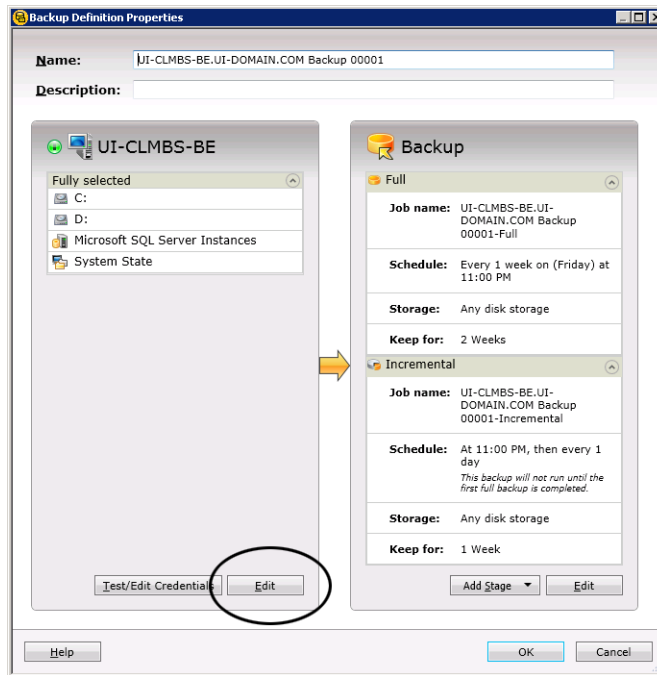
You can also redirect a restore job to a native vhd if you use Microsoft Windows 2008 R2 or later. When you redirect a restore job to a native vhd, Backup Exec creates a vhd file that expands dynamically as you save data to it. The file expands until it reaches 2040 GB, which is the maximum size for a native vhd file. You can create one vhd file with data from all redirected backup sets or you can create a vhd file for each backup set.

Backup Exec's Agent for Hyper-V supports Microsoft vhd files:

See [“Using the Agent for Hyper-V with Windows Server 2012/2012 R2”](#) on page 865.

About selecting data to back up

When you back up a server, Backup Exec includes all of the data on the server in the backup selections by default. If you want to modify the backup selections, you can click **Edit** in the **Selections** box on the **Backup Definitions Properties** dialog box.

Figure 4-2 Selections box on the **Backup Definition Properties** dialog box

See [“Backing up data”](#) on page 158.

Instead of backing up all of the data on a server, you can select drives, folders, files, System State, network shares, or databases on the **Browse** tab.

To expand or collapse the view for an item, click the arrow next to it or double-click the item's name. To view the contents of an item, double-click the item's icon. The item's contents appear in the right frame of the backup selections view. You can traverse file levels from either side of the window by clicking folders and subfolders as they appear.

You can explicitly select any item on the **Browse** tab by right-clicking the item and selecting **Make Item a Required Selection**. Explicitly selected backup sources are included in the backup selections. You can explicitly select critical resources to ensure that you can perform an SDR recovery, conversion to a virtual machine, or a complete online restore. Explicitly selected backup sources are usually backed up before any backup sources that are dynamically included in the backup.

When you browse remote selections, Backup Exec requires a valid logon account to expand the computer contents. If the default logon account does not enable access to a remote selection, Backup Exec prompts you to select another existing

logon account. You can also create a new logon account that can access the selection.

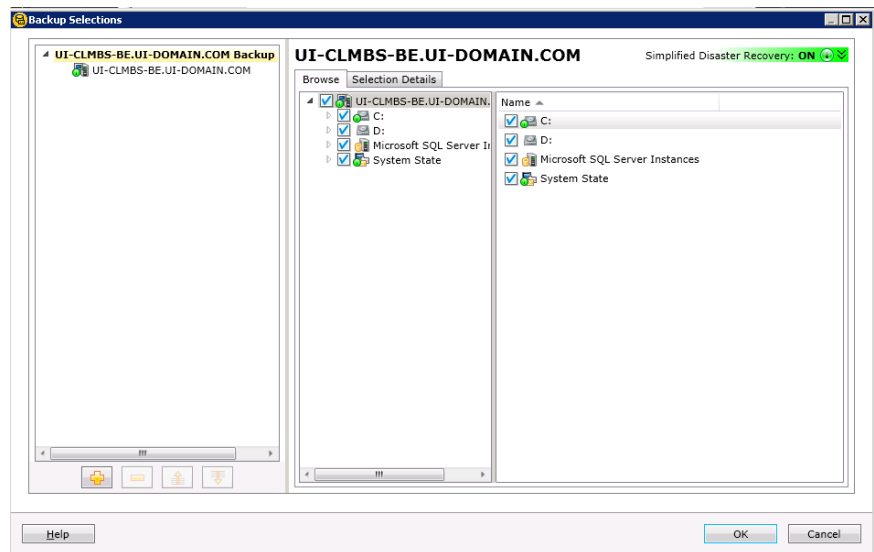
See [“Testing credentials for backup sources”](#) on page 211.

See [“Replacing the credentials for a backup source”](#) on page 212.

See [“Creating new credentials for a backup source”](#) on page 213.

To include data in the backup, select the check box next to the drive or directory that you want to back up.

Figure 4-3 Selecting data for a single server



This topic includes the following information:

[Including and excluding specific files or subdirectories](#)

[Selecting multiple servers or applications](#)

[Selecting critical system components](#)

[Using fully qualified domain names in backup selections](#)

Including and excluding specific files or subdirectories

The **Selection Details** tab lets you include or exclude files for backups by specifying file attributes. Exclusions apply to all of the jobs in a backup definition.

You can do any of the following on the **Selection Details** tab:

- Include or exclude subdirectories. For example, you can choose to back up a parent folder without backing up any folders that reside inside it.
- Include only modified files. For example, you can choose to back up only the files that have changed since the last backup job.
- Include only read-only files.
- Include or exclude files by file name attributes. For example, you can select only files with .txt extensions, or exclude files with .exe extensions from a backup. If you exclude files by an attribute that does not exist, all files of that type are excluded. For example, excludes based on SQL database dates result in global SQL excludes since SQL databases do not have date attributes.
- Select only any files that fall within a specified date range. For example, you can select any files that were created or modified during the month of December.
- Specify the files that have not been accessed in a specified number of days. For example, you can select the files that have not been accessed in 30 days from your "My Documents" folder. Then, run a full backup job for which you select the method to back up and delete the files.
The Backup Exec Archive Option offers more features for data archiving.
See ["About the Archiving Option"](#) on page 1313.

See ["Excluding files from backups"](#) on page 178.

See ["Including specific files with a backup job's backup selections"](#) on page 181.

Selecting multiple servers or applications

If you select to back up data from multiple servers or applications, you can create separate backup definitions for each server or application. Alternatively, you can create one backup definition that includes all servers and applications. If you create separate backup definitions, it is easier to identify problems when backup jobs fail. Also, if an issue arises with one server that causes a backup job to fail, the other backup jobs can still complete successfully. If you create one backup definition that includes all of your servers and applications, it takes less work to monitor the job's status. You can select a default method for backing up multiple servers in the **Backups** dialog of the Backup Exec Settings.

See ["Configuring the default setting for backing up multiple servers or applications"](#) on page 566.

Selecting critical system components

When all the critical system components are included in your backup job selections, the **Simplified Disaster Recovery** indicator on the selections pane reads **ON**. If you deselect one or more critical system component files, the indicator changes to **OFF**.

Figure 4-4 Simplified Disaster Recovery indicator is ON in the **Backup Definition Properties** dialog box

Simplified Disaster
Recovery indicator

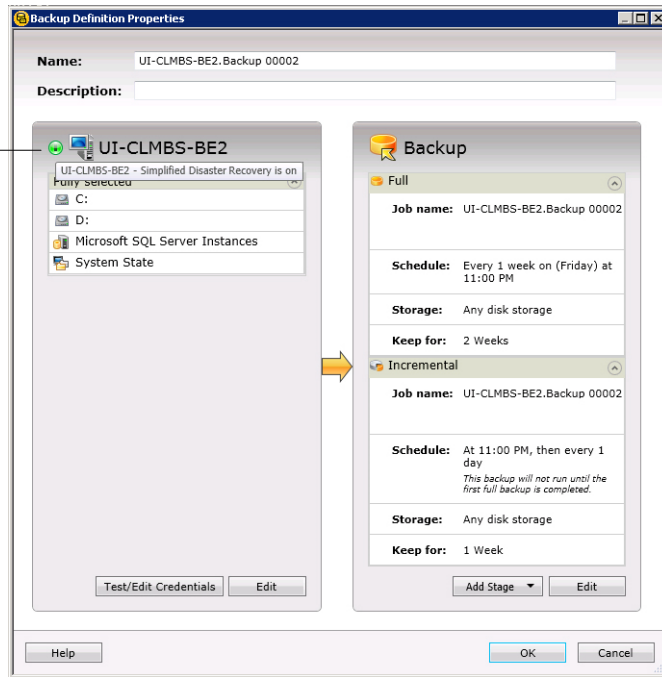
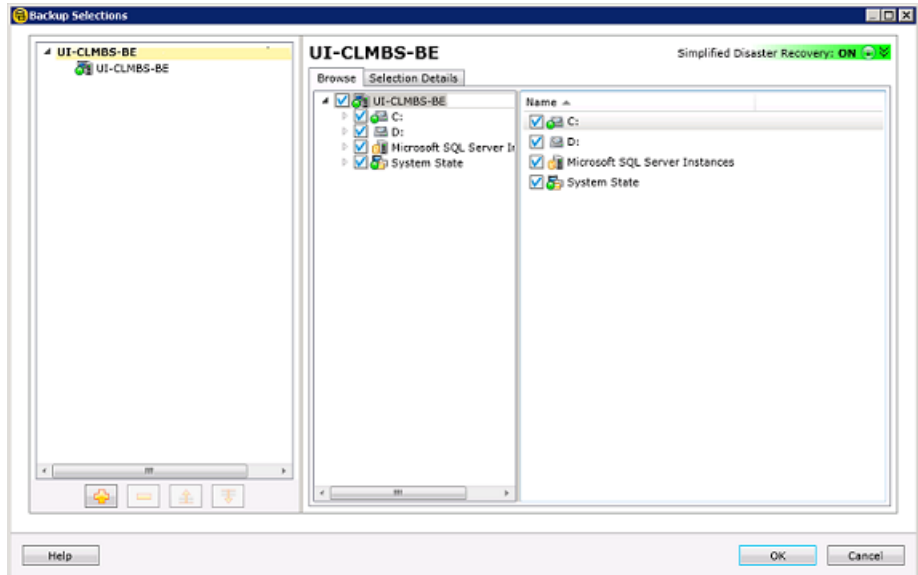


Figure 4-5 Simplified Disaster Recovery indicator is ON in the **Backup Selections** dialog box



See [“About backing up critical system components”](#) on page 183.

If you deselect any critical system components, it can disqualify your backup data from being used in certain types of restore scenarios.

You must include all critical system components in your backup selections if you intend to use any of the following restore scenarios:

- Simplified Disaster Recovery
- Conversion to virtual machines
- Complete online restore of a Microsoft Windows computer

Using fully qualified domain names in backup selections

You can enter fully qualified domain names in Backup Exec anywhere that you can enter a computer name. In addition, Backup Exec can show fully qualified domain names where computer names are listed.

For fully qualified domain names, the following rules apply:

- The maximum number of characters for each label (the text between the dots) is 63
- The maximum total number of characters in the fully qualified name is 254, including the dots, but excluding the \\

- The name cannot include the following characters: * | < > ?

Symantec does not recommend using both fully qualified domain names and non-qualified domain names. Symantec recommends using fully qualified domain names.

For example, if you have a computer named Test_Computer, you can have two selections for it. One selection is called Test_Computer. The fully qualified selection is called Test_Computer.domain.company.com. In this case, Backup Exec treats each selection as a separate computer, even though both selections are for the same computer. For any backup jobs that use the short computer name, the catalog contains the short computer name. For any backup jobs that use the fully qualified name, the catalog contains the fully qualified name.

See [“Changing the order in which backup sources are backed up”](#) on page 175.

Changing the order in which backup sources are backed up

After you make selections for a backup job, you can configure Backup Exec to process those selections in a certain order. You can create backup jobs where your most important backup sources are backed up first, for example.

You must select a backup source to reorder the source or any of its children. If a backup source is included dynamically in the backup, you cannot specify the order in which any of its children are backed up. Any time that you include or exclude a backup source, it creates an entry in the backup's selection details. You can reorder any of the entries that appear on the **Selection Details** tab.

You can explicitly select any item on the **Browse** tab by right-clicking the item and selecting **Make Item a Required Selection**. Explicitly selected backup sources are included in the backup selections. You can explicitly select critical resources to ensure that you can perform an SDR recovery, conversion to a virtual machine, or a complete online restore. Explicitly selected backup sources are usually backed up before any backup sources that are dynamically included in the backup.

There are some limitations to which backup sources can be reordered:

- Any backup sources that reside on branches in the browse tree must be backed up together. Branches are containers for backup sources in the browse tree. Individual applications appear in the Backup Exec browse tree as branches. If you have a server which contains SharePoint, SQL, and Exchange data, each application appears as a branch. You can change the backup order of individual items on a branch, but you cannot change the order of items across multiple branches.

For example, if you want to back up a volume and a SQL database, each of the SQL instances must be backed up in succession. You can change the order in which the SQL instances are backed up. You cannot back up a SQL instance and then the volume before backing up the other SQL instances, however.

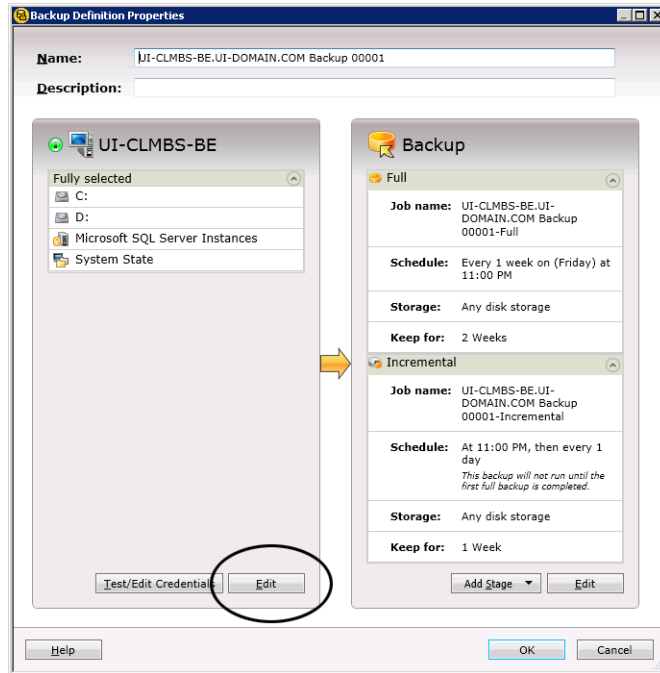
- All of the data that resides on a specific backup source is treated as a single item. You can change the order in which the backup sources are backed up in relation to one another. You can also change the order in which the data on a backup source is backed up. However, you cannot reorder the data across multiple backup sources.

For example, if you back up two volumes, you can select which volume should be backed up first. You can also select the order in which the data on each volume is backed up. However, you cannot back up some of the data from one volume and then the other volume before backing up the rest of the data on the first volume.

- Some system resources cannot be reordered. For example, Shadow Copy Components and System State backup sources must always be backed up last.

To change the order in which backup sources are backed up

- 1 Create or edit a backup definition.
- 2 In the **Selections** box, click **Edit**.



- 3 Do one of the following:

To reorder the servers in a backup definition with multiple servers

In the left pane of the **Backup Selections** dialog box, use the up and down arrows to configure the order of the servers or applications.

Backup Exec backs up the servers or applications at the top first.

To reorder the backup sources within a server

Do the following:

- In the left pane of the **Backup Selections** dialog box, select the server that contains the resources that you want to reorder.
- In the right pane of the **Backup Selections** dialog box, select the **Selection Details** tab.
- To reorder the backup sources when a server is selected in its entirety, you must insert the individual selections that you want to reorder:
 - Click **Insert**, and then select **Insert Selection**.
 - Select the backup source that you want to reorder, and then click **OK**.

Repeat this step for each backup source that you want to reorder.

- In the right pane of the **Backup Selections** dialog box, use the up and down arrows to configure the order of the backup sources.
Backup Exec backs up the backup sources at the top first.
- Click **OK**.

4 Click **OK**.

See [“Backing up data”](#) on page 158.

See [“About selecting data to back up ”](#) on page 169.

Excluding files from backups

You can exclude specific files or types of files from backups if you want to ensure that they do not get backed up. For example, you may not want to back up any mp3 files, read-only files, or files from specific directories.

You can exclude files from backups in any of three ways:

- Exclude files from backup jobs
When you exclude files from individual backup jobs, the exclusions do not affect any of the other backup jobs in the related backup definition. The job-level exclusions also do not affect any of your other backup definitions. The exclusion applies only to the backup job to which you apply it. You cannot apply a job-level exclusion to Full backups.
- Exclude files from backup definitions

When you exclude files from backup definitions, the exclusions do not affect any of your other backup definitions. The exclusions apply only to any jobs in that backup definition.

- **Exclude files globally from all backups**
When you globally exclude files from all backups, the exclusions apply to all of your backup definitions. Backup Exec automatically skips any global exclusions when you run backup jobs.

To exclude files from backups

1 Do any of the following:

To exclude files from a backup job

Complete the following steps:

- Create or edit a backup definition.
- In the **Backup** box, click **Edit**.
- In the left pane, select **Exclusions**.
- Click **Insert**.

To exclude files from a backup definition

Complete the following steps:

- Create or edit a backup definition.
- In the **Selections** box, click **Edit**.
- On the **Selection Details** tab, click **Insert** and then select **Add Backup-Level Exclusion**.

- To globally exclude files from all backups

Complete the following steps:

 - Click the Backup Exec button.
 - Select **Configuration and Settings**, and then select **Job Defaults**.
 - Click **Exclude Selections**.
 - Click **Insert**.
- 2

Complete any of the following fields to identify the files that you want to exclude:

Resource name	Enter the name of the volume or drive that you want to exclude from backups.
Path	Enter the path of the folder and/or subfolder that contains the files that you want to exclude. You can use wildcard characters. Use a question mark (?) to represent any single character. Use two asterisks (**) to represent any number of characters.
Name	<p>Enter the name of a specific file that you want to exclude from the backup. You can use wildcard characters. Use a question mark (?) to represent any single character. Use two asterisks (**) to represent any number of characters.</p> <p>For example, to exclude all files with a .exe extension, type "***.exe".</p>

3 Complete any of the following options, as necessary:

Apply to subdirectories	Select this option to exclude the contents of all of the subfolders when you select a directory.
Only modified files	Select this option to exclude only the files that have been modified in the directory that you specified.
Only read-only files	Select this option to exclude only read-only files in the directory that you specified.
Files dated	Select this option to exclude the files that were created or modified during a specific time period. You must select the beginning and ending dates for the time period.
Files not accessed in X days	Select this option to exclude any files that have not been accessed in a specific number of days. You must enter the number of days.

4 Click **OK**.

See [“About selecting data to back up ”](#) on page 169.

See [“Including specific files with a backup job's backup selections”](#) on page 181.

Including specific files with a backup job's backup selections

When you create a backup definition, you create a list of backup selections. The backup selections are the files and data that you want to back up when the backup jobs run. You can include additional specific files or backup sources with the rest of your backup selections. When you include files with a job's backup selections, you can select certain conditions which must be met for the file to be included. For example, you may want to include any read only files that reside in a specific directory. When you include a file with a backup job's backup selections, Backup Exec backs up the regular backup selections and it includes the file or files that you specifically selected.

To include specific files with a backup job's backup selections

- 1 Create or edit a backup definition.
- 2 In the **Selections** box, click **Edit**.
- 3 On the **Selection Details** tab, click **Insert** and then select **Insert Selection**.

4 Complete any of the following fields to identify the files that you want to include:

Resource name	Enter the name of the volume or drive that you want to include in the backup.
Path	Enter the path of the folder and/or subfolder that contains the files that you want to include. You can use wildcard characters. Use a question mark (?) to represent any single character. Use two asterisks (**) to represent any number of characters.
Name	<p>Enter the name of a specific file that you want to include in the backup. You can use wildcard characters. Use a question mark (?) to represent any single character. Use two asterisks (**) to represent any number of characters.</p> <p>For example, to include all files with a .exe extension, type "***.exe".</p>

5 Complete any of the following options as necessary:

Apply to subdirectories	Select this option to include the contents of all of the subfolders when you select a directory.
Only modified files	Select this option to include only the files that have been modified in the directory that you specified.
Only read-only files	Select this option to include only read-only files in the directory that you specified.
Files dated	Select this option to include the files that were created or modified during a specific time period. You must select the beginning and ending dates for the time period.
Files not accessed in X days	Select this option to include any files that have not been accessed in a specific number of days. You must enter the number of days.

6 Click **OK**.

See [“About selecting data to back up ”](#) on page 169.

See [“Excluding files from backups”](#) on page 178.

About backing up critical system components

Backup Exec is configured to automatically back up the critical system components that you need to perform a full system restore. Backing up critical system components ensures that you are capable of recovering your computers in the event of a disaster.

When all the critical system components are included in your backup job selections, the **Simplified Disaster Recovery** indicator on the selections pane reads **ON**. If you deselect one or more critical system component files, the indicator changes to **OFF**.

Figure 4-6 Simplified Disaster Recovery indicator is ON in the **Backup Definition Properties** dialog box

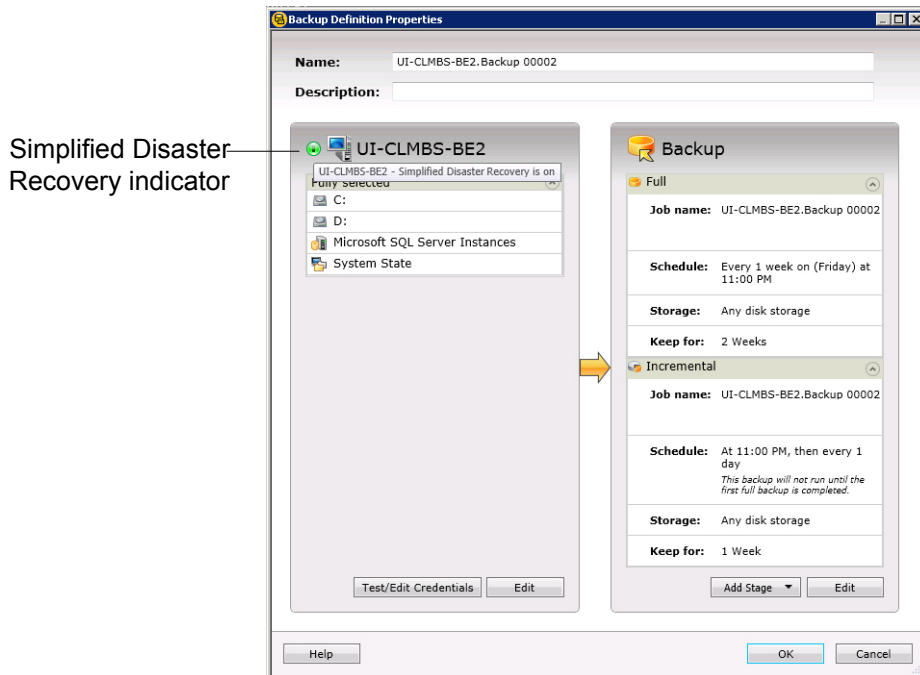
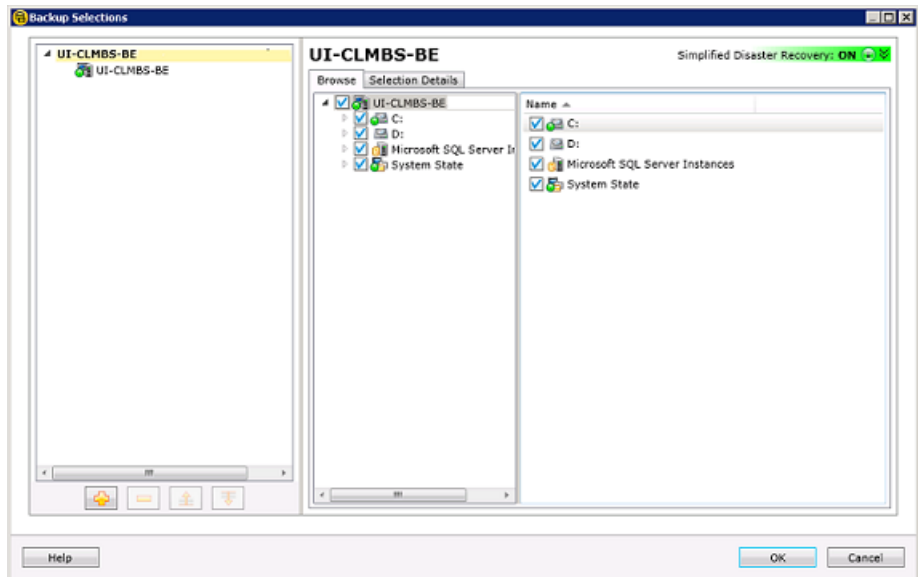


Figure 4-7 Simplified Disaster Recovery indicator is ON in the **Backup Selections** dialog box



You must include all critical system components in your backup selections if you intend to use any of the following restore scenarios:

- Simplified Disaster Recovery
See [“About Simplified Disaster Recovery”](#) on page 736.
- Conversion to virtual machines
See [“How conversion of physical computers to virtual machines works in Backup Exec”](#) on page 479.
- Complete online restore of a Microsoft Windows computer
See [“Performing a complete online restore of a Microsoft Windows computer”](#) on page 239.

Some restore scenarios are only available for certain data types and Backup Exec agents. Additionally, you must select a specific backup method for some data types if you intend to use a physical-to-virtual conversion or a backup-to-virtual conversion.

Table 4-2 Restore scenarios by data type

File System or Agent Name	Operating System and Applications Are Installed and Functional	Complete Online Restore Can Be Performed	Physical-to-Virtual Conversion Can Be Performed	Backup-to-Virtual Conversion Can Be Performed	Simplified Disaster Recovery Can Be Performed
Active Directory	Yes	Yes	Yes	Yes	Yes
Active Directory Lightweight	Yes	No	No	No	No
Archiving Option	Yes	No	No	No	No
CSV	Yes	Yes	No	No	Yes
EFI	Yes	Yes	No	No	Yes
Enterprise Vault	Yes	No	No	No	No
Exchange Agent	Yes	No	Yes, any incremental or differential backups must use the block-level backup method	Yes, any incremental or differential backups must use the block-level backup method	No
FAT	Yes	Yes	Yes	Yes	Yes
Hyper-V Agent	Yes	No	No	No	No
Lotus Agent	Yes	No	No	No	No
NTFS	Yes	Yes	Yes	Yes	Yes
Oracle RMAN Windows Agent	Yes	No	No	No	No
ReFS	Yes	Yes	No	No	Yes

Table 4-2 Restore scenarios by data type (*continued*)

File System or Agent Name	Operating System and Applications Are Installed and Functional	Complete Online Restore Can Be Performed	Physical-to-Virtual Conversion Can Be Performed	Backup-to-Virtual Conversion Can Be Performed	Simplified Disaster Recovery Can Be Performed
Shadow Copy Components	Yes	Yes	Yes	Yes	Yes
SharePoint Agent	Yes	No	No	No	No
SQL Agent	Yes	No	Yes, any incremental or differential backups must use the block-level backup method	Yes, any incremental or differential backups must use the block-level backup method	No
System State	Yes	Yes	Yes	Yes	Yes
Utility Partition	Yes	Yes	No	No	Yes
VMware Agent	Yes	No	No	No	No
Windows Deduplication	Yes	Yes	No	No	Yes

You cannot individually select critical system components as backup selections. You must select the entire server to ensure that all critical system components are included in the backup. When you select to back up a server, Backup Exec includes all of the server's system devices and application agents. Backup Exec dynamically discovers and protects all critical and non-critical system devices and application agents.

You can explicitly exclude any non-critical devices or application data from the backup selections without affecting your ability to perform a full system restore. You can exclude Microsoft Exchange data from your backup, for example, and still use the backup sets to perform a disaster recovery.

The following system resources are considered critical and they must be included in backups if you want to be able to use the backup sets to perform a full system restore:

- System volume (including EFI and utility partitions)
- Boot volume (excluding operating system)
- Services application volumes (boot, system, and automatic startup)
- System State devices and volumes (including Active Directory, System Files, etc.)
- Windows Recovery Partition (WinRE) on Windows 8/Server 2012/8.1/Server 2012 R2

Backup Exec Shadow Copy Components file system

The Backup Exec Shadow Copy Components file system uses Microsoft's Volume Shadow Copy Service to protect critical operating system and application service data, and third-party application and user data on Windows Server 2003/2008 computers.

Volume Shadow Copy Service allows a computer to be backed up while applications and services are running by providing a copy of a volume when a backup is initiated. Applications do not need to be shut down to ensure a successful volume backup. Volume Shadow Copy Service enables third-party vendors to create snapshot plug-ins, or Writers, for use with this shadow copy technology.

A Writer is specific code within an application that participates in the Volume Shadow Copy Service framework to provide point-in-time, recovery-consistent operating system and application data. Writers appear as Shadow Copy Components, which are listed as data in backup and restore selections.

Only Writers that have been tested for use with Backup Exec are available for selection in the backup selections. Other Writers may be displayed in the selections, but they cannot be selected for backup.

If you select a volume that contains Shadow Copy data for backup, Backup Exec determines which Shadow Copy files should not be included in a volume-level backup. These files will be automatically excluded for backup by a feature called Active File Exclusion. If this exclusion did not happen during a non-snapshot backup, these files would appear as "in use - skipped." If this exclusion did not happen during a snapshot backup, the files would be backed up in a possible inconsistent state, which could create restore issues.

The Windows SharePoint Services feature pack utilizes a SQL (MSDE) instance called SHAREPOINT as a repository for shared information and collaboration data.

On Windows Server 2003/2008, in the absence of a Symantec SQL Agent installation, the SQL SHAREPOINT instance can be protected by the Shadow Copy Components file system. If the SQL Agent is installed, then the SQL SHAREPOINT instance can be protected by the SQL Agent.

Note: If Windows SharePoint Services is installed using an instance name other than the default SHAREPOINT instance name, then it cannot be protected by the Shadow Copy Components file system. In that case, the Symantec SQL Agent must be used to protect the SQL SHAREPOINT instance.

Windows Small Business Server 2003 Standard and Premium contain a SQL (MSDE) instance called SBSMONITORING as a repository for server-related activity data. In the absence of a Symantec SQL Agent installation, the SQL SBSMONITORING instance can be protected by the Shadow Copy Components file system. If the SQL Agent is installed, then the SQL SBSMONITORING instance can be protected by the SQL Agent.

Backup methods in Backup Exec

When you create a backup definition, you must select backup methods for each backup job. The backup method determines what data Backup Exec backs up. The standard backup methods are full, differential, and incremental. However, you can also select specific types of backup methods for some agents and types of data.

Each backup definition must contain one backup job that uses the full backup method. This initial full backup job establishes a baseline by backing up all of the data that you selected. Then, you can add additional differential or incremental backup jobs to the backup definition. You can select to use data-specific backup methods instead of the standard backup methods by selecting the appropriate backup method on the options page for that agent.

This topic includes the following information:

[Full backup method](#)

[Differential backup method](#)

[Incremental backup method](#)

[Backup method advantages and disadvantages](#)

[Backup methods for specific types of data](#)

Full backup method

Full backups include all of the data that was selected for backup. Backup Exec detects that the server was backed up.

Note: You should perform a full backup of your server to establish a baseline for disaster recovery.

Duplicate backups, which include all selected data, are a type of full backup. Duplicate backups do not affect any tape media rotation strategy because the archive bit is not reset.

Symantec recommends that you always run full backups before and after you upgrade, update, or modify any backup sources. This recommendation applies to any significant operating system and application configuration or modifications.

Table 4-3 Backup sources and scenarios for which a full backup is recommended

Backup source	Scenarios for which a full backup is recommended
Operating system	<p>You should run a full backup when you:</p> <ul style="list-style-type: none"> ■ Upgrade to a new version of the operating system. ■ Update the existing operating system using Windows Update. <p>Note: You should run a full backup before you apply the update. Apply the update and then restart the computer. Then you should run another full backup.</p> <ul style="list-style-type: none"> ■ Add, modify, or remove any server roles or features.
Applications	<p>You should run a full backup when you:</p> <ul style="list-style-type: none"> ■ Install new applications. ■ Upgrade applications to new versions. ■ Update the existing version of applications.
Backup Exec	<p>You should run a full backup when you:</p> <ul style="list-style-type: none"> ■ Upgrade to a new version of Backup Exec. <p>Note: Any recurring jobs that are migrated to a new version of Backup Exec retain their existing schedules. You should manually run a full backup for any existing jobs before any incremental backups or differential backups run.</p> <ul style="list-style-type: none"> ■ Update the existing version of Backup Exec using LiveUpdate. ■ Make configuration changes.

Running full backups before and after each of these scenarios helps to ensure that you can restore back to the previous configuration, if necessary.

Differential backup method

Differential backups include all files that have changed since the last full backup or incremental backup. The difference between differential and incremental backups is that differential backups are cumulative. After a differential backup, each subsequent differential backup backs up the same files as the previous differential backup. It also backs up as any new files or changed files dating back to the last full backup or incremental backup.

Note: In a backup definition that includes a differential task, all of the backup tasks must use storage devices that the same Backup Exec server can access.

By default, Backup Exec uses the Windows Change Journal to determine if files were previously backed up. You can also configure Backup Exec to use a file's modified time or archive bit to determine if the file was backed up.

Note: You cannot use Backup Exec catalogs to determine if files were backed up for any differential backups.

See [“Configuring file and folder options for backup jobs”](#) on page 552.

Differential backups allow much easier restoration of an entire device than incremental backups since fewer backups are required. Using fewer media also decreases the risk of having a restore job fail because of media errors.

Incremental backup method

Incremental backups include only the files that have changed since the last full or incremental backup. The difference between incremental and differential backups is that incremental backups are not cumulative. Each incremental backup creates a baseline. After the incremental backup, the subsequent incremental backup or differential backup backs up only any new files or changed files dating back to the baseline.

Note: In a backup definition that includes an incremental task, all of the backup tasks must use storage devices that the same Backup Exec server can access.

By default, Backup Exec uses the Windows Change Journal to determine if files were previously backed up. You can also configure Backup Exec to use a file's modified time, archive bit, or the Backup Exec catalogs to determine if the file was backed up.

See [“Configuring file and folder options for backup jobs”](#) on page 552.

Incremental backups take much less time than full or differential backups to complete. They also require less storage space for backed up data since only any files that have changed since the last backup are backed up.

Backup method advantages and disadvantages

Each backup method has advantages and disadvantages.

Table 4-4 Backup method Advantages and Disadvantages

Method	Advantages	Disadvantages
Full	<ul style="list-style-type: none"> Files are easy to find Full backups include all the data that you selected to back up. Therefore, you don't have to search through several backup sets to find a file that you need to restore. A current backup of your entire system is available on one backup set If you run a full backup of your entire system and then need to restore it, all of the most current information is located in one place. 	<ul style="list-style-type: none"> Redundant backups Most of the files on your file server do not change. Each full backup that follows the first is merely a copy of what has already been backed up. Full backups require more storage. Full backups take longer to perform Full backups can be time consuming, especially when you have other servers on the network that need to be backed up (for example, agent workstations, remote servers).

Table 4-4 Backup method Advantages and Disadvantages (*continued*)

Method	Advantages	Disadvantages
Differential	<ul style="list-style-type: none"> Files are easy to find Restoring a system that is backed up with a differential method requires fewer backups. Differentials require the latest full backup, any subsequent incremental backups, and the latest differential backup. Restoring differentials is less time consuming than restoring incrementals. Restoring incrementals requires the latest full backup and all incremental backups that were created since the full backup. Less time is required for backup and restore Differential backups take less time to restore than full backups. Faster recovery is possible in disaster situations because you only need backup sets from the latest full backup, any subsequent incremental backups, and the latest differential backup to fully restore a server. 	<ul style="list-style-type: none"> Redundant backups All of the files that were created or modified since the last full backup are included; thus creating redundant backups.
Incremental	<ul style="list-style-type: none"> Better use of storage Only the files that have changed since the last backup are included, so much less data storage space is required. Less time is required for backup Incremental backups take much less time than full and differential backups to complete. 	<ul style="list-style-type: none"> Backups are spread across multiple backup sets Since multiple backup sets are required in a disaster situation, recovering a server can take longer. In addition, the backup sets must be restored in the correct order to effectively bring the system up to date.

Consider the following backup strategy scenario:

You want to implement a backup strategy for the office file server. All backup strategies begin with a full backup (the backup of an entire server using the full backup method). So you create and submit a full backup job to run at the end of the day on Friday.

Most files on the server, such as operating system files and application files seldom change. Therefore, you decide that you can save time and storage by using incremental backups or differential backups. You opt to use incremental backups. You schedule a job to run at the end of each day, Monday through Thursday, with the incremental backup method.

On Friday, your backup sets contain all of the data on the file server. Backup Exec changes all of the files' statuses to backed up. At the end of the day on Monday, the incremental job runs and only the files that were created or changed are backed up. When the incremental job completes, Backup Exec turns off the archive bit, showing that the files have been backed up. On Tuesday through Thursday, the same events happen.

If your file server then crashed on Thursday, you would restore each backup in the order in which it was created. You would begin with Friday's backup and proceed through Wednesday's backup.

If you had decided to perform differential backups on Monday through Thursday, you would have only needed Friday's and Wednesday's backup sets. Friday's backup sets would have included all of the data from the original backup. Wednesday's backup sets would have included every file that had been created or changed since Friday's backup.

Backup methods for specific types of data

Agents and options may have specific types of backup methods.

By default, each backup definition contains a backup job that limits the available backup method to the initial full backup for most data types. You can configure additional backup jobs that contain special backup methods for specific types of data. You can name the job templates for those backup jobs so that they have additional meaning to your organization's overall backup strategy.

See [“Configuring backup methods for backup jobs”](#) on page 196.

When you create a backup definition that contains more than one backup method for multiple data types, it is called a mixed backup in the Job Monitor.

Table 4-5 Available Backup Methods By Data Type

Data type	Job type and backup method
Files and Folders	<p>Initial full:</p> <ul style="list-style-type: none"> ■ Full - Back up files <p>Additional backup methods for files and folders:</p> <ul style="list-style-type: none"> ■ Full - Back up files ■ Full Copy - Back up files (copy) ■ Differential - Back up changed files since the last full ■ Incremental - Back up changed files since the last full or incremental
Enterprise Vault	<p>Initial full:</p> <ul style="list-style-type: none"> ■ Full - Back up components <p>Additional backup methods for Enterprise Vault:</p> <ul style="list-style-type: none"> ■ Full- Back up components ■ Differential - Back up component changes since the last full ■ Incremental - Back up component changes since the last full or incremental
Lotus Domino	<p>Initial full:</p> <ul style="list-style-type: none"> ■ Full - Back up databases and logs (reset archive bit) <p>Additional backup methods for Lotus Domino:</p> <ul style="list-style-type: none"> ■ Full - Back up databases and logs (reset archive bit) ■ Incremental - Back up changed databases and logs (reset archive bit)
Microsoft Exchange	<p>Initial full:</p> <ul style="list-style-type: none"> ■ Full - Back up databases and logs (truncate logs) ■ Full Copy - Back up databases and logs <p>Additional backup methods for Microsoft Exchange:</p> <ul style="list-style-type: none"> ■ Full - Back up databases and logs (truncate logs) ■ Full Copy - Back up databases and logs ■ Differential - Back up logs ■ Incremental - Back up logs (truncate logs)

Table 4-5 Available Backup Methods By Data Type (*continued*)

Data type	Job type and backup method
Virtual Machines	<p>Initial full:</p> <ul style="list-style-type: none"> ■ Full - Back up virtual machines <p>Additional backup methods for virtual machines:</p> <ul style="list-style-type: none"> ■ Full - Back up virtual machines ■ Differential - Back up virtual machine changes since the last full ■ Incremental - Back up virtual machine changes since the last full or incremental
Microsoft SharePoint	<p>Initial full:</p> <ul style="list-style-type: none"> ■ Full - Back up databases ■ Full Copy - Back up databases (copy) <p>Additional backup methods for Microsoft SharePoint:</p> <ul style="list-style-type: none"> ■ Full - Back up databases ■ Full Copy - Back up databases (copy) ■ Differential - Back up database changes since the last full ■ Differential (block-level) - Back up database changes since the last full - use with convert to virtual machine job ■ Incremental (block-level) - Back up database changes since the last full or incremental - use with convert to virtual machine job ■ Log - Back up and truncate transaction log
Microsoft SQL	<p>Initial full:</p> <ul style="list-style-type: none"> ■ Full - Back up databases ■ Full Copy - Back up databases (copy) <p>Additional backup methods for Microsoft SQL:</p> <ul style="list-style-type: none"> ■ Full - Back up databases ■ Full Copy - Back up databases (copy) ■ Automatic - Back up transaction log if enabled and then back up database changes since the last full or incremental ■ Log - Back up and truncate transaction log ■ Log No Truncate - Back up without truncating transaction log ■ Differential - Back up database changes since the last full ■ Differential (block-level) - Back up database changes since the last full - use with convert to virtual machine job ■ Incremental (block-level) - Back up database changes since the last full or incremental - use with convert to virtual machine job ■ Database Snapshot - Read-only point-in-time copy of databases

Table 4-5 Available Backup Methods By Data Type (*continued*)

Data type	Job type and backup method
NDMP (all)	<p>Initial full:</p> <ul style="list-style-type: none"> Level 0 - Full backup <p>Additional backup methods for NDMP:</p> <ul style="list-style-type: none"> Level 0 - Full backup Level 1 - Incremental (backs up new or modified files since level 0) Level 2 - Incremental (backs up new or modified files since level 1) Level 3 - Incremental (backs up new or modified files since level 2) Level 4 - Incremental (backs up new or modified files since level 3) Level 5 - Incremental (backs up new or modified files since level 4) Level 6 - Incremental (backs up new or modified files since level 5) Level 7 - Incremental (backs up new or modified files since level 6) Level 8 - Incremental (backs up new or modified files since level 7) Level 9 - Incremental (backs up new or modified files since level 8)
Oracle	<p>Initial full:</p> <ul style="list-style-type: none"> Full - Back up selections <p>Additional backup methods for Oracle:</p> <ul style="list-style-type: none"> Full - Back up selections Differential - Back up changes since the last full Incremental - Back up changes since the last full or incremental
Archive Options	<p>Initial full:</p> <ul style="list-style-type: none"> Full - Back up components <p>Additional backup methods for Archive Option:</p> <ul style="list-style-type: none"> Full - Back up components Differential - Back up component changes since the last full Incremental - Back up component changes since the last full or incremental

Configuring backup methods for backup jobs

When you create a backup definition, you must select backup methods for each backup job. The backup method determines what data Backup Exec backs up. The standard backup methods are full, differential, and incremental. However, you can also select specific backup methods for some agents and types of data.

See [“Backup methods in Backup Exec”](#) on page 188.

By default, each backup definition contains one backup job with an initial full backup method that cannot be changed and a backup job with the incremental backup method. You can replace the incremental job with a differential job, add additional incremental or differential backup jobs, or change the backup methods to data-specific backup methods for certain types of data. Each backup method can be given a unique name to help make it more easily identifiable.

To configure backup methods for backup jobs

- 1 Create or edit a backup definition.
- 2 In the **Backup** box, click **Edit**.
- 3 In the left pane, click **Schedule**.
- 4 (Optional) To add additional backup jobs to the backup definition, complete the following steps:
 - Click **Add a Backup Job**.
 - Select the type of backup method that you want to use for the new backup job.

Note: You can change the backup method to a data specific one for specific agents and types of data later.


- 5 Type a name for each backup job that you want to configure a backup method for in the **Job template name** field.

The screenshot shows the 'Backup Options' dialog box. On the left is a sidebar with a list of categories: Schedule, Storage, Network, Notification, Test Run, Verify, Advanced Open File, Advanced Disk-based Backup, Security, Pre/Post Commands, Files and Folders, Microsoft SQL, and Exclusions. The 'Schedule' category is selected and highlighted in yellow.

The main area of the dialog is divided into two sections. The top section is for a job template named 'Full'. It includes a text field for 'Job template name' with the value 'Full'. Below this is a text field for 'Job name' with the value 'dheath-w2k8r2.HROUS.sen.symantec.com Backup 00001-Full'. The 'Schedule' section has three radio buttons: 'Recurrence' (selected), 'Run now with no recurring schedule', and 'Create without a schedule'. The 'Recurrence' dropdown is set to 'Every 1 week on (Friday) at 11:00 PM'. There is also a checkbox for 'Submit job on hold'.

The bottom section is for a job template named 'SQL truncate logs'. It has a text field for 'Job template name' with the value 'SQL truncate logs'. The 'Job name' field contains 'dheath-w2k8r2.HROUS.sen.symantec.com Backup 00001-SQL truncate logs'. The 'Schedule' section has three radio buttons: 'Recurrence' (selected), 'Create without a schedule', and 'Submit job on hold'. The 'Recurrence' dropdown is set to 'At 11:00 PM, then every 1 day'. At the bottom of the dialog are 'Help', 'OK', and 'Cancel' buttons.

- 6 (Optional) To change any of the backup methods that you selected into a data-specific backup method, complete the following steps:
 - In the left pane, select the agent or data type for which you want to configure the data-specific backup method.
For example, if you want to select a data-specific backup method for SQL data, select **Microsoft SQL**.
 - In the **Backup method** field, select the data-specific backup method for each applicable backup job.



7 Click **OK**.

How Backup Exec determines if a file has been backed up

If you use the incremental or the differential backup method as part of your backup strategy, Backup Exec must know when a file has been modified. Full backups include all of the data that you selected to back up. Subsequent incremental and differential backups back up only new files and any files that have changed.

Whenever a file is created or changed, a computer's file system notes and records the change. You can specify the method by which Backup Exec determines if a file needs to be backed up in the **Files and Folders** options when you create a backup job. Backup Exec uses the method that you choose to determine when a file is created or changed.

This topic includes the following information:

[Modified time](#)

[Archive bit](#)

[Catalogs](#)

[How Backup Exec uses the change journal to determine which files have changed](#)

[Resolving change journal errors](#)

Modified time

When Backup Exec runs a full backup or incremental backup job, the time that the backup job starts is recorded in the Backup Exec Database. Backup Exec adds the time of the backup job to the Backup Exec Database only if the full backup job completes successfully. The next time that you run an incremental backup job or a differential backup job, Backup Exec compares the file system time to the backup time. If the file system time is later than the time that is recorded in the database, the file is backed up. If the file's modified time is older than the previous backup's modified time, that file is not backed up. If the job does not complete successfully, subsequent differential or incremental backup jobs back up all of the data instead of only the data that has changed.

Note: A file's last modified date and timestamp do not change when the file is copied or moved. To ensure that the files are protected, run a full backup after you copy or move files.

When you run an incremental backup job, Backup Exec records a new time in the Backup Exec Database. The database time is not updated for differential backup jobs.

When you select the modified time method, Backup Exec uses the Windows change journal to determine if a file has changed since the last time it was backed up. If the change journal is not available, Backup Exec compares the file information to the previous backup time to determine if the file has changed.

Using modified time lets Backup Exec run more accurate incremental backups or differential backups even if other processes have modified files' archive bits.

Archive bit

Backup Exec uses the archive bit from the file system to determine if a file has changed since the last time it was backed up.

When you use the archive bit, Backup Exec turns the archive bit off when a file is backed up. Turning off the archive bit indicates to Backup Exec that the file has been backed up. If the file changes again before the next backup job, the bit is turned on again. Backup Exec backs up the file in the next backup.

If the next backup job is a full backup job or an incremental backup job, the bit is turned off when the backup job completes. If the next backup job is a differential backup job, the archive bit is left intact.

Catalogs

Backup Exec compares path names, modified time, deleted and renamed files and folders, and other attributes. When you select the catalog method, Backup Exec uses the Windows change journal to determine if a file has changed since the last time it was backed up. If the change journal is not available, Backup Exec compares the file information to previous catalogs to determine if it has changed.

The catalog method is only available if the Advanced Disk-based Backup Option (ADBO) is installed.

Note: The off-host backup feature of ADBO does not support the catalog method.

See [“Configuring file and folder options for backup jobs”](#) on page 552.

How Backup Exec uses the change journal to determine which files have changed

When you choose the modified time backup method or the catalog backup method for files, Backup Exec uses the update sequence number (USN) change journal. Backup Exec scans the change journal to get a list of the changed files to back up, rather than scanning all files. Use of the change journal reduces the amount of time that Backup Exec requires to perform an incremental backup or a differential backup.

You cannot disable the use of the change journal for either the modified time backup method or the catalog backup method. If the change journal is not already enabled on the volume, then Backup Exec automatically enables it. The volume must support change journals. For example, NTFS and ReFS volumes support change journals, but FAT volumes do not.

When Backup Exec enables the change journal, it sets the change journal defaults as shown in [Table 4-6](#). Backup Exec does not modify settings for a change journal that already exists.

Table 4-6 Change journal defaults set by Backup Exec

Volume size	Defaults
128 GB or larger	<ul style="list-style-type: none">Change Journal Max Size: 32 MBAllocation Delta: 4 MB
Over 64 GB - 127 GB	<ul style="list-style-type: none">Change Journal Max Size: 16 MBAllocation Delta: 2 MB

Table 4-6 Change journal defaults set by Backup Exec (*continued*)

Volume size	Defaults
Over 4 GB - 63 GB	<ul style="list-style-type: none">■ Change Journal Max Size: 8 MB■ Allocation Delta: 1 MB
4 GB or less	<ul style="list-style-type: none">■ Change Journal Max Size: 1 MB■ Allocation Delta: 256 KB

Note: The Windows Boot Volume is preconfigured by Windows with default values.

Resolving change journal errors

If change journal errors occur, then Backup Exec uses the modified time backup method or the catalog backup method without using the change journal.

When change journal errors occur, you can try the following solutions:

- Select the **Using archive bit** backup method for files.
- Select the file backup method that was not in use when the change journal errors occurred. For example, if you selected the **Modified time** method and change journal errors occurred, then select the **Using catalogs** method and run the backup again.
- Increase the size of the change journal database.
Refer to your Microsoft documentation for information on how to use fsutil to query, create, modify, or delete the change journal.
- Avoid using the local Backup Exec server to back up itself when the Backup Exec server is installed on the Windows Boot volume. Instead, use a remote Backup Exec server to back up the local server, or install the Backup Exec server on a data volume.

Note: The off-host backup feature of ADBO does not support the catalog backup method for files.

Configuring Backup Exec to automatically delete files after a backup

When you run a full backup, you can select to back up the files and then delete them. The **Delete selected files and folders after successful backup** option lets you free disk space on your server by deleting files and folders from the server after

they are successfully backed up. Backup Exec backs up the selected data, verifies the backup sets, and then deletes the data from the server. You can back up and delete files for full backups only.

See [“Configuring file and folder options for backup jobs”](#) on page 552.

Backup Exec performs a verify operation after the data is backed up. If the verify operation fails, the job stops and you are notified. If you get a verification failure, view the job log. Try to correct the problem, and then retry the job. After the data is backed up and verified, Backup Exec deletes the selected data. The job log contains a list of the data that is deleted.

You can enable the checkpoint restart option for a full backup job that uses the **Delete selected files and folders after successful backup** option. If the job fails and is resumed, the files are not deleted from the source volume after the backup completes.

The Backup Exec Archive Option offers more features for data archiving.

See [“About the Archiving Option”](#) on page 1313.

To configure Backup Exec to automatically delete files after a backup

1 Do one of the following:

To configure Backup Exec to automatically delete files after all backup jobs

Complete the following steps:

- Click the Backup Exec button, and then select **Configuration and Settings**.
- Select **Job Defaults**, and then select the type of backup for which you want to configure Backup Exec to automatically delete files.

To configure Backup Exec to automatically delete files for specific backup jobs

Complete the following steps:

- Create a new backup definition or edit an existing backup definition.
- In the **Backup** box, click **Edit**.

2 In the left pane, click **Files and Folders**.

3 Select **Delete selected files and folders after successful backup**.

Note: You can select to delete files and folders only for full backups.

4 Click **OK**.

Configuring network options for backup jobs

You can configure options for how Backup Exec works with your network. Backup Exec contains global network and security settings that apply to all jobs.

See [“Changing network and security options for Backup Exec”](#) on page 576.

You can override the global network settings when you create backup jobs if the global settings do not apply in a particular instance. Complete the steps in the following procedure to configure network options for individual backup jobs.

Note: Network options differ in CASO environments.

To configure network options for backup jobs

- 1 Create a new backup definition or edit an existing backup definition.
- 2 In the **Backup** box, click **Edit**.
- 3 In the left pane, click **Network**.
- 4 Complete the following options:

Network interface	Select the name of the network interface card that connects the Backup Exec server to the network that you want to use for this backup job. The list includes all available network interfaces on the Backup Exec server.
Protocol	<div>Select the protocol you want to use for this backup job.</div> <div>The options are as follows:</div> <div><ul style="list-style-type: none">■ Use any available protocol■ IPv4■ IPv6</div>
Subnet	Select the 32-bit number that determines the subnet to which the network interface card belongs.
Allow use of any available network interface, subnet, or protocol for Backup Exec agents not bound to the above network interface, subnet, or protocol	<div>Select this option to let Backup Exec use any available network if the remote system that you selected for backup or restore is not part of the specified backup network.</div> <div>If you do not select this option and the remote system is not part of the specified backup network, the job fails. Backup Exec cannot access the data from the remote system.</div>

Interface Details

Click this option to view the Media Access Control (MAC) address, adapter type, description, IP addresses, and subnet prefixes for the interface that you selected for the backup network.

**Allow managed Backup
Exec server to use any
network interface to access
Backup Exec agents**

Select this option to let a job use any network interface to access Backup Exec agents if the selected network interface is unavailable. Enabling this option lets the managed Backup Exec server use an alternate network interface to run any important backup jobs that would otherwise fail.

This option is available only if the Central Admin Server Option (CASO) is installed.

See [“About the Central Admin Server Option”](#) on page 1096.

5 Click **OK**.

See [“Backup networks”](#) on page 574.

Running the next scheduled backup job before its scheduled time

You can run the next scheduled backup job in a backup definition at any time. You may want to run a scheduled backup job early to ensure that important data gets backed up or to make sure that a scheduled job completes successfully. Running a scheduled backup job early does not affect its regular schedule. The job still runs normally as scheduled.

To run the next scheduled backup job

- 1 On the **Backup and Restore** tab, do one of the following:
 - To run the next scheduled backup for a single server's backup jobs, right-click the server name.
 - To run the next scheduled backup for multiple servers' backup jobs, Shift + click or Ctrl + click the server names, and then right-click one of the selected servers.

2 Click **Run Next Backup Now**.

3 Click **Yes** to confirm that you want to run the job or jobs now.

See [“Backing up data”](#) on page 158.

Editing backup definitions

You can edit existing backup definitions. You can modify any existing backup definition's backup selections or the backup job settings for any backup jobs that it contains. The backup selections include any servers, volumes, or data that you have selected to back up. Backup job settings can include scheduling options, storage device options, or backup methods for selected types of data, for example.

You can also add a stage to an existing backup definition to duplicate your backup data, create virtual machines with your backup data, or archive your backup data.

See [“Adding a stage to a backup definition”](#) on page 219.

If you choose to edit more than one backup definition at once, you can edit only the properties that the definitions have in common. For example, if you choose to edit two backup definitions at once and the definitions use different schedules, you cannot edit the schedules. If you do not see the settings that you want to edit, repeat this procedure, but select only one definition to edit at a time.

Note: You cannot edit a backup definition while one of its backup jobs is running.

This topic includes the following procedures:

[To edit a backup definition's backup selections or select user shares to back up](#)

[To edit a backup definition's job settings](#)

To edit a backup definition's backup selections or select user shares to back up

1 Do one of the following:

To edit backups from the
Backup and Restore tab

Complete the following steps:

- On the **Backup and Restore** tab, do one of the following:
 - To edit backups for a single server, right-click the server name.
 - To edit backups for multiple servers, Shift + click or Ctrl + click the server names, and then right-click one of the selected servers.
- Click **Edit Backups**.

Note: If the server or servers that you selected have multiple backup definitions, select the definitions that you want to edit on the **Backup Job Selection** dialog box and then click **OK**.

To edit backups from the **Job** Complete the following steps:

Monitor tab

- On the **Job Monitor** tab, do one of the following:
 - To edit a single backup, right-click the job name.
 - To edit multiple backups at once, Shift + click or Ctrl + click the job names, and then right-click one of the selected jobs.
- Click **Edit**.

To edit backups from the **Storage** tab

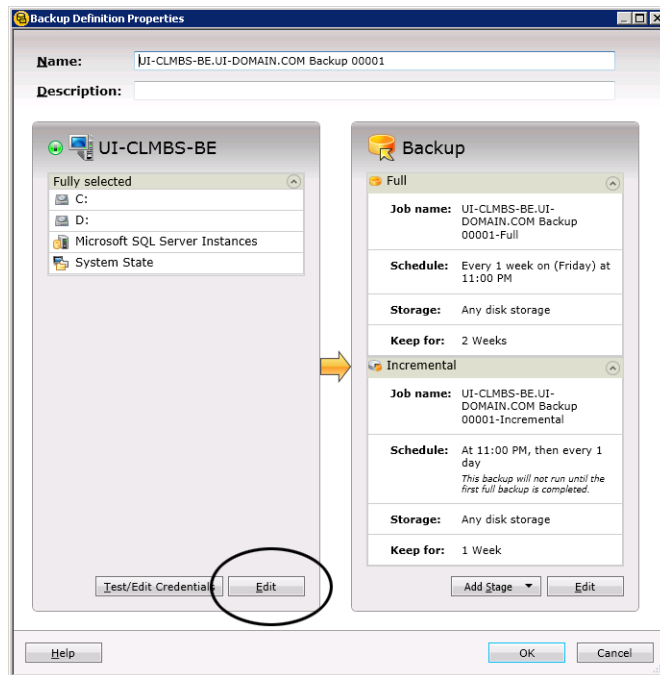
Complete the following steps:

- On the **Storage** tab, double-click the storage device or storage pool for the backup that you want to edit.
- On the **Job History** pane, do one of the following:
 - To edit a single backup, right-click the backup.
 - To edit multiple backups, Shift + click or Ctrl + click the backups and then right-click one of the selected backups.

Note: You can only edit any backup jobs that have previously run from the **Storage** tab.

- Click **Edit Backup**.

2 In the **Selections** box, click **Edit**.



3 To add additional servers to the backup definition:

- Click the **Add (+)** button.
- Select the server or servers that you want to add to the backup definition. You can also click **New Server** to add a new server to Backup Exec.

- Click **OK**.
- 4 Select the data that you want to back up.

Servers are selected in their entirety by default. If you do not want to back up the entire server, double-click the server name to see all of the contents. Select the check boxes for each item that you want to back up.

Note: Deselecting a server's critical system components creates backup sets that cannot be used for some restore scenarios.

See [“About backing up critical system components”](#) on page 183.

- 5 Click **OK**.
- 6 When you are finished editing the backup definition, click **OK** on the **Backup Properties** dialog box.

To edit a backup definition's job settings

- 1 Do one of the following:

To edit backups from the **Backup and Restore** tab

Complete the following steps:

- On the **Backup and Restore** tab, do one of the following:
 - To edit backups for a single server, right-click the server name.
 - To edit backups for multiple servers, Shift + click or Ctrl + click the server names, and then right-click one of the selected servers.
- Click **Edit Backups**.

Note: If the server or servers that you selected have multiple backup definitions, select the definitions that you want to edit on the **Backup Job Selection** dialog box and then click **OK**.

To edit backups from the **Job Monitor** tab

Complete the following steps:

- On the **Job Monitor** tab, do one of the following:
 - To edit a single backup, right-click the job name.
 - To edit multiple backups at once, Shift + click or Ctrl + click the job names, and then right-click one of the selected jobs.
- Click **Edit**.

To edit backups from the **Storage** tab

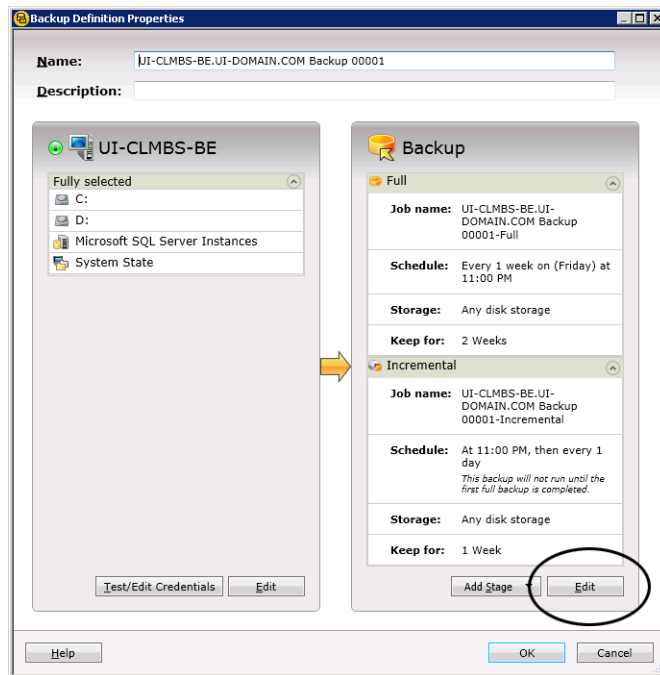
Complete the following steps:

- On the **Storage** tab, double-click the storage device or storage pool for the backup that you want to edit.
- On the **Job History** pane, do one of the following:
 - To edit a single backup, right-click the backup.
 - To edit multiple backups, Shift + click or Ctrl + click the backups and then right-click one of the selected backups.

Note: You can only edit any backup jobs that have previously run from the **Storage** tab.

- Click **Edit Backup**.

2 In the **Backup** box, click **Edit**.



3 In the left pane, select the backup job setting that you want to modify.

The options vary depending on what you selected to back up.

4 Make any necessary changes.

- 5 When you are finished modifying any applicable options for the backup, click **OK**.
- 6 When you are finished editing the backup definition, click **OK** on the **Backup Properties** dialog box.

See [“Backing up data”](#) on page 158.

Testing credentials for backup sources

Backup Exec lists the backup sources and their associated credentials for each server that you monitor on the **Credentials** pane. If Backup Exec does not have the correct credentials to access content, any attempts to back up that content fail.

Symantec recommends that you test to make sure that you have the appropriate credentials to access the content that you want to back up. If a credentials test fails, you can enter new credentials for the content so that Backup Exec can access it.

You can test the credentials that are associated with a backup source at any time from the **Credentials** pane. You can also test the credentials when you create a backup definition.

Note: You cannot test credentials for virtual machines, but the job runs if you provide the correct credentials. If the job fails, you may need to retry the job with different credentials.

Complete any of the following procedures to test credentials for backup sources:

[To test the credentials that are associated with a backup source](#)

[To test the credentials for all of the backup sources in a specific backup definition](#)

[To test the credentials for a specific backup source in a specific backup definition](#)

To test the credentials that are associated with a backup source

- 1 On the **Backup and Restore** tab, double-click the server whose credentials you want to test.
- 2 In the left pane, click **Credentials**.
- 3 Select the backup source that you want to test.
- 4 In the **Credentials** group, click **Test Credentials**.

The **Credential Status** field displays the results of the test.

To test the credentials for all of the backup sources in a specific backup definition

- 1 Create a new backup or edit an existing backup.
- 2 In the **Selections** box, click **Test/Edit Credentials**.
- 3 Click **Test All**.

The **Credential Status** field displays the results of the test.

- 4 Click **OK**.

To test the credentials for a specific backup source in a specific backup definition

- 1 Create a new backup or edit an existing backup.
- 2 In the **Selections** box, click **Test/Edit Credentials**.
- 3 Select the backup source.
- 4 Click **Test Selected**.

The **Credential Status** field displays the results of the test.

- 5 Click **OK**.

See [“Replacing the credentials for a backup source”](#) on page 212.

See [“Creating new credentials for a backup source”](#) on page 213.

See [“Deleting retired or unused backup sources from the Credentials pane”](#) on page 214.

Replacing the credentials for a backup source

Backup Exec lists the backup sources and their associated credentials for each server that you monitor on the **Credentials** pane. If Backup Exec does not have the correct credentials to access content, any attempts to back up that content fail.

Symantec recommends that you test to make sure that you have the appropriate credentials to access the content that you want to back up. If a credentials test fails, you can enter new credentials for the content so that Backup Exec can access it. If you need to change the credentials that are associated with a backup source, you can do so at any time on the **Backup and Restore** tab from the **Credentials** pane. You can also change a backup source's credentials when you create or edit backups.

To replace the credentials for a backup source

- ◆ Do one of the following:

To replace the logon account that is associated with a backup source on the **Credentials** pane

Complete the following steps:

- On the **Backup and Restore** tab, double-click the server whose credentials you want to view.
- In the left pane, click **Credentials**.
- In the **Logon Account** field, select the logon account that you want to use for the backup source.
- Click **Apply**.

To replace the logon account that is associated with a backup source in a backup definition

Complete the following steps:

- Create a new backup or edit an existing backup.
- In the **Selections** box, click **Test/Edit Credentials**.
- In the **Logon Account** field, select the logon account that you want to use for the backup source.
- Click **OK**.

See [“Testing credentials for backup sources”](#) on page 211.

See [“Creating new credentials for a backup source”](#) on page 213.

See [“Deleting retired or unused backup sources from the Credentials pane”](#) on page 214.

Creating new credentials for a backup source

Backup Exec lists the backup sources and their associated credentials for each server that you monitor on the **Credentials** pane. If Backup Exec does not have the correct credentials to access content, any attempts to back up that content fail.

If the credentials for a backup source change, you must enter the new credentials in Backup Exec and associate them with the backup source. You can create new credentials for a backup source at any time on the **Credentials** pane. You can also create new credentials for a backup source when you create or edit backup definitions.

To create new credentials for a backup source

- 1 Do one of the following:

To create new credentials for a backup source on the **Credentials** pane

Complete the following steps:

- On the **Backup and Restore** tab, double-click the server for which you want to create new credentials.
- In the left pane, click **Credentials**.

To create new credentials for a backup source in a backup definition

Complete the following steps:

- Create a new backup or edit an existing backup.
- In the **Selections** box, click **Test/Edit Credentials**.

- 2 In the **Logon Account** field next to the backup source, select **<new logon account>**.
- 3 In the **User name** field, type the user name for the new logon account.
- 4 In the **Password** field, type the password for the new logon account.
- 5 In the **Confirm password** field, type the password again to confirm it.
- 6 In the **Account name** field, type a unique name for the new logon account.
- 7 In the **Notes** field, type any optional notes to explain how the Backup Exec logon account is used.
- 8 Select **This is a restricted logon account** if you want the Backup Exec logon account to be used only by the owner of the logon account and those who know the password.

If this option is not selected, the Backup Exec logon account is created as a common account. Common accounts are the shared accounts that all users can access.
- 9 Select **This is my default account** to make this account your default Backup Exec logon account, which is used to browse, make selections, or restore data on your local computers and remote computers
- 10 Click **OK**.

See [“Testing credentials for backup sources”](#) on page 211.

See [“Replacing the credentials for a backup source”](#) on page 212.

See [“Deleting retired or unused backup sources from the Credentials pane”](#) on page 214.

Deleting retired or unused backup sources from the Credentials pane

Backup Exec lists the backup sources and their associated credentials for each server that you monitor on the **Credentials** pane. If you remove one of the backup sources from your environment, it still displays in the list of backup sources on the **Credentials** pane. You can delete retired or unused backup sources from the pane to help make it easier to manage. Backup Exec browses and discovers resources during normal operations. If the backup source that you delete still exists in your

environment, Backup Exec will add it to the list of resources again the next time that it is discovered. You should only delete backup sources from the **Credentials** pane if they no longer exist in your environment.

To delete retired or unused backup sources from the Credentials pane

- 1 On the **Backup and Restore** tab, double-click the server that contains the backup source that you want to delete from the **Credentials** pane.
- 2 In the left pane, click **Credentials**.
- 3 Select the backup source that you want to delete from the **Credentials** pane.
- 4 In the **Credentials** group, click **Delete Selected Resource**.

Backup Exec removes the backup source from the list.

Note: If you try to delete a backup source that is used in a backup definition, Backup Exec does not delete that backup source. When you delete a backup source that still exists in your environment, Backup Exec repopulates the list with the backup source the next time it browses your environment.

See [“Testing credentials for backup sources”](#) on page 211.

See [“Replacing the credentials for a backup source”](#) on page 212.

See [“Creating new credentials for a backup source”](#) on page 213.

How job scheduling works in Backup Exec

Backup Exec lets you configure the time and the frequency for which you want to run backup jobs. You can run jobs immediately, once on a specific day and time, or more than once according to a schedule. Backup Exec lets you use minutes, hours, days, weeks, months, or years as measurements of time to create a recurring pattern for the schedule. Alternatively, you can select specific days of the month to create a recurring schedule on which jobs should run.

This topic includes the following information:

[Availability window](#)

[Scheduling conflicts](#)

[Including and excluding dates](#)

Availability window

Backup Exec has additional scheduling options that you can use to ensure that the job does not run outside of your availability window. The availability window is the

time range when your backup sources are available to be backed up. You can configure how long you want an incomplete job to remain scheduled before Backup Exec reschedules the job and changes its completion status to Missed. You can also configure the job to automatically be canceled if it is running for too long after its scheduled start time. These options help ensure that backup jobs do not affect your system resources during critical hours.

See [“Configuring schedules for backup jobs”](#) on page 519.

Scheduling conflicts

Backup Exec resolves scheduling conflicts between two backup jobs by running the less common job and skipping the more common job. If a full backup job and an incremental or differential backup job are scheduled to run at the same time, Backup Exec runs the full backup. The incremental backup or differential backup is skipped and it runs again at its next scheduled time. Full backup jobs always supersede incremental and differential backup jobs. If two backup jobs of the same type are scheduled to run at the same time, Backup Exec runs the less frequently scheduled job. The more frequently scheduled job is skipped. The skipped backup job runs again at its next scheduled time. For example, if a monthly job and a daily job are scheduled to run at the same time, Backup Exec runs the monthly job. The daily job is skipped and it runs on the following day per its normal schedule.

Including and excluding dates

The **Exclude Dates** option lets you exclude specific dates from the schedule. For example, you can exclude holidays from your normal backup schedule.

You can use the **Include Dates** option to include dates with your backup schedule. When you include a date with your backup schedule, the backup job runs on the included date, even if it is not normally scheduled to run on that date. You may want to schedule an include date to run an extra backup job, outside of the job's normal schedule.

See [“Preventing backup jobs from running on a specific date”](#) on page 218.

See [“Including a specific date in the schedule for a backup job”](#) on page 216.

Including a specific date in the schedule for a backup job

You can include specific dates with your backup schedule for a backup job. When you include a date with your backup schedule, the backup job runs on the included date, even if it is not normally scheduled to run on that date. You may want to schedule an include date to run an extra backup job, outside of the job's normal

schedule. Scheduling an include date does not affect a backup job's normal schedule.

Note: Included dates are applied to the job schedule before excluded dates. When any excluded dates are applied, they may overwrite the included dates if there are any conflicts. Therefore, if you select the same date as both an included date and an excluded date, Backup Exec excludes the date from your backup schedule.

See [“Preventing backup jobs from running on a specific date”](#) on page 218.

To include a specific date in the schedule for a backup job

- 1 Create or edit a backup definition.
- 2 In the **Backup** box, click **Edit**.
- 3 In the left pane, click **Schedule**.
- 4 Click the **Recurrence** field for the job to which you want to include a date.
- 5 On the **Include/exclude dates** tab, click **Include dates**.
- 6 Do one of the following:

To manually enter the date

Complete the following steps:

- In the **Select Date** field, type the date that you want to include with the backup schedule.
- Click **Add**.

Note: You can add only one date at a time.

To select the date from the calendar

Click the date that you want to include with the backup schedule.

The calendar displays 3 months at a time. You can navigate forward and backward to view additional months by clicking the arrows.

Note: You can select only one date at a time.

- 7 Click **OK**.

See [“How job scheduling works in Backup Exec”](#) on page 215.

Preventing backup jobs from running on a specific date

You can exclude specific dates, such as holidays, from your backup schedule for individual backup jobs.

When you exclude dates for a specific backup job, only that backup job is affected by the exclude date. The job does not run on the excluded date, even if it is normally scheduled to run. After the exclude date, the job resumes running on its normal schedule.

Note: Included dates are applied to the job schedule before excluded dates. When any excluded dates are applied, they may overwrite the included dates if there are any conflicts. Therefore, if you select the same date as both an included date and an excluded date, Backup Exec excludes the date from your backup schedule.

See [“Including a specific date in the schedule for a backup job”](#) on page 216.

To prevent backup jobs from running on a specific date

- 1 Create or edit a backup definition.
- 2 In the **Backup** box, click **Edit**.
- 3 In the left pane, click **Schedule**.
- 4 Click the **Recurrence** field for the job to which you want to add an exclude date.
- 5 On the **Include/exclude dates** tab, click **Exclude dates**.
- 6 Do one of the following:

To manually enter the date

Complete the following steps:

- In the **Select Date** field, type the date that you want to exclude from the backup schedule.
- Click **Add**.

Note: You can add only one date at a time.

To select the date from the calendar

Click the date that you want to exclude.

The calendar displays 3 months at a time. You can navigate forward and backward to view additional months by clicking the arrows.

Note: You can select only one date at a time.

7 Click **OK**.

See [“How job scheduling works in Backup Exec”](#) on page 215.

Viewing all scheduled backup jobs on a calendar

You can view all of your scheduled backup jobs for a month, for a week, or for a day on a calendar. It can be helpful to view your backup jobs in the calendar format to make sure there are no scheduling conflicts. You may want to check the calendar before you create a new job.

To view all scheduled backup jobs on a calendar

- 1 On the **Backup and Restore** tab, in the **Backups** group, click **Backup Calendar**.
- 2 When you are finished viewing the calendar, click **Close**.

See [“How job scheduling works in Backup Exec”](#) on page 215.

Adding a stage to a backup definition

Stages are the additional tasks that you can run with backup jobs as part of the backup definition. You may choose to add stages to the backup definition to customize it. You can add one or more stages for virtualization, duplication, and archiving when you create a backup definition. Stages can also be added to existing backup definitions.

For example, you may create a backup job that backs up any important data that must be sent off-site. You can add a duplicate stage to the backup definition that contains that job. The duplicate stage automatically sends the backup data to tape storage when the backup job is complete. Then you can take the tape off-site to ensure that your data is safe.

Table 4-7 Types of stages

Stage	Description
Duplicate to Disk	Creates a duplicate copy of your backup and sends it to disk storage.
Duplicate to Tape	Creates a duplicate copy of your backup and sends it to tape storage.
Convert to Virtual Machine After Backup	Creates a virtual machine from your backup sets after the backup job is complete. See “How conversion of physical computers to virtual machines works in Backup Exec” on page 479.
Convert to Virtual Machine Simultaneously with Backups	Creates a virtual machine from your backup sets while the backup job is running. See “How conversion of physical computers to virtual machines works in Backup Exec” on page 479.
Archive	Creates an archive job to archive the selected file system shares and folders and Exchange mailboxes to a vault store. See “About the Archiving Option” on page 1313.

To add a stage to a backup definition

- 1 Create or edit a backup definition.

Note: You cannot add a stage to one-time backups.

- 2 In the **Backup** box, click **Add Stage**.
- 3 Select the type of stage that you want to add.

You can add multiple stages to a backup definition.
- 4 In the stage box, click **Edit**.
- 5 In the left pane, click **Schedule**, and then select when you want the stage to run.
- 6 In the left pane, click **Storage**, and then select the storage device that you want to use for this stage.

- 7 Select any additional options that apply to this stage.
 - 8 When you are finished selecting all applicable options for this stage, click **OK**.
Repeat this procedure for each stage that you want to add to the backup definition.
- See [“Backing up data”](#) on page 158.
- See [“Editing a stage”](#) on page 221.

Editing a stage

You can edit a stage that is part of a backup definition.

To edit a stage

- 1 Do one of the following:

To edit a stage from the **Backup and Restore** tab

Complete the following steps:

- On the **Backup and Restore** tab, right-click the server that contains the backup definition with the stage that you want to edit.
- Click **Edit Backups**.
- If the server has more than one backup definition, select the definition that contains the stage that you want to edit, and then click **OK**.

To edit a stage from the **Job Monitor** tab

Complete the following steps:

- On the **Job Monitor** tab, right-click the job that contains the stage that you want to edit.
- Click **Edit**.

- 2 Click **Edit** in the box that contains the stage that you want to edit.
- 3 Make any necessary changes.
- 4 When you are finished making changes to the stage, click **OK** on the **Backup Properties** dialog box.

See [“Adding a stage to a backup definition”](#) on page 219.

Duplicating backup sets or a job history manually

You can configure a duplicate stage to automatically duplicate backup data after a backup job runs. You may want to duplicate data to have an extra copy to send off-site, for example.

See [“Adding a stage to a backup definition”](#) on page 219.

You can also manually duplicate backup data from completed jobs at any time. When you duplicate backed up data from completed jobs, you select the backup sets or job history that you want to duplicate. The data that you select is read from the source and written to the selected destination, such as a drive, drive pool, or backup folder. You can encrypt the duplicated data. You can schedule when this type of job runs, but it only runs one time.

You can select to duplicate one or more individual backup sets or you can duplicate an entire job history. You should duplicate backup sets if you want to duplicate only the data that was backed up in a specific backup job instance. When you duplicate a job history, Backup Exec includes all of a backup definition's dependent backup sets. For example, if you select to duplicate an incremental backup job, Backup Exec automatically duplicates all incrementals dating back to, and including, the last full backup job.

This topic includes the following information:

[Duplicating data from a virtual device to a physical device](#)

[To duplicate backup sets](#)

[To duplicate job history](#)

Duplicating data from a virtual device to a physical device

You can use a duplicate backup job to copy data directly from a virtual device to a physical device. Software encryption cannot be applied to a duplicate backup job when you copy data directly from a virtual device to a physical device. You must either disable DirectCopy or select not to encrypt the job.

See [“Copying data from a virtual tape library to a physical tape device using DirectCopy to tape”](#) on page 229.

If you duplicate any Oracle backup sets that were created with multiple data streams, note the following:

- Backup Exec converts the multiple data streams to a sequential data stream during the duplication job.
- A restore job from the duplicated copy may be slower than a restore job from the original media.

To duplicate backup sets

- 1 On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the backup sets that you want to duplicate.
- 2 In the left pane, click **Backup Sets**.

- 3 Do one of the following:
 - To duplicate a single backup set, right-click the backup set.
 - To duplicate multiple backup sets, Shift + click or Ctrl + click the backup sets, and then right-click one of the selected backup sets.
- 4 Select **Duplicate**.
- 5 In the **Schedule** group box, select when you want Backup Exec to run the duplicate job:

To run the duplicate job immediately	Select Run now .
To schedule the job for a later time	Select Run on , and then enter the date and time.
To schedule the job to run later using an external scheduling tool	Select Create without a schedule . You can schedule the job to run later using an external scheduling tool.
To submit the job on hold	Click Submit job on hold . The job is created with an On Hold status. It remains on hold until you remove the hold on it.
- 6 In the **Storage** field, select the storage device to which you want to copy the backup sets.
- 7 In the **Keep for** field, select the amount of time that you want Backup Exec to keep the backup sets.

The media is protected from being overwritten for the amount of time that you specify.

Select **Use source retention** to keep the backup sets for the same amount of time as the source backup sets that you want to duplicate.
- 8 Do any of the following:

To enable compression for the duplicate backup sets	In the Compression field, select the type of compression.
---	--

To enable encryption for the duplicate backup sets

Complete the following steps:

- In the **Encryption type** field, select the type of encryption.
- In the **Encryption key** field, select the encryption key that you want to use or select **Manage keys** to create a new key.

To run a verify operation on the duplicate backup sets

Select **Verify at the end of the job**.

9 On the **Duplicate Job** dialog box, click **OK**.

To duplicate job history

1 Do one of the following:

To duplicate job history from the **Backup and Restore** tab or the **Storage** tab

Complete the following steps:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the job history that you want to duplicate.
- In the left pane, click **Job History**.
- Do one of the following:
 - To duplicate a single job history, right-click the job history.
 - To duplicate multiple job histories, Shift + click or Ctrl + click the job histories, and then right-click one of the selected job histories.

To duplicate job history from the **Job Monitor** tab

Complete the following steps:

- On the **Job Monitor** tab, do one of the following:
 - To duplicate a single job history, right-click the job history.
 - To duplicate multiple job histories, Shift + click or Ctrl + click the job histories, and then right-click one of the selected job histories.

2 Click **Duplicate**.

- 3 In the **Schedule** group box, select when you want Backup Exec to run the duplicate job:

To run the duplicate job immediately	Select Run now .
To schedule the job for a later time	Select Run on , and then enter the date and time.
To schedule the job to run later using an external scheduling tool	Select Create without a schedule . You can schedule the job to run later using an external scheduling tool.
To submit the job on hold	Click Submit job on hold . The job is created with an On Hold status. It remains on hold until you remove the hold on it.

- 4 In the **Storage** field, select the storage device to which you want to copy the job history.

- 5 In the **Keep for** field, select the amount of time that you want Backup Exec to keep the job history.

The media is protected from being overwritten for the amount of time that you specify.

Select **Use source retention** to keep the job history for the same amount of time as the source job history that you want to duplicate.

- 6 Do any of the following:

To enable compression for the duplicate job history	In the Compression field, select the type of compression.
To enable encryption for the duplicate job history	Complete the following steps: <ul style="list-style-type: none"> In the Encryption type field, select the type of encryption. In the Encryption key field, select the encryption key that you want to use or select Manage keys to create a new key.
To run a verify operation on the duplicate job history	Select Verify at the end of the job .

- 7 On the **Duplicate Job** dialog box, click **OK** for each duplicate job that you create.

Running a test run job manually

Test run jobs attempt to determine if a scheduled backup could possibly fail when you run it. When you run a test job, no data is backed up. Instead, Backup Exec checks your storage capacity, credentials, and media to find potential errors. If there is an error, the job continues to run until it is completed. The error appears in the job log. You can also configure Backup Exec to send a notification to a designated recipient.

During a test run job, the following things may cause a job to fail:

- Logon credentials are incorrect.
- Storage capacity is not sufficient.
- Tape cartridge media or disk cartridge media is not available.
- Overwritable media is not available for an overwrite job.
- Appendable media is not available for an append job.

A test run job checks the media capacity that is available for the selected job. However, you can check if there is enough available media for multiple test run jobs in the Test Run Results Report.

See [“Test Run Results report”](#) on page 683.

You can configure test run jobs to run automatically before your scheduled backup jobs. Or you can manually run a test run job at any time.

Before you run a test run job, Symantec recommends that you run backup jobs to your storage devices first. Backup Exec does not recognize the capacity of a storage device until an actual backup job sends data to the device. If you create a test run job before any other jobs, Backup Exec cannot check that the device has sufficient capacity to perform the backup job. After at least one backup job has sent data to a device, Backup Exec can determine the capacity.

To run a test job manually

- 1 Do one of the following:

To run a test job from the
Backup and Restore tab

Complete the following steps:

- On the **Backup and Restore** tab, double-click the server that contains the job you want to test.
- In the left pane, click **Jobs**.

To run a test job from the **Job Monitor** tab, select the **Job Monitor** tab.

- 2 Right-click the job that you want to test, and then click **Test Run**.
 - 3 Click **Yes** to confirm that you want to run the test job now.
- See [“Configuring automatic test run jobs for backup jobs”](#) on page 533.

Verifying backed up data manually

Backup Exec can perform a verify operation to make sure that the media can be read after a backup job has been completed. Symantec recommends that you verify all backed up data to ensure the integrity of the collection of data and the media on which it resides. By default, Backup Exec automatically verifies backed up data at the end of a backup job. However, you can also schedule the verify operation to take place at a later time or disable the verify operation altogether. You can change Backup Exec's verify options as part of the default backup settings or for individual backup jobs.

See [“Configuring automatic verify operations for backup jobs”](#) on page 534.

You can also choose to manually run a verify operation on a backup set or a job history at any time. You can verify backup sets if you want to verify only the data that was backed up in a specific backup job instance. If you want to verify a backup definition and all of its dependent backup sets, you can verify a job history. For example, if you want to verify a backup definition that used incremental backups, Backup Exec verifies all incrementals dating back to, and including, the last full backup.

This topic includes the following procedures:

[To verify specific backup sets](#)

[To verify a job history](#)

To verify specific backup sets

- 1 On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the backup set or backup sets that you want to verify.
- 2 In the left pane, click **Backup Sets**.
- 3 Do one of the following:
 - To verify a single backup set, right-click the backup set.

- To verify multiple backup sets, Shift + click or Ctrl + click the backup sets, and then right-click one of the selected backup sets.

4 Click **Verify**.

5 In the **Schedule** group box, select when you want Backup Exec to run the verify operation:

To run the verify operation immediately Select **Run now**.

To schedule the operation for a later time Select **Run on**, and then enter the date and time.

To schedule the job to run later using an external scheduling tool Select **Create without a schedule**.
You can schedule the verify operation to run later using an external scheduling tool.

6 Click **OK**.

To verify a job history

1 Do one of the following:

To verify a job history from the **Backup and Restore** tab or the **Storage** tab

Complete the following steps:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the job history that you want to verify.
- In the left pane, click **Job History**.
- Do one of the following:
 - To verify a single job history, right-click the job history.
 - To verify multiple job histories, Shift + click or Ctrl + click the job history, and then right-click one of the selected job histories.

To verify a job history from the **Job Monitor** tab

On the **Job Monitor** tab, do one of the following:

- To verify a single job history, right-click the job history.
- To verify multiple job histories, Shift + click or Ctrl + click the job histories, and then right-click one of the selected job histories.

2 Click **Verify**.

Copying data from a virtual tape library to a physical tape device using DirectCopy to tape

- 3 In the **Schedule** group box, select when you want Backup Exec to run the verify operation:

To run the verify operation immediately Select **Run now**.

To schedule the operation for a later time Select **Run on**, and then enter the date and time.

To schedule the job to run later using an external scheduling tool Select **Create without a schedule**.
You can schedule the verify operation to run later using an external scheduling tool.

- 4 Click **OK**.

Copying data from a virtual tape library to a physical tape device using DirectCopy to tape

Backup Exec's **DirectCopy to tape** option enables data to be copied from a virtual tape library directly to a physical tape device during a duplicate backup job. The Backup Exec server coordinates the duplicate job, but it does not copy the data. Instead, the virtual tape library copies the virtual tape image directly to the physical device. The Backup Exec server records information about the data in the catalog. Because the information about the copied data is in the catalog, you can restore data from either the virtual tape library or the physical device. The job log for the duplicate backup job indicates that **DirectCopy to tape** is enabled.

To use **DirectCopy to tape**, both the source device and the destination device must be NDMP-enabled. If the devices are not NDMP-enabled, then Backup Exec performs a regular duplicate backup job.

Both hardware encryption and software encryption are supported with **DirectCopy to tape**. For software encryption, both the source backup set and the destination backup set must use software encryption.

Note: If you select disk storage as the destination device for a duplicate job with **DirectCopy to tape** enabled, Backup Exec performs a regular duplicate job.

Copying data from a virtual tape library to a physical tape device using DirectCopy to tape

Table 4-8 How to use DirectCopy to tape to copy data from a virtual tape library to a physical device

Step	Notes	For more information
Create a regular backup job.	Select a virtual tape library as the storage destination.	See “Backing up data” on page 158. See “Configuring storage options for backup jobs” on page 527.
Create a duplicate backup job.	In the DBA-initiated job settings: <ul style="list-style-type: none"> ■ Select a physical tape device as the destination. ■ Select Enable DirectCopy to tape. 	See “DBA-initiated job templates” on page 597. See “Storage options for DBA-initiated jobs” on page 600. See “Duplicate job settings for DBA-initiated jobs” on page 606.

Restores

This chapter includes the following topics:

- [Methods for restoring data in Backup Exec](#)
- [Searching for data to restore](#)
- [Restoring data from a server, a backup set, a backup job, or a storage device](#)
- [Restoring file system data](#)
- [Performing a complete online restore of a Microsoft Windows computer](#)
- [Restoring System State](#)
- [Installing a new Windows Server domain controller into an existing domain by using a redirected restore](#)
- [Restoring Backup Exec Shadow Copy Components](#)
- [Restoring utility partitions or Unified Extensible Firmware Interface system partitions](#)
- [About restoring encrypted data](#)
- [About restoring NetWare SMS volume backups to non-SMS volumes with Backup Exec](#)
- [Restoring data from ARCserve tapes](#)
- [Canceling a restore job](#)
- [How Backup Exec catalogs work](#)
- [Cataloging backup sets](#)

Methods for restoring data in Backup Exec

Backup Exec provides guided **Search** and **Restore** methods to assist you when you search for or restore backed up data. If you have installed the Archiving Option, you can also search for and restore archived data.

From **Search** or **Restore** on the **Backup and Restore** tab, you can do the following:

- Restore data to the location from which it was originally backed up or redirect the restore to another location.
- Start the restore job immediately or schedule it to run at a future time.

Note: The **Restore Wizard** displays only up to 30,000 items. If you want to restore data from a folder that contains more than 30,000 items, you should search for the item that you want to restore. You can use search criteria such as the type of data and the date of the original backup to help reduce the number of items that displays.

Table 5-1 Guided methods to search for and restore data

Method	Description
Search	<p>Lets you select multiple servers on which to search for backup sets. Then, you can choose to restore the data, or you can copy and save the search criteria and the results to the clipboard. You can then email the results to the person who requested the restore to ensure that you have found the correct data before you restore it.</p> <p>To start the Search Wizard, on the Backup and Restore tab, select one or more servers, and then in the Restores group, click Search.</p> <p>See “Searching for data to restore” on page 234.</p>

Table 5-1 Guided methods to search for and restore data (*continued*)

Method	Description
Restore	<p>Lets you browse the backup sets from a single server, and then restore the data. You can restore file system data, System State data, Backup Exec Shadow Copy Components, utility partitions or UEFI system partitions, and more.</p> <p>You can also perform a complete online restore of a Windows computer if the computer was fully selected for a backup. By default, backup jobs include all necessary components that are required for a complete restore.</p> <p>To start the Restore Wizard, on the Backup and Restore tab, select a server, and then in the Restores group, click Restore.</p> <p>See “Restoring data from a server, a backup set, a backup job, or a storage device” on page 234.</p> <p>See “Performing a complete online restore of a Microsoft Windows computer” on page 239.</p>
Simplified Disaster Recovery	<p>Lets you recover Windows computers after a hard drive failure. The Simplified Disaster Recovery wizards guide you in preparing for disaster recovery, and in recovering a local computer or a remote computer to its pre-disaster state.</p> <p>See “About Simplified Disaster Recovery” on page 736.</p>

See [“Restoring file system data”](#) on page 238.

See [“Restoring System State ”](#) on page 240.

See [“Restoring Backup Exec Shadow Copy Components”](#) on page 244.

See [“Restoring utility partitions or Unified Extensible Firmware Interface system partitions”](#) on page 245.

See [“Restoring Exchange data”](#) on page 962.

See [“Restoring SQL databases and transaction logs”](#) on page 921.

See [“About restoring Oracle resources”](#) on page 1012.

See [“Restoring VMware virtual machines and vmdk files”](#) on page 852.

See [“Restoring a deduplication disk storage device or deduplicated data”](#) on page 827.

See [“About restoring items from the archives”](#) on page 1374.

See [“About restoring an Archiving Option component”](#) on page 1382.

See [“Restoring Microsoft SharePoint data”](#) on page 981.

See [“About restoring Enterprise Vault”](#) on page 1031.

See [“About restoring data to Linux computers”](#) on page 1207.

See [“About restoring Lotus Domino databases”](#) on page 1073.

Searching for data to restore

You can select one server or multiple servers on which to search for backup sets. Then, you can choose to restore the data, or you can copy and save the search criteria and the results to the clipboard. You can then email the results to the person who requested the restore to ensure that you have found the correct data before you restore it.

Backup Exec creates separate restore jobs for each server that you restore data to.

The **Search Wizard** supports only the following types of data:

- Files and folders
- Exchange and SharePoint backup sets for which Granular Recovery Technology was enabled
- Archived file system data and archived Exchange mailbox data

Note: You cannot search the backup sets for any data that was backed up from virtual machines.

To search for data to restore

- 1 On the **Backup and Restore** tab, right-click the server or servers on which you want to search for data, and then click **Search**.
- 2 Follow the **Search Wizard** prompts to search for and restore data.

Restoring data from a server, a backup set, a backup job, or a storage device

You can restore data by launching the **Restore Wizard** directly from a server, from backup sets, from completed backup jobs, and from storage device media.

You can also perform a complete online restore of a Windows computer if the computer was fully selected for a backup. By default, backup jobs include all necessary components that are required for a complete restore.

See [“Performing a complete online restore of a Microsoft Windows computer”](#) on page 239.

Table 5-2 Restore methods

Restore method	More information
To restore from a server	Restoring data from a server
To restore from a backup set	To restore data from a backup set
To restore from a completed backup job	Restoring data from a completed backup job
To restore from storage device media	Restoring data from storage device media

Restoring data from a server

You can browse the backup sets from a single server, and then restore the data.

Note: If you back up and then rename a server, the new server name and the old server name both appear on the **Backup and Restore** tab. The status next to the new server name indicates that it is backed up. However, you should select the icon with the old server name to restore any data that you backed up before you changed the server name.

To restore data from a server

- 1 On the **Backup and Restore** tab, in the list of servers, or on the **Job Monitor** tab, right-click a server that has been backed up.
- 2 Click **Restore**.
- 3 Follow the **Restore Wizard** prompts to restore the data.

Restoring data from a completed backup job

You can restore data from a completed backup job. When you restore data from backup jobs, you can only choose data that is contained in the backup sets that the backup job produced.

To restore data from a completed backup job

- 1 Do one of the following:

To restore data from the **Backup and Restore** tab

- On the **Backup and Restore** tab, in the list of servers, double-click a server that has been backed up.
- In the **Jobs** view, expand a backup definition.
- Under the backup definition, right-click the backup job from which you want to restore data.
- Select **Restore backup sets created by this job**.

To restore data from the **Job Monitor** tab

- On the **Job Monitor** tab, in the **Jobs** pane, right-click a server that has been backed up.
- Select **Restore backup sets created by this job**.

- 2 Follow the **Restore Wizard** prompts to restore the data.

Restoring data from a backup set

You can restore data from the backup sets that appear on the **Backup and Restore** tab. When you restore data from backup sets, you can only choose data that is contained in the backup sets.

To restore data from a backup set

- 1 On the **Backup and Restore** tab, in the list of servers, double-click a server that has been backed up.
- 2 In the **Backup Sets** view, expand a backup set from which you want to restore.
- 3 Right-click the backup set from which you want to restore data
- 4 Click **Restore**.
- 5 Follow the **Restore Wizard** prompts to restore the data.

Restoring data from storage device media

You can restore data from the media that are contained in the storage devices that appear on the **Storage** tab. When you restore data directly from storage device media, you can only restore data that is contained in the backup sets on the media. The Restore Wizard prompts you to select options for each type of data, and then submits a separate job for each type of data.

When you restore from storage device media, you can restore multiple types of data. Separate restore jobs are submitted for each type of data.

To restore data from storage device media

- 1 Select from the following:

To restore from a disk storage device Do the following in the order listed:

- On the **Storage** tab, in the **All Storage** view, navigate to the storage device from which you want to restore.
- Double-click the disk storage device, and then click **Backup Sets** on the left.
- Right-click the backup set from which you want to restore, and then click **Restore**.

To restore from tape drive media or a disk cartridge media Do the following in the order listed:

- On the **Storage** tab, in the **All Storage** view, expand **Tape and Disk Cartridge Media**.
- Navigate to the media that contains the data that you want to restore.
- Right-click the media from which you want to restore, and then click **Restore**.

To restore from media within a robotic library Do the following in the order listed:

- On the **Storage** tab, in the **All Storage** view, navigate to the robotic library from which you want to restore.
- Expand the robotic library.
- Double-click **Slots**.
- In the Slots view, right-click the slot that contains the media from which you want to restore, and then click **Restore**.

- 2 Follow the **Restore Wizard** prompts to restore the data.
- 3 (Optional) To restore multiple types of data, proceed through the **Restore Wizard** and select the appropriate options for the type of data that you select.

Then, do the following:

- On the summary page, click **Continue** to submit the job and return to the page that lists the different types of data.
The **Job Submitted** column displays an icon to indicate that the restore job for the first type of data has been submitted.
- Select the next type of data and the appropriate restore options for that type of data. If you do not want to restore the other types of data, you can click **Cancel** to exit the **Restore Wizard** without affecting any of the jobs that you already submitted.

- After you configure jobs for each type of data, the **Continue** button changes to a **Finish** button. When you click **Finish**, Backup Exec submits the final restore job, and then closes the **Restore Wizard**.

Restoring file system data

When you restore files, folders, or volumes, you can restore to a point-in-time or you can restore from a backup set.

If you know which backup set, completed backup job, or storage device media contains the data that you want to restore, you can start the **Restore Wizard** from there. Otherwise, you can select a server and then start the **Restore Wizard**.

See [“Restoring data from a server, a backup set, a backup job, or a storage device”](#) on page 234.

To restore file system data

- 1 On the **Backup and Restore** tab, right-click the server for which you want to restore data, and then click **Restore**.
- 2 Select **Files, folders, or volumes**, and then click **Next**.

3 Do one of the following:

To restore the data to a point-in-time at which the backup set was created

Select **File and folder backups to a point-in-time**.

Note: When you restore file system data to a point-in-time, you select a backup set to restore. The backup set represents the file system data at the specific point-in-time at which it was backed up. Backup Exec automatically restores any related backup sets that are required to restore the file system data to its state at the point-in-time.

To restore file and folder backups from a backup set

Select **File and folder backups from a backup set**.

Note: When you restore file system data from a backup set, you select a backup set to restore. Backup Exec restores the file system data exactly as it exists in the backup set that you select. No dependent backup sets are selected for the restore.

To search for files and folders

Select **Files and folders located through Search**.

4 Click **Next**, and then follow the **Restore Wizard** prompts to restore the data.

Performing a complete online restore of a Microsoft Windows computer

You can perform a complete online restore of a Microsoft Windows computer if the computer was fully selected for a backup. You select the backup set time from which you want to recover the computer. All required backup sets are automatically selected. You can select additional backup sets to restore as appropriate. You cannot redirect the restore of the computer.

To run a complete online restore of a Microsoft Windows computer

- 1 On the **Backup and Restore** tab, right-click the computer that you want to restore, and then click **Restore**.
- 2 Select **A Microsoft Windows computer that was fully selected for a backup**, and then click **Next**.

3 Select **Complete online restore of a computer, or restore system components**, and then click **Next**.

4 Follow the **Restore Wizard** prompts to restore the data.

See [“About backing up critical system components”](#) on page 183.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

Restoring System State

Depending on the version of Microsoft Windows, service pack levels, and features that are installed, you can restore the following system state data:

- Active Directory
- Automated system recovery
- Background Intelligent Transfer Service
- COM+ Class Registration database
- Dynamic Host Configuration Protocol
- Event logs
- File Server Resource Manager
- Internet Information Service (IIS)
- Microsoft Search Service
- Network Policy Server
- Registry
- Remote Storage
- Removable Storage Manager
- Shadow Copy Optimization Writer
- System files
- Terminal Server Licensing
- Terminal Services Gateway
- Windows Deployment Services
- Windows Management Instrumentation

If the server is a certificate server, then System State includes the Certificate Services database.

If the server is a domain controller, then System State includes the Active Directory services database and the SYSVOL directory.

See [“About the Agent for Microsoft Active Directory”](#) on page 1086.

You must restart the computer after you restore System State data.

Warning: You should not cancel a System State restore job. Canceling this job can leave the server unusable.

If you know which backup set, completed backup job, or storage device media contains the data that you want to restore, you can start the **Restore Wizard** from there. Otherwise, you can select a server and then start the **Restore Wizard**.

See [“Restoring data from a server, a backup set, a backup job, or a storage device”](#) on page 234.

To restore System State

- 1 On the **Backup and Restore** tab, right-click the computer for which you want to restore System State, and then click **Restore**.
- 2 Select **Complete online restore of a computer, or restore system components**, and then click **Next**.
- 3 Follow the **Restore Wizard** prompts to restore the data.
- 4 After you restore System State data, restart the computer.

See [“About backing up critical system components”](#) on page 183.

See [“Performing a complete online restore of a Microsoft Windows computer”](#) on page 239.

Restoring System State to a domain controller

To restore System State to a computer that is a domain controller, you must start the computer in safe mode. Then, use the Directory Services Restore Mode to perform the restore.

To replicate Active Directory to the other domain controllers that exist in the domain, you must perform an authoritative restore of the Active Directory. An authoritative restore ensures that the restored data is replicated to all of the servers. Performing an authoritative restore includes running Microsoft's Ntdsutil utility after Backup Exec restores System State, but before you restart the server. For more information about authoritative restore and the Ntdsutil utility, see your Microsoft documentation.

To restore System State to a domain controller

- 1 Retart the destination server in Directory Services Restore Mode.
See your Microsoft documentation for details on how to do this.
- 2 Open the services.
- 3 For each Backup Exec service listed, do the following in the order listed:
 - Click **Properties**.
 - On the **Log On** tab, click **This account**.
 - Enter a user account with local administrator's rights, and then click **OK**.
 - Right-click the service, and then click **Start**.
- 4 After the Backup Exec services have started, run the **Restore Wizard** to restore System State.
See [“Methods for restoring data in Backup Exec”](#) on page 232.
- 5 In the **Restore Wizard**, enable the option **Mark this server as the primary arbitrator for replication when restoring SYSVOL in System State**.
- 6 Restart the server before you restore more data.

Installing a new Windows Server domain controller into an existing domain by using a redirected restore

To install a new Windows Server domain controller into an existing domain, the Active Directory and SYSVOL data must be replicated to the new domain controller. If there is a large amount of data to be replicated or if the connection between the domain controllers is slow, the replication time can be lengthy. The amount of data to be replicated and the connection speed also affects the Active Directory Application Mode replication time. To decrease the replication time for Active Directory and Active Directory Application Mode, you can use the Microsoft Windows feature called **Install from Media**.

For Active Directory, use the **Install from Media** feature. Restore the system state backup sets of an existing domain controller in the domain in which you want to add a new domain controller. Then, perform a redirected restore of the system state backup sets to the destination domain controller.

For Active Directory Application Mode, you can back up data using the ADAM Writer. Then, you can perform a redirected restore of the data from the ADAM backup to the destination computer.

See [“About the Agent for Microsoft Active Directory”](#) on page 1086.

For more information, refer to your Microsoft documentation.

Table 5-3 How to install a new Windows Server domain controller into an existing domain by using a redirected restore

Step	Description
Step 1	<p>Back up the System State data of an active Windows Server domain controller that is in the target domain. You should back up the data to some type of removable storage, such as a disk cartridge device or a tape.</p> <p>See "Backing up data" on page 158.</p>
Step 2	<p>Attach the storage that contains the System State data to the computer that you want to install into the destination domain.</p> <p>Note: Symantec recommends that you encrypt the storage. Use caution when transporting it to the location of the destination domain.</p>
Step 3	<p>Inventory and catalog the storage.</p> <p>See "Inventorying and cataloging a storage device" on page 454.</p>
Step 4	<p>Redirect the restore of the system state backup sets to a temporary location on a volume or directory on the destination computer.</p> <p>See "Restoring System State " on page 240.</p>
Step 5	<p>Start the domain controller installation by doing the following in the order listed:</p> <ul style="list-style-type: none">■ On the destination computer, click Start, and then click Run.■ Type <code>dcpromo /adv</code>, and then click OK.■ Click Next when the Active Directory Installation Wizard appears.■ Select Additional domain controller for an existing domain, and then click Next.■ Select From these restored backup files, enter the temporary location to which you redirected the restore of the System State data, and then click Next.■ Complete the Active Directory Installation Wizard by following the prompts on the screen.
Step 6	<p>Complete the domain controller installation.</p>
Step 7	<p>Restart the computer that has the new domain controller.</p>

Table 5-3 How to install a new Windows Server domain controller into an existing domain by using a redirected restore (*continued*)

Step	Description
Step 8	Expire any remaining system state backup sets that you redirected to the temporary location. See “Changing the expiration date of backup sets on disk-based storage” on page 345.

Restoring Backup Exec Shadow Copy Components

The Backup Exec Shadow Copy Components file system uses Microsoft’s Volume Shadow Copy Service to protect third-party application and user data on Windows computers. You can restore the items in Backup Exec Shadow Copy Components individually or together.

The following items are contained in Backup Exec Shadow Copy Components:

- Backup Exec Deduplication Disk Storage
- Distributed File System Replication (DFSR)
- OSISoft PI Server data

When you restore domain-based Microsoft Distributed File System (DFS) namespaces, you must also restore the Active Directory to the same point-in-time. Domain-based DFS namespaces reference information that resides in the Active Directory. If you restore the namespaces without restoring Active Directory to the same point-in-time, that information is not restored and you may receive errors in the DFS Management Console.

If you know which backup set, completed backup job, or storage device media contains the data that you want to restore, you can start the **Restore Wizard** from there. Otherwise, you can select a server and then start the **Restore Wizard**.

See [“Restoring data from a server, a backup set, a backup job, or a storage device”](#) on page 234.

To restore Backup Exec Shadow Copy Components

- 1 On the **Backup and Restore** tab, right-click the server, and then click **Restore**.
- 2 Select **Shadow Copy Components**, and then click **Next**.
- 3 Follow the **Restore Wizard** prompts to restore the data.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

See [“Restoring System State ”](#) on page 240.

Restoring utility partitions or Unified Extensible Firmware Interface system partitions

You can select utility partitions or Unified Extensible Firmware Interface (UEFI) partitions for restore. Utility partitions are small partitions that OEM vendors such as Dell, Hewlett-Packard, and IBM install on the disk. These partitions contain system diagnostic and configuration utilities. UEFI partitions are the small partitions that the operating system creates. The UEFI system partitions contain the critical system files, such as bootmgr and BOOT\BCD files.

Requirements for restoring utility partitions are as follows:

- You must recreate the utility partitions before you restore any data.
- You must have Administrator rights.
- You cannot redirect the restore of a utility partition to another computer.
- You can only restore the utility partitions that belong to the same vendor. For example, you cannot restore Dell utility partitions to a Compaq utility partition.
- The size of the utility partition to which you restore the data must be equal to or greater in size than the utility partition that was backed up.

If you know which backup set, completed backup job, or storage device media contains the data that you want to restore, you can start the **Restore Wizard** from there. Otherwise, you can select a server and then start the **Restore Wizard**.

See [“Restoring data from a server, a backup set, a backup job, or a storage device”](#) on page 234.

To restore utility partitions or UEFI system partitions

- 1 On the **Backup and Restore** tab, right-click the computer for which you want to restore utility partitions or UEFI system partitions, and then click **Restore**.
- 2 Select one of the following, and then click **Next**:
 - **Utility partition**
 - **Unified Extensible Firmware Interface system partition**
- 3 Follow the **Restore Wizard** prompts to restore the data.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

About restoring encrypted data

Encrypted backup sets are identified in the restore selection list by an icon with a lock on it. When you select encrypted data to restore, Backup Exec automatically

validates the encryption key for the data. If the encryption key that was used to back up the data is still in the Backup Exec Database, then Backup Exec selects that encryption key automatically. However, if the encryption key cannot be located, Backup Exec prompts you to provide the pass phrase for the encryption key that was used to back up the data. If you enter the correct pass phrase, Backup Exec recreates the key.

When you use a restricted encryption key to back up data, users other than the key owner must enter the pass phrase to restore data.

See [“Using encryption with Backup Exec”](#) on page 583.

See [“Encryption key management”](#) on page 586.

About restoring NetWare SMS volume backups to non-SMS volumes with Backup Exec

Backup Exec supports restoring NetWare SMS volume backups to non-SMS volumes. For example, the data that is backed up with Backup Exec for NetWare Servers or Novell's SBackup can be restored to the Backup Exec server or to another network share.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Restoring data from ARCserve tapes

You can restore data from ARCserve tapes.

The following types of data cannot be restored from ARCserve tapes:

- Databases, such as Microsoft SQL and Exchange Server and NetWare Directory Services
- Windows registry
- Interleaved files
- Compressed files
- Encrypted files
- Long filenames and Extended Attributes for OS/2 files
- Long filenames and resource forks for Macintosh files

Storage that contains ARCserve backups can be overwritten; however, backup append jobs are not supported. All Backup Exec storage utility functions can be performed on ARCserve media.

Note: If the ARCserve backup spans multiple tapes, you must have all the tapes that were included in the ARCserve backup available. Make sure you start both the catalog job and the restore job with the first tape used in the ARCserve backup.

Table 5-4 Restoring data from ARCserve tapes

Step	Action
Step 1	<p>Inventory and catalog all of the storage from the ARCserve backup.</p> <p>See “Inventorying and cataloging a storage device” on page 454.</p> <p>During cataloging, Backup Exec reports file formats that it can read. Files that cannot be read do not appear in the catalogs. The media description that appears in the Backup Exec catalog comes from the session description that ARCserve uses.</p> <p>Storage-based catalogs are not supported on storage that is created by other vendors’ backup products. Therefore, cataloging ARCserve storage takes considerably longer than cataloging Backup Exec storage.</p>
Step 2	<p>Restore the selected data to a computer.</p> <p>See “Methods for restoring data in Backup Exec” on page 232.</p> <p>Due to the naming conventions that ARCserve uses for some computers, it may be necessary to redirect the data that you want to restore to another computer.</p>

Canceling a restore job

Warning: Canceling a restore job while it is in progress may result in unusable data and can leave the disk in an unusable state.

To avoid canceling a restore job, you can redirect the restore to a noncritical destination. Then, copy the data to a final destination when the job completes successfully.

You should not cancel a System State restore job. Canceling a System State restore job can leave the system unusable.

To cancel a restore job

- 1 On the **Job Monitor**, right-click the job that you want to cancel.
- 2 Click **Yes** when you are prompted if you are sure that you want to cancel the job.

See [“Canceling an active job”](#) on page 257.

How Backup Exec catalogs work

While backing up data, Backup Exec creates a catalog that contains information about the backup sets and about the storage device on which the backup sets are stored. When you select data to restore, Backup Exec uses the catalog information to find the restore selections and the storage devices on which they reside.

When a storage device is fully cataloged, you can do the following:

- View information on all the directories and files that are contained in each backup set.
- Search for files to restore.

Backup Exec fully catalogs each backup job. However, if the catalogs are truncated, only backup set information is listed. You cannot view files or file attributes. The amount of information in the catalog is determined by whether you choose to truncate the catalogs after a specific amount of time.

Catalogs reside on the Backup Exec server and on the storage device to which you sent the backup job.

To restore the data that was backed up by another installation of Backup Exec, you must first run a catalog operation on the storage device on the local Backup Exec server. The catalog for a backup job that was run on one installation of Backup Exec does not exist on another installation of Backup Exec.

When you enable Granular Recovery Technology (GRT) for Exchange, SharePoint, or virtual machine backups, a catalog operation runs after the backup operation by default. The delayed catalog operation runs once every 24 hours, even if you schedule more than one GRT-enabled backup to run in the same period. Because the catalog operation runs at a different time, the GRT-enabled backup job does not block another scheduled GRT-enabled backup job from starting on time. You can schedule a delayed catalog operation.

See [“Configuring default options for catalogs”](#) on page 249.

See [“Cataloging a storage device”](#) on page 452.

See [“About cataloging tape or disk cartridge media that contains encrypted backup sets”](#) on page 409.

See [“Configuring delayed catalog operations to improve backup performance”](#) on page 536.

Configuring default options for catalogs

You can configure the default options for catalogs to specify the defaults that are best suited for your environment.

See [“How Backup Exec catalogs work”](#) on page 248.

To configure default options for catalogs

- 1 Click the Backup Exec button, click **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, click **Catalog**.
- 3 Edit any of the following options:

Request all media in the sequence for catalog operations

Catalogs the media in tape drives and disk cartridges by starting with the lowest known tape number in the tape family. For example, if you don't have the first tape, the catalog job starts with the second tape. If you uncheck this option, the catalog job begins on the tape that you specify.

If you uncheck **Request all media in the sequence for catalog operations**, then you cannot select the option **Use storage-based catalogs**.

This option is enabled by default.

Use storage-based catalogs Lets Backup Exec read the catalog information from the storage device.

Storage-based catalogs allow quick cataloging of the backup sets that are not included in the Backup Exec server-based catalog. An example is when you want to catalog backup sets that another installation of Backup Exec creates.

Storage-based catalogs enable backup sets to be cataloged in minutes, rather than the hours that are required with traditional file-by-file cataloging methods.

To create a new catalog by having Backup Exec read each file block, clear this option. You should clear this option only if normal catalog methods are unsuccessful.

Note: Symantec recommends that you always attempt to use storage-based catalogs first before clearing this option.

If you uncheck **Request all media in the sequence for catalog operations**, then the option **Use storage-based catalogs** is unavailable.

This option is enabled by default.

The **Use storage-based catalogs** option must be enabled for the following restore scenarios:

- If you use the NDMP option, this option must be enabled so that NAS server backup sets can be cataloged.
See [“About restoring and redirecting restore data for NDMP servers”](#) on page 1179.
- If you use the Advanced Disk-based Backup Option, this option must be enabled to use the true image restore capability.
See [“About true image restore for synthetic backups”](#) on page 1161.
- If you use Simplified Disaster Recovery, this option must be enabled to use the backup sets as part of an SDR recovery operation.
See [“Preparing computers for use with Simplified Disaster Recovery”](#) on page 739.
- If you back up a virtual machine to tape, this option must be enabled to perform a Granular Recovery Technology enabled restore with the backup sets.
See [“Granular Recovery Technology”](#) on page 590.

Truncate catalogs after

Retains only the header information and removes all file details and directory details after the specified amount of time. This option reduces the size of the catalogs considerably. After the catalogs have been truncated, the files and directories cannot be restored until you recatalog the storage.

See “[Cataloging a storage device](#)” on page 452.

The last access date is not reset when catalogs are truncated.

You can perform a full restore of backup sets from truncated catalogs.

This option does not apply to synthetic backup jobs.

This option is not enabled by default.

Catalog path

Designates a path on the volume for the catalog files. If the path does not exist, you are prompted to create the path.

The default path is C:\Program Files\Symantec\Backup Exec\Catalogs.

Note: Symantec recommends that you use Backup Exec Utility to change the catalog location. If you change the location using the **Catalog path** field, you must also manually copy the existing catalogs to the new location, and then restart the Backup Exec services. Refer to the following URL for more information:

<https://www.symantec.com/docs/TECH210578>

- 4 Click **OK**.

Cataloging backup sets

Before you can restore or verify data, the data must be cataloged. While backing up data, Backup Exec creates a catalog that contains information about the backup sets and about the storage device on which the backup sets are stored. However, you can manually catalog backup sets at any time.

To catalog backup sets

- 1 On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the backup sets that you want to catalog.
- 2 In the left pane, click **Backup Sets**.

- 3 Do one of the following:
 - To catalog a single backup set, right-click the backup set.
 - To catalog multiple backup sets, Shift + click or Ctrl + click the backup sets, and then right-click one of the selected backup sets.
- 4 Click **Catalog**.

Backup Exec runs the catalog operation.

See [“Backup sets ”](#) on page 342.

See [“How Backup Exec catalogs work”](#) on page 248.

Job management and monitoring

This chapter includes the following topics:

- [How to monitor and manage jobs in Backup Exec](#)
- [About the Job Monitor](#)
- [About the Job History](#)
- [Viewing the job log](#)
- [Error-handling rules for failed or canceled jobs](#)
- [List of job statuses in Backup Exec](#)
- [Setting job status and recovery options](#)

How to monitor and manage jobs in Backup Exec

Backup Exec provides many ways to monitor and manage your backup, restore, and storage operation jobs.

Job monitoring

You can monitor your active jobs and scheduled jobs from the **Job Monitor** tab or from the **Jobs** list on the **Backup and Restore** tab or the **Storage** tab. You can monitor the types of jobs that are active and scheduled, the state and status of the jobs, the schedule, and other details. The status of reports can be monitored from the **Reports** tab.

Table 6-1 Places where you can monitor jobs

Location where you can monitor jobs	More information
The Job Monitor tab	<p>The Job Monitor provides a single location to monitor and manage all of your backup, restore, installation, and storage operation jobs. The Job Monitor is divided into two panes. The Jobs pane displays the details of all of your active jobs and scheduled jobs. The Job Histories pane displays the details about all of your jobs that ran recently. By default, all recent jobs appear in the Job Histories pane. You can change the default by filtering the Job Histories view.</p> <p>Note: Report jobs do not appear on the Job Monitor. To monitor and manage report jobs, go to the Reports tab.</p>
The Backup and Restore tab	<p>When you double-click a server name on the Backup and Restore tab, details for that server become available. You can view the backup and restore jobs that are scheduled and active for the selected server. You can also see the history of backup and restore jobs that were processed on the selected server.</p>
The Storage tab	<p>When you double-click a storage device name on the Storage tab, details for that storage device become available. You can view the storage operation jobs that are scheduled and active for the selected storage device. You can also see the history of storage operation jobs that were processed for the selected storage device.</p>
The Reports tab	<p>On the Reports tab, you can see the scheduled reports in the See Upcoming report group, and you can see a list of reports that completed in the See Completed report group.</p>

Job management

Backup Exec includes many features that enable you to manage backup jobs, restore jobs, and storage operation jobs.

You can manage jobs in the following ways:

- Edit scheduled jobs.
- Delete scheduled jobs.
- Cancel active jobs.
- Change the priority of scheduled jobs.
- Run a scheduled job immediately.
- Hold a job or the job queue.
- Run a test of the backup job.
- View the job activity details, such as job statistics and system information, for an active job.
- View the history of a job.

See [“Editing backup definitions”](#) on page 206.

See [“Viewing job activity details for active jobs”](#) on page 257.

See [“Deleting scheduled jobs”](#) on page 263.

See [“Canceling an active job”](#) on page 257.

See [“Changing the priority for a scheduled job”](#) on page 262.

See [“Running a scheduled job immediately”](#) on page 261.

See [“Holding jobs”](#) on page 258.

See [“Holding the job queue”](#) on page 260.

See [“Running a test run job manually”](#) on page 226.

About the Job Monitor

The **Job Monitor** provides a single location to monitor and manage all of your backup, restore, installation, and storage operation jobs. The **Job Monitor** is divided into two panes. The **Jobs** pane displays the details of all of your active jobs and scheduled jobs. The **Job History** pane displays the details about all of your jobs that ran recently. By default, all recent jobs appear in the **Job History** pane. You can change the default by filtering the **Job History** view.

Note: Report jobs do not appear on the **Job Monitor**. To monitor and manage report jobs, go to the **Reports** tab.

You can manage active and scheduled jobs in the **Jobs** pane on the **Job Monitor** in the following ways:

- Edit scheduled jobs.
- Delete scheduled jobs.
- Cancel active jobs.
- Change the priority of scheduled jobs.
- Run a scheduled job immediately.
- Hold a job or the job queue.
- Run a test backup job.
- View the job activity details for an active job.
- View the history of a job.

In the **Job History** pane of the **Job Monitor**, you can manage completed jobs in the following ways:

- Run the job again immediately.
- View the job log.
- Delete the job.
- View the job history details.
- Enable error-handling rules.
- Duplicate job histories.
- Run a verify backup job.
- Edit the settings for a backup job.

From the **Job Monitor**, you can restore data from the backup sets of a completed backup job. Additionally, you can view the backup calendar.

See [“How to monitor and manage jobs in Backup Exec”](#) on page 253.

See [“Running a test run job manually”](#) on page 226.

See [“Duplicating backup sets or a job history manually”](#) on page 221.

See [“Verifying backed up data manually”](#) on page 227.

See [“Editing backup definitions”](#) on page 206.

Viewing job activity details for active jobs

When a job is running, you can view details for the job, such as the percent complete, job rate, and byte count.

To view details for active jobs

- 1 Do one of the following:

To view job activity from the **Job Monitor** tab Select the **Job Monitor** tab.

To view job activity from the **Backup and Restore** tab or the **Storage** tab Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device on which the job is running.
- In the left pane, click **Jobs**.

- 2 Right-click the job, and then click **View Job Activity**.

Canceling an active job

You can cancel a job that is in progress. If the job is scheduled, it runs again at the next scheduled time.

It may take several minutes for a job to cancel. While Backup Exec processes the cancelation of a job, the Cancel Pending status appears in the **Job Status** column.

To cancel an active job

- 1 Do one of the following:

To cancel the job from the **Job Monitor** tab Select the **Job Monitor** tab.

To cancel the job from the **Backup and Restore** tab or the **Storage** tab Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job is running.
- In the left pane, click **Jobs**.

- 2 Right-click the active job that you want to cancel, and then click **Cancel**.
- 3 Click **Yes** to confirm the cancelation of the job.

See [“List of job statuses in Backup Exec”](#) on page 277.

Holding jobs

Active and scheduled jobs can be placed on hold. When you place an active job on hold, the job continues to run until it is complete. However, the next scheduled occurrence of that job is placed on hold. When you place a job on hold, the status in the **Job Status** column changes to On Hold.

To hold jobs

- 1 Do one of the following:

To place a job on hold from the Job Monitor tab	Select the Job Monitor tab.
To place a job on hold from the Backup and Restore tab or the Storage tab	<div>Do the following:</div> <div><div>■ On the Backup and Restore tab or the Storage tab, double-click the server or the storage device where the job is running or is scheduled to run.</div><div>■ In the left pane, click Jobs.</div></div>

- 2 Do one of the following:

To hold a single job	Select the job from the list of jobs, and then in the Jobs group, click Hold . Then select Hold Job .
To hold all active jobs	<div>In the Jobs group, click Hold, and then click Hold All Active Jobs. You may need to click Yes to confirm that you want to hold all of the active jobs.</div> <div>Note: This option is available only from the Job Monitor.</div>
To hold all scheduled jobs	<div>In the Jobs group, click Hold, and then click Hold All Scheduled Jobs. You may need to click Yes to confirm that you want to hold all of the scheduled jobs.</div> <div>Note: This option is available only from the Job Monitor.</div>

See [“Removing the hold on jobs”](#) on page 259.

Removing the hold on jobs

You can remove the hold on a job at any time. When you remove the hold on a job, the status in the **Job Status** column changes to the job's original status, such as Active or Scheduled.

To remove the hold on jobs

1 Do one of the following:

To remove the hold on jobs from the **Job Monitor** tab Select the **Job Monitor** tab.

To remove the hold from jobs from the **Backup and Restore** tab or the **Storage** tab Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job is on hold.
- In the left pane, click **Jobs**.

2 Do one of the following:

To remove the hold on a single job

Select the job from the list of jobs, and then in the **Jobs** group, click **Hold**. Then select **Hold Job** to remove the check mark.

Note: If the job that you want to hold is part of a backup definition, you may need to double-click the job definition to view the job.

To remove the hold on all active jobs

In the **Jobs** group, click **Hold**, and then click **Hold All Active Jobs** to remove the check mark. You may need to click **Yes** to confirm that you want to remove the hold on all of the active jobs.

Note: This option is available only from the **Job Monitor**.

To remove the hold on all scheduled jobs

In the **Jobs** group, click **Hold**, and then click **Hold All Scheduled Jobs** to remove the check mark. You may need to click **Yes** to confirm that you want to remove the hold on all of the scheduled jobs.

Note: This option is available only from the **Job Monitor**.

See [“Holding jobs”](#) on page 258.

Holding the job queue

You can place the entire job queue on hold to make changes to your environment. The server is paused to place the job queue on hold. When the job queue is on hold, only active jobs continue to run unless you choose to cancel them. No other jobs can run until the job queue is taken off hold.

To place the job queue on hold

- 1 Do one of the following

To place the job queue on hold from the **Job Monitor** tab Select the **Job Monitor** tab.

To place the job queue on hold from the **Backup and Restore** tab or the **Storage** tab Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job is running or scheduled to run.
- In the left pane, click **Jobs**.

- 2 In the **Jobs** group, click **Hold**, and then click **Hold Job Queue**.
- 3 Click **Yes**.
- 4 If active jobs are running, select the active jobs that you want to cancel, and then click **OK**.

See [“Removing the hold on the job queue”](#) on page 260.

Removing the hold on the job queue

When you remove the hold on the job queue, the server is unpaused and jobs then run according to the schedule.

To remove the hold on the job queue

- 1 Do one of the following

To remove the hold on the job queue from the **Job Monitor** tab Select the **Job Monitor** tab.

To remove the hold on the job queue from the **Backup and Restore** tab or the **Storage** tab Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job queue is on hold.
- In the left pane, click **Jobs**.

- 2 In the **Jobs** group, click **Hold**, and then click **Hold Job Queue** to remove the check box.

See [“Holding the job queue”](#) on page 260.

Running a scheduled job immediately

You can run a scheduled job immediately. The job will also run on the next scheduled occurrence.

To run a scheduled job immediately

- 1 Do one of the following:

To run the job from the **Job Monitor** tab Select the **Job Monitor** tab.

To run the job from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job is scheduled to run.
- In the left pane, click **Jobs**.

- 2 Right-click the scheduled job that you want to run, and then click **Run Now**.

Note: If the scheduled job is part of a backup definition, you may need to double-click the job definition to view the scheduled job.

See [“Changing the priority for a scheduled job”](#) on page 262.

See [“Deleting scheduled jobs”](#) on page 263.

Changing the priority for a scheduled job

The priority determines the order that jobs run. If two jobs are scheduled to run at the same time, the priority you set determines which job runs first. The priority is changed for all occurrences of the scheduled job.

The priority of the job is displayed in the **Priority** column in the **Jobs** list.

This option is most useful if there are limited storage devices in your environment, but you want certain jobs to have priority access to the devices. A ready job that has a higher priority runs before a ready job that has a lower priority. A ready job that has a higher priority also runs before a ready job that has an earlier scheduled start time.

If multiple jobs are ready to run but must wait for a storage device to become available, then Backup Exec determines which jobs to run first. Backup Exec reviews the job priority and the scheduled start time of the job.

You can set the job priority when you do the following tasks:

- Select storage options when you create or edit a backup job
- View scheduled jobs in the **Job Monitor** tab
- View a server's scheduled jobs from the **Jobs** pane on the **Backup and Restore** tab

If you change the job priority in the storage options or on the **Jobs** pane, you can choose from the following levels of priority:

- **Highest**
- **High**
- **Medium**
- **Low**
- **Lowest**

If you change the job priority from the **Job Monitor**, you can only increase or decrease the job priority. You cannot specify a level.

To change the priority for a scheduled job

- 1 Do one of the following:

To change the job's priority from the **Job Monitor** tab: Select the **Job Monitor** tab.

To change the job's priority from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job is scheduled to run.
- In the left pane, click **Jobs**.

- 2 Right-click the scheduled job, and then click **Change Priority**.

Note: If the job is part of a backup definition, you may need to double-click the job definition to view the job.

- 3 Select the new priority.

Deleting scheduled jobs

Deleting a scheduled job removes all scheduled occurrences of the job. To delete only the occurrence of a scheduled job on a specific date, you can edit the schedule to remove that date.

Note: If a backup definition includes more than one type of job, then you cannot use the **Delete** option to delete an individual job from the definition. Instead, you must edit the job definition to remove the scheduled job.

To delete a scheduled job

- 1 Do one of the following:

To delete a scheduled job from the **Job Monitor** tab

Select the **Job Monitor** tab.

To delete a scheduled job from the **Backup and Restore** tab or the **Storage** tab Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job is scheduled to run.
- In the left pane, click **Jobs**.

- 2 Right-click the scheduled job, and then click **Delete**.

Note: If the job is part of a backup definition, you may need to double-click the job definition to view the job.

- 3 Click **Yes**.

See [“How to monitor and manage jobs in Backup Exec”](#) on page 253.

About the Job History

The **Job History** displays a list of completed and failed backup, restore, and storage operation jobs. The **Job History** appears on the lower pane of the **Job Monitor** tab. It also appears when you select a server on the **Backup and Restore** tab, and when you select a storage device on the **Storage** tab.

From the **Job History**, you can do any of the following:

- View the job log.
- Delete a job.
- Rerun a job.
- Duplicate the data from a completed backup job.
- Verify a backup job.
- Enable error-handling rules for a failed job.

See [“Running a job from the Job History”](#) on page 266.

See [“Viewing the job log ”](#) on page 267.

See [“Deleting a job from the Job History”](#) on page 265.

See [“Enabling an error-handling rule for a failed job”](#) on page 275.

See [“Duplicating backup sets or a job history manually”](#) on page 221.

See [“Verifying backed up data manually”](#) on page 227.

Viewing the history of a job

The job history shows statistics for all occurrences of a job.

To view the history of a job

- 1 Do one of the following:

To view the history of a job from the **Job Monitor** tab

On the **Job Monitor** tab, locate the job in the **Job Histories** pane.

To view the history of a job from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

- 2 Right-click the job, and then click **View Job History**.

Deleting a job from the Job History

You can delete a job from the **Job History**, or have Backup Exec automatically delete the job history using database maintenance.

If you delete a job, it is removed from the computer and cannot be recovered.

To delete a job from the Job History

- 1 Do one of the following:

To delete a job from the Job History on the **Job Monitor** tab

On the **Job Monitor** tab, locate the job in the **Job Histories** pane.

To delete a job from the Job History on the **Backup and Restore** tab or the **Storage** tab Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

- 2 Right-click the job that you want to delete, and then click **Delete**.

You can select multiple jobs by selecting a job, and then pressing the <Ctrl> or <Shift> keys while you click other jobs that you want to select. This lets you perform tasks such as Delete on more than one job at a time, as long as the jobs are of similar type.

You can delete up to 2500 jobs from the **Job History**. If you attempt to delete more than 2500 jobs, you are prompted to continue with the deletion.

- 3 Click **Yes**.

See [“About the Job History”](#) on page 264.

See [“Configuring database maintenance”](#) on page 567.

Running a job from the Job History

After a job runs, the job moves to the **Job History**. You can run a completed job again from the **Job History**.

To run a job from the Job History

- 1 Do one of the following:

To run a job from the **Job History** on the **Job Monitor** tab On the **Job Monitor** tab, locate the job in the **Job Histories** pane.

- To run a job from the **Job History** on the **Backup and Restore** tab or the **Storage** tab
- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
 - In the left pane, click **Job History**.

- 2 Right-click the job that you want to run, and then click **Run Now**.

See [“About the Job History”](#) on page 264.

Viewing the job log

You can view detailed job-related properties for each job that has been processed. You can save a copy of the job log to a location of your choice or you can print the job log.

To view the job log

- 1 Do one of the following:

To view the job log from the Job Monitor tab	On the Job Monitor tab, locate the job in the Job Histories pane.
---	---

To view the job log from the Backup and Restore tab or the Storage tab	Do the following:
--	-------------------

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

- 2 Right-click the job, and then select **View Job Log**.
- 3 Do any of the following:
 - Click **Find** to search for a particular word, phrase, or file name.
 - Click **Save As** to save a copy of the job log to the location of your choice.
 - Click **Print** to print the job log.

Finding text in the job log

You can search for specific text in the job log. Backup Exec searches only the sections that are expanded. To search the entire job log, you should select the **Expand All** option.

To find text in the job log

- 1 Do one of the following:

To find text in the job log from the Job Monitor tab	On the Job Monitor tab, in the Job Histories pane, locate the job for which you want to search the job log.
---	---

To find text in the job log from the **Backup and Restore** tab or the **Storage** tab Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

- 2 Right-click the job, and then click **View Job Log**.
- 3 Click **Find**.
- 4 In the **Find** field, type the text that you want to find.
- 5 Select any of the following additional options to facilitate your search:

Match whole word only

Check this check box if you want Backup Exec to search for the whole word you typed. If you do not select this option, Backup Exec finds the text that includes part of the word. For example, if you search for the word "file" and do not select this option, Backup Exec finds all occurrences of "file", "files", "filed", and any other words that contain "file". If you do select this option, Backup Exec finds only the occurrences of "file".

Match case

Check this check box if you want Backup Exec to search for words using the exact capitalization that you typed. For example, if you search for the word "File" and select this option, Backup Exec finds all occurrences of "File", but does not find any occurrences of "file".

Highlight all matches

Check this check box if you want Backup Exec to highlight the text that matches the search criteria. The option is selected by default.

- 6 Click **Next** to find the next occurrence of the text.

See [“Viewing the job log”](#) on page 267.

Printing the job log

If your Backup Exec server is connected to a printer, you can print the job log for a completed job.

To print the job log

- 1 Do one of the following:

To access the job log from the **Job Monitor** tab

On the **Job Monitor** tab, locate the job in the **Job Histories** pane.

To access the job log from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

- 2 Right-click the job, and then select **View Job Log**.
- 3 Click **Print**.
- 4 Select the printer that you want to use, and then click **Print**.

See [“Saving the job log”](#) on page 269.

Saving the job log

Backup Exec provides the ability to save the job log to a location on your hard drive or network. In addition, you can select the format in which to save the file; as a complete webpage, a web archive, an HTML-only webpage, or a text file.

To save the job log

- 1 Do one of the following:

To access the job log from the **Job Monitor** tab

On the **Job Monitor** tab, locate the job in the **Job Histories** pane.

To access the job log from the **Backup and Restore** tab or the **Storage** tab

Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

- 2 Right-click the job, and then select **View Job Log**.

- 3 Click **Save As**.
 - 4 Select the location where you want to save the job log.
- See [“Printing the job log”](#) on page 269.

Linking from the job log to the Symantec Technical Support website

Errors that are reported in the job log have a unique code, called a Unique Message Identifier (UMI). These codes contain hyperlinks that you can click to go to the Symantec Technical Support website. From the website, you can access technical notes and troubleshooting tips that are related to a specific message. UMI codes establish unique message codes across all Symantec products.

Some alerts also contain a UMI. For example, if a Warning alert appears when a job fails, the alert includes the UMI code.

You can create or enable an error-handling rule for errors. These rules let you set options to retry or stop a job when the error occurs.

See [“Error-handling rules for failed or canceled jobs”](#) on page 272.

To link from the job log to the Symantec Technical Support website

- 1 Do one of the following:

To link to the job log from the **Job Monitor** tab Select the **Job Monitor** tab.

To link to the job log from the **Backup and Restore** tab or the **Storage** tab Do the following:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

- 2 Right-click a job, and then select **View Job Log**.
- 3 Scroll to the **Job Completion Status** section.
- 4 Click the UMI code, which appears as a blue hyperlink.

How to use job logs with vertical applications

The Backup Exec Administration Console provides a view of the job logs in HTML format. If necessary, you can convert the job logs to a text format for use with vertical applications.

To convert a job log file to a text format, load the Backup Exec Management Command Line Interface, and then type the following at a command prompt:

Get-BEJobLog "pathname\job log filename"

For example, to display the job log C:\program files\Symantec\Backup Exec\Data\bex00001.xml in text format to the command prompt, you would type:

Get-BEJobLog "C:\program files\Symantec\Backup Exec\Data\bex00001.xml"

To redirect the job log to a file, you would type one of the following:

Get-BEJobLog "C:\program files\Symantec\Backup Exec\Data\bex00001.xml" > bex00001.txt

See ["Viewing the job log"](#) on page 267.

Configuring default job log options

You can configure default options for job logs that specify the amount of detail you want to include in the completed job log. For the jobs that produce large job logs, you may want to reduce the amount of detail in the job log. The size of the job log increases proportionally to the level of detail that is configured for the job log.

To configure default job log options

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, click **Job Logs**.
- 3 Select the appropriate options.

Summary information only

Select this option to include the following information in the job log:

- Job name
- Job type
- Job log name
- Backup Exec server name
- Storage device
- Starting date and time
- Errors encountered
- Ending date and time
- Completion statistics

This option also includes the names of files that were skipped, the name of the media set, the backup type and results of the verify operation if one was performed.

Summary information and directories processed	Select this option to include summary information and a list of all processed subdirectories in the job log.
Summary information, directories, and files processed	Select this option to include summary information, processed subdirectories, and a list of all the file names that were processed in the job log.
Summary information, directories, files, and file details	Select this option to include summary information, processed subdirectories, a list of all the file names and their attributes in the job log. This option increases the job log size significantly.
Prefix for the job log file name	Enter a prefix to add to the job logs that are processed. The default prefix is BEX. The job log file name consists of Prefix_ServerName_Count. Prefix is the label that you enter in this field, ServerName is the name of the Backup Exec server that ran the job, and Count is the number of job logs that this job has produced
Attach job logs as HTML	Select this option to attach the job logs in an HTML format when an email notification is sent.
Attach job logs as text	Select this option to attach the job logs in a text format when an email notification is sent.
Job log path	Shows the current location of the job log. To change the path you can use BE Utility.

See [“Viewing the job log ”](#) on page 267.

Error-handling rules for failed or canceled jobs

You can enable default rules or create custom rules to set retry options and final job disposition for failed or canceled jobs. Retry options let you specify how often to retry a job if it fails and the time to wait between retry attempts. The final job disposition lets you either place the job on hold until you can fix the error, or reschedule the job for its next scheduled service.

Each default error-handling rule applies to one category of errors, such as Network Errors or Security Errors. Default error-handling rules are disabled by default, so you must edit a rule and enable the rules that you want to use. You cannot delete default error-handling rules, add specific error codes to a category, or add new error categories. Before the error-handling rules will apply, the final error code must be in an error category that is associated with a rule, and the rule must be enabled.

To apply an error-handling rule for a specific error code that is in an error category, you can create a custom error-handling rule. You can select up to 28 error codes in an error category that a custom error-handling rule can apply to. You can also add an error code to an existing custom rule.

A custom error-handling rule named "Recovered Jobs" is created when Backup Exec is installed and is enabled by default. This rule applies retry options and a final job disposition to jobs that fail and that are not scheduled to run again.

See ["Creating a custom error-handling rule"](#) on page 273.

If both a custom error-handling rule and a default error-handling rule apply to a failed job, the settings in the custom rule are applied to the job.

Note: If the server on which Backup Exec is installed is in a cluster environment, the Cluster Failover error-handling rule is displayed on the list of error-handling rules. This rule is enabled by default.

See ["About the cluster failover error-handling rule"](#) on page 276.

Creating a custom error-handling rule

You can create custom rules to set retry options and final job disposition for failed or canceled jobs.

See ["Error-handling rules for failed or canceled jobs"](#) on page 272.

To create a custom error-handling rule

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Error-handling rules**.
- 2 Click **New**.
- 3 Check **Enable error-handling rule**.
- 4 In the **Name** field, type a name for this rule.
- 5 In the **Error category** list, select the category of error that you want this rule to apply to.

- 6 Under **Available errors**, check the check box of the error code that you want this rule to apply to. You can select up to 28 error codes.
- 7 If you want Backup Exec to retry the job after it fails, check **Retry job**, and then enter the number of times you want to retry the job and how often you want to retry the job.

The maximum number of retries is 99. The maximum retry interval is 1440 minutes.

- 8 Under **Final job disposition**, select the way you want to handle the job after the maximum number of retries.

If you choose the option to place the job on hold until the error condition is cleared manually, you must manually remove the hold from the job after you manually clear the error condition.

- 9 Click **OK**.

See [“Custom error-handling rule for recovered jobs”](#) on page 275.

See [“About the cluster failover error-handling rule”](#) on page 276.

Enabling or disabling error-handling rules

Follow these steps to enable or disable specific error-handling rules.

To enable or disable error-handling rules

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Error-handling rules**.
- 2 Select the rule that you want to enable or disable, and then click **Edit**.
- 3 Do one of the following:
 - To enable the rule, check **Enable error-handling rule**.
 - To disable the rule, clear the **Enable error-handling rule** check box
- 4 Click **OK**.

See [“Error-handling rules for failed or canceled jobs”](#) on page 272.

Deleting a custom error-handling rule

A custom error-handling rule can be deleted at any time. A default error-handling rule cannot be deleted.

To delete a custom error-handling rule

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Error-handling rules**.
- 2 Select the custom rule that you want to delete, and then click **Delete**.
- 3 Click **Yes** to confirm that you want to delete the rule.

See [“Error-handling rules for failed or canceled jobs”](#) on page 272.

Enabling an error-handling rule for a failed job

You can create custom rules to set retry options and final job disposition for failed jobs.

To enable an error-handling rule for a failed job

- 1 Do one of the following:

To enable an error-handling rule from the **Job Monitor** tab Select the **Job Monitor** tab.

To enable an error-handling rule from the **Backup and Restore** tab or the **Storage** tab On the **Backup and Restore** tab or the **Storage** tab, in the left pane, select **Job History**.

- 2 Right-click the failed job, and then select **Error Handling**.
- 3 Check the **Enable error-handling rule** check box.
- 4 Complete the remaining options for this rule.

Custom error-handling rule for recovered jobs

Backup Exec includes a custom error-handling rule called "Recovered Jobs" to recover the jobs that failed with specific errors. This rule is created when Backup Exec is installed and is enabled by default.

The retry options for this rule are to retry the job twice, with an interval of five minutes between the retry attempts. The final job disposition is to place the job on hold until you have manually cleared the error condition.

The following table describes the error codes that are selected by default for the Recovered Jobs custom error-handling rule.

Table 6-2 Error codes for recovered jobs custom error-handling rule

Error code	Description
0xE00081D9 E_JOB_ENGINE_DEAD	The displayed error message is: The Backup Exec job engine system service is not responding.
0xE0008820 E_JOB_LOCAL RECOVERNORMAL	The displayed error message is: The local job has been recovered. No user action is required.
0xE000881F E_JOB_REMOTE RECOVERNORMAL	The displayed error message is: The remote job has been recovered. No user action is required.
0xE0008821 E_JOB_STARTUP RECOVERY	The displayed error message is: Job was recovered as a result of Backup Exec RPC service starting. No user action is required.

Note: If the Central Admin Server Option is installed, additional error codes are selected.

See [“Error-handling rules for failed or canceled jobs”](#) on page 272.

See [“About the cluster failover error-handling rule”](#) on page 276.

About the cluster failover error-handling rule

If the Backup Exec server is installed in a cluster environment, the cluster failover error-handling rule is displayed on the list of error-handling rules. This rule is enabled by default.

You cannot configure any options for this rule. You can only enable or disable the cluster failover error-handling rule.

The cluster failover error-handling rule and the **Enable checkpoint restart** option in **Advanced Open File** backup options work together to enable you to resume jobs from the point of failover. The **Enable checkpoint restart** option is dependent on the cluster failover error-handling rule; if you disable the rule, the option is automatically disabled to match the rule's setting.

See [“Error-handling rules for failed or canceled jobs”](#) on page 272.

List of job statuses in Backup Exec

Backup Exec displays various job statuses for active, completed, and scheduled jobs.

See [the section called “Active job statuses”](#) on page 277.

See [the section called “Completed job statuses”](#) on page 278.

See [the section called “Scheduled job statuses”](#) on page 280.

Active job statuses

The following statuses may appear for the jobs that are active:

Table 6-3 Active job statuses

Status	Description
Running	The operation is underway.
Queued	The job has been initiated, but Backup Exec is actively looking for a suitable drive or media.
Cancel Pending	Backup Exec cannot process the Cancel request immediately. This status is displayed until the job is actually canceled. The job is then displayed in the job history with a status of Canceled.
Loading Media	The media is being loaded and positioned on the target device.
Pre-processing	<p>This status can indicate any or all of the following:</p> <ul style="list-style-type: none">■ Backup Exec is calculating the amount of data that will be backed up, if the Display progress indicators for backup jobs option is enabled in the Preferences section of Backup Exec settings. See “Changing the default preferences” on page 565.■ Backup Exec is waiting for a pre-command or a post-command to complete.■ Backup Exec is retrieving the set maps and is positioning the tape to the append point location for an append job.
Snapshot processing	Backup Exec is processing a snapshot operation.
Device Paused	<p>The device that the job was sent to is paused.</p> <p>See “Pausing and unpausing a storage device” on page 455.</p>

Table 6-3 Active job statuses (*continued*)

Status	Description
Server Paused	The Backup Exec server is paused. See “Pausing or resuming a managed Backup Exec server” on page 1144.
Stalled	The Backup Exec services have become unresponsive. See “Setting job status and recovery options” on page 282.
Media Request	You must insert media for the job to continue.
Communication Stalled	Communications between the managed Backup Exec server and the central administration server have not occurred within the configured time threshold. See “Enabling or disabling communications between the managed Backup Exec server and the central administration server” on page 1132.
No Communication	No communication about jobs is being received at the central administration server from the managed Backup Exec server. The configured time threshold has been reached. See “Enabling or disabling communications between the managed Backup Exec server and the central administration server” on page 1132.
Consistency check	Backup Exec is running a consistency check of the databases before backup.
Updating Catalogs	Backup Exec is updating the catalog information.

Completed job statuses

The following statuses may appear for the jobs that were completed:

Table 6-4 Job completion status

Status	Description
Successful	The job completed without errors.
Completed with exceptions	The job completed, but some files were in use, skipped, or corrupted.

Table 6-4 Job completion status (*continued*)

Status	Description
Failed over	The job ran in a cluster environment and was active on one computer, and then the cluster performed a failover and the job was restarted on another computer in the cluster. Two separate sets of job history are available when a job is failed over. The first job history includes the failed over status and the second job history includes the status that is appropriate for the completed job.
Resumed	The status is the same as the failed over status, however the Enable checkpoint restart option was selected.
Canceled	The administrator terminated the operation as it was running.
Canceled, timed out	The Cancel the job if it is still running x hours after its scheduled start feature was enabled and the job was not completed within the specified timeframe.
Failed	<p>The operation took place, but one or more significant errors occurred. The job log should indicate what caused the errors so that you can decide if you want to run the job again. For example, if a job failure occurred due to a lost connection during job processing, you could choose to resubmit the job when the connection is restored.</p> <p>If a drive loses power during a backup operation, you should restart the backup job using a different tape. You can restore the data that was written to the tape up to the point of the power loss, but you should not reuse the tape for subsequent backup operations.</p> <p>A failed job has an error message in the Errors section of the job log with a link to the Symantec Technical Support website.</p> <p>A job may fail for the following reasons:</p> <ul style="list-style-type: none"> ■ The storage device that was selected for the job was not available when the job ran. ■ The logon account that was used in the backup job is incorrect. Verify that the logon account information is valid for the resource being backed up. ■ A problem occurred with the storage device when the job ran. ■ The computer being backed up was shut down before or during the backup job.
Recovered	The job was active when the status of the managed Backup Exec server was changed from Communication Stalled to No Communication. The custom error-handling rule for Recovered Jobs was applied to the job.

Table 6-4 Job completion status (*continued*)

Status	Description
Missed	The job did not run during the scheduled time window. The job is rescheduled to run based on the time window that you configured.

Scheduled job statuses

The following statuses may appear for the jobs that are scheduled:

Table 6-5 Scheduled job statuses

Status	Description
Invalid Schedule	The scheduled job will not run because of a scheduling issue. See “Setting default schedule options for rule-based jobs and run now jobs” on page 559.
Not in time window	The job was ready to be sent for processing, but the time window for the job closed. See “Setting default schedule options for rule-based jobs and run now jobs” on page 559.
On Hold	The job has been placed on hold.
Queued	A temporary state that displays when Backup Exec is applying an error-handling rule that is enabled to retry the job. See “Custom error-handling rule for recovered jobs” on page 275.

Table 6-5 Scheduled job statuses (*continued*)

Status	Description
Ready	<p>The job is ready to run, but cannot for one of the following reasons:</p> <ul style="list-style-type: none"> ■ Internal error. No devices are available, but the cause is unknown. ■ Invalid job. The job type is unknown; there may be an internal error or the database is corrupted. ■ Invalid target. This type of storage device no longer exists. ■ Backup Exec server not available. ■ No license for option name. A license must be purchased and installed on the Backup Exec server. ■ No Backup Exec servers are available. ■ No Backup Exec servers are available in Backup Exec server pool. ■ Specified destination storage device pool is empty. ■ Specified destination device is not in Backup Exec server pool. ■ Specified destination device not on local Backup Exec server. ■ Specified destination storage device pool on local Backup Exec server is empty. ■ The destination storage device cannot be a storage pool. ■ The destination storage device cannot be a Backup Exec server. ■ Another job is running in the system that is blocking execution of this job. This job will run after the other job completes. ■ Invalid input. ■ Incompatible Resumes. ■ No server license available. ■ No multi-server license available. ■ No Windows license. ■ No Windows server. ■ Need local Backup Exec server. ■ Local server is not a Backup Exec server. ■ No idle storage devices are available. ■ No eligible storage devices within the storage pool are available. ■ Blocked by an active, linked duplicate backup sets job.

Table 6-5 Scheduled job statuses (*continued*)

Status	Description
Scheduled	The job is scheduled to run in the future. The scheduled jobs that are linked to another job, such as a job to duplicate backup sets, will not display a scheduled job status.
Server Paused	<p>The job is ready, but the Backup Exec server has been paused. No jobs are dispatched while the Backup Exec server is paused.</p> <p>See “Pausing or resuming a managed Backup Exec server” on page 1144.</p>
To Be Scheduled	A state that the scheduled job transitions through as it is being sent for processing.
Rule Blocked	<p>The job cannot run because one or more of the settings in the backup definition cannot be satisfied.</p> <p>The Rule Blocked status may appear for any of the following reasons:</p> <ul style="list-style-type: none">■ A job cannot run until another job has completed. Example: If you added a duplicate stage to a backup definition and the source backup job has not yet completed, the duplicate job is blocked until the source backup job completes. Example: If a job definition includes both a full backup and an incremental backup, the full backup must run first. If you try to run the associated incremental backup job before the full backup job has completed, the incremental backup job is blocked until the full backup completes.■ A server for a job cannot be changed until the linked jobs complete.■ A server for a job cannot be changed until another job completes.■ A job cannot run because multiple jobs are scheduled to run at the same time.

Setting job status and recovery options

If the Backup Exec services become unresponsive or jobs no longer run, you can set the threshold at which Backup Exec changes the status of active jobs to stalled. You can also set the threshold at which Backup Exec fails the jobs that were stalled, and then recovers them.

By lowering the amount of time before Backup Exec reaches the threshold for changing a job's status to stalled, you can receive an earlier notification that jobs have stalled. A shorter time between the stalled and recovered thresholds also allows Backup Exec to fail and then recover the stalled jobs earlier. However, setting the thresholds too low may force a job to be recovered when it is not necessary.

Backup Exec recovers the jobs by using the custom error-handling rule named Recovered Jobs. This custom error-handling rule is created and enabled when Backup Exec is installed, and specifies that stalled, failed, and recovered jobs are retried two times, with an interval of five minutes between the retries.

Jobs that are stalled and then failed and recovered by Backup Exec because of unresponsive Backup Exec services are displayed differently in Backup Exec than the jobs that fail because of errors in normal daily activities. The stalled/failed/recovered jobs are not indicated in red text in the job history as other failed jobs are. Instead, these jobs are displayed in gray text with a job status of **Recovered**.

In the job history, the error category is listed as Job Errors. The job history indicates the type of internal communication error that occurred and that the job was recovered. Based on the type of error that occurred, a log file may or may not be associated with the recovered job.

To set job status and recovery options

- 1 Click the Backup Exec button, click **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, click **Job Status and Recovery**.
- 3 Set the thresholds for stalled and recovered jobs.

Stalled

Enter the amount of time you want to wait before Backup Exec changes an unresponsive job's status to Stalled.

Recovered

Enter the amount of time you want to wait before Backup Exec fails jobs that stalled and then recovers them. A custom error-handling rule named Recovered Jobs is applied to recovered jobs. If this rule is disabled, then any other error-handling rules that have been enabled will apply to the recovered jobs. If no error-handling rules apply to the job, then the job fails.

- 4 Click **OK**.

See [“List of job statuses in Backup Exec”](#) on page 277.

See [“Custom error-handling rule for recovered jobs”](#) on page 275.

Alerts and notifications

This chapter includes the following topics:

- [Alerts and notifications in Backup Exec](#)
- [Where to find alerts in Backup Exec](#)
- [Enabling active alerts and alert history to display on the Home tab](#)
- [Viewing the alert history for a server or a storage device](#)
- [Deleting an alert from alert history](#)
- [Copying alert text to a document or email](#)
- [Filtering alerts](#)
- [Viewing the job log from an alert](#)
- [Responding to active alerts](#)
- [Clearing all informational alerts manually](#)
- [Setting up notification for alerts](#)
- [Configuring email or text message notification for alerts](#)
- [Adding recipients for notification](#)
- [Adding a recipient group for alert notifications](#)
- [Disabling email or text message alert notification for a recipient](#)
- [Editing recipient notification properties](#)
- [Deleting recipients](#)
- [Configuring alert categories](#)

- [Assigning recipients to receive notifications for specific alert categories](#)
- [Sending a notification when a job completes](#)
- [Disabling notifications for a specific alert category](#)
- [Configuring default alert settings](#)
- [Enabling or disabling pop-up alerts](#)
- [SNMP traps for Backup Exec alerts](#)

Alerts and notifications in Backup Exec

An alert is any event or condition in Backup Exec that is important enough to display a message or require a response from you. Backup Exec includes many alert categories and four alert types. Alert categories are the events or the conditions that cause alerts. Alert categories encompass many circumstances or problems that affect the system, jobs, media, or storage sources. Each alert category can include one or more events that generate an alert. For example, a Job Failed error may occur for many reasons. The alert severity levels can help you to determine which alerts need immediate attention and which alerts require a response.

The following alert severity levels are used in Backup Exec:

Table 7-1 Alert severity

Item	Description
Attention required	Indicates the issues that require a response before the job or operation can continue.
Error	Indicates the issues that affect job processing or the integrity of your backup. These alerts cannot be disabled and cannot be configured to be cleared automatically. You must respond to them manually.
Warning	Indicates the conditions that may or may not cause jobs to fail. You should monitor the conditions and take actions to resolve them.
Informational	Provides status messages for the conditions that you might want to know about.

By default, most alerts are enabled, which means that they appear in the **Active Alerts** pane when they occur. You can disable warning alerts and informational alerts by editing alert category properties. However, error alerts and attention

required alerts cannot be disabled. You can filter the alerts so that only specific alerts appear.

See [“Configuring alert categories”](#) on page 302.

From the **Home** tab, you can view all active alerts or filter the alerts to view only specific alert severities or only the alerts that occurred on certain dates. On the **Backup and Restore** tab, when you double-click a server, you can see the active alerts that are specific to that server. Similarly, on the **Storage** tab, when you double-click a type of storage, you can see the active alerts that are specific to that storage device.

Alerts remain in the **Active Alerts** pane until they receive a response. You can respond to an alert manually or you can configure Backup Exec to respond to some alerts automatically after a specified length of time. Depending on the alert severity, a response might not be required, such as with informational alerts. After you respond to an alert, Backup Exec moves it to the alert history. Alert history is available on the **Home** tab, the **Backup and Restore** tab, and the **Storage** tab. In addition, an Alert History report is available from the **Reports** tab.

See [“Alert History report”](#) on page 652.

See [“Responding to active alerts”](#) on page 293.

You can configure notifications to inform recipients when alerts occur. For example, you can notify a backup administrator by email or cell phone text message when a critical alert occurs.

See [“Setting up notification for alerts”](#) on page 294.

To assist with hardware troubleshooting, Backup Exec displays alerts for SCSI event ID 9 (storage timeout), ID 11 (controller error), and ID 15 (storage not ready).

Where to find alerts in Backup Exec

You can find alerts in multiple locations on the Backup Exec Administration Console.

Table 7-2 Where alerts appear on the Backup Exec Administration Console

Location of alerts	Details
Home tab	<p>To see a list of active alerts on the Home tab, ensure that the Active Alerts check box is checked in the System Health group. To see a list of alerts in alert history, ensure that the Alert History check box is checked in the System Health group.</p> <p>From the Active Alerts widget on the Home page, you can respond to alerts, view the job log, clear all informational alerts, or view the details of specific alerts.</p>
Backup and Restore tab	<p>On the Backup and Restore tab, active alerts are listed for each server on the list of servers. Also, when you double-click a server, you can select Active Alerts in the left pane to display all of the active alerts for that server. When you select Active Alerts in the left pane, the Show Alert History option also becomes available in the Alerts group at the top of the dialog box.</p> <p>Additionally, you can click the alert icon next to a server name to view the alerts for that server.</p>
Storage tab	<p>On the Storage tab, active alerts are listed for each storage device on the list of storage devices. Also, when you double-click a storage device, you can select Active Alerts in the left pane to display all of the active alerts for that device. When you select Active Alerts in the left pane, the Show Alert History option also becomes available in the Alerts group at the top of the dialog box.</p> <p>Additionally, you can click the alert icon next to a storage device name to view the alerts for that device.</p>

Table 7-2

Where alerts appear on the Backup Exec Administration Console
(continued)

Location of alerts	Details
Backup Exec status bar	The Backup Exec status bar shows the number of active alerts for each type of alert. Double-click in the Alerts section of the Backup Exec status bar to see a list of all the active alerts and the alert history, along with details of those alerts.
Windows status bar	If you have minimized Backup Exec, the Backup Exec icon on the Windows status bar flashes when a Backup Exec alert is triggered. When you maximize Backup Exec, the pop-up alerts appear.
Pop-up window	Alert pop-up windows appear when an alert is triggered. The color of the alert pop-up corresponds to the type of alert; red for error alerts, yellow for warning alerts, purple for attention-required alerts, and blue for informational alerts. If more than three alerts are available, the pop-up lists the number of alerts that are pending and the color of the pop-up indicates the most severe of the alert types. For example, if three informational alerts and two error alerts have occurred, the pop-up alert indicates that five alerts have arrived. Additionally, the pop-up color will be red to indicate that error alerts are part of the group of new alerts. You can respond to attention required alert types directly from the pop-up.

Enabling active alerts and alert history to display on the Home tab

The **Active Alerts** pane appears on the **Home** tab by default. If it does not appear, follow these steps to show the alert details. Optionally, you can also enable a history of all alerts for the server to appear on the **Home** tab.

To view active alerts and alert history on the Home tab

- ◆ On the **Home** tab, in the **System Health** group, do any of the following:

- Check the **Active Alerts** check box to see a list of active alerts.
- Check the **Alert History** check box to see a list of all alerts that occurred on the server.

See [“Active alerts properties ”](#) on page 290.

Active alerts properties

Properties for active alerts can be viewed on the **Home** tab or in the details for a backup job, a restore job, or a type of storage.

See [“Enabling active alerts and alert history to display on the Home tab”](#) on page 289.

Table 7-3 Properties for active alerts

Item	Description
Severity	<p>Indicates the severity of the alert. The severity helps you determine how quickly you want to respond.</p> <p>The following alert severity levels may appear:</p> <ul style="list-style-type: none"> ■ Error ■ Warning ■ Information ■ Attention Required
Category	Indicates the condition that caused the alert. Categories include Database Maintenance, General Information, Device Error, or Job Failed.
Message	Indicates the text of the error message.
Date and Time	Shows the date and time when the alert was received.
Job Name	Indicates the name of the job that triggered the alert. This column is blank if a job triggered the alert, such as for general information alerts.
Storage	Shows the name of the storage device on which the alert occurred.
Server	Shows the name of the server on which the alert occurred.

Table 7-3 Properties for active alerts (*continued*)

Item	Description
Source	<p>Indicates the cause of the alert.</p> <p>Alerts can originate from one of the following sources:</p> <ul style="list-style-type: none">■ System■ Job■ Storage■ Media

Viewing the alert history for a server or a storage device

After you respond to an alert, Backup Exec moves it to the alert history.

To view the alert history for a server or a storage device

- 1 On the **Backup and Restore** tab or the **Storage** tab, double-click the item for which you want to view the alert history.
- 2 In the left pane, click **Active Alerts**.
- 3 In the **Alerts** group, select **Show Alert History**.

Deleting an alert from alert history

Follow these steps to delete an alert from the alert history.

To delete an alert from alert history

- 1 Access the alert history from any of the following locations:
 - In the Backup Exec status bar, double-click in the **Alerts** area.
 - On the **Backup and Restore** tab, double-click a server. In the left pane, select **Active Alerts**, and then click **Show Alert History**.
 - On the **Storage** tab, double-click a storage device. In the left pane, select **Active Alerts**, and then click **Show Alert History**.
- 2 On the **Alerts** dialog box, select the **Alert History** tab.
- 3 Right-click the alert that you want to remove from the alert history, and then click **Delete**.
- 4 Click **Yes** to confirm that you want to delete the alert.

Copying alert text to a document or email

Backup Exec gives you the ability to copy alert information to a document, such as Notepad or Microsoft Word, or to an email. When you copy alert text to a word-processing application or to an email, Backup Exec formats the alert text in a table. Alert text that is copied to Notepad displays in plain text.

To copy the text of an alert to a document or email

- 1 Access the list of alerts from any of the following locations:
 - In the Backup Exec status bar, double-click in the **Alerts** area.
 - On the **Backup and Restore** tab, double-click a server. In the left pane, select **Active Alerts**.
 - On the **Storage** tab, double-click a storage device. In the left pane, select **Active Alerts**.
- 2 On the **Alerts** dialog box, select either the **Active Alerts** tab or the **Alert History** tab, depending on whether the alert you want to copy is active or in the alert history.
- 3 Right-click the alert that you want to copy, and then click **Copy**.

Note: You can also select an alert, and then press Ctrl + C as a shortcut to copy the alert text. In addition, you can copy multiple alerts by using Shift + Click or Ctrl + Click.

- 4 Open the document or email to which you want to copy the alert text, right-click in the document or email, and then select **Paste**.

Filtering alerts

You can filter the alerts that appear in the **Active Alerts** pane on the **Home** tab. Filters are useful when you have many alerts and you want to only view specific alert types. Alerts can be filtered by severity, time, and source. For example, you can choose to view only the error alerts that occurred during the last 12 hours for jobs.

To filter alerts

- 1 On the **Home** tab, locate the **Active Alerts** pane.
If the **Active Alerts** pane does not appear, you must enable the alert details.
See [“Enabling active alerts and alert history to display on the Home tab”](#) on page 289.
- 2 Use any combination of the following options to filter the alerts list:
 - In the **Source** field, select the source of the alerts that you want to view.
 - In the **Time** field, select the time frame for which you want to view alerts.
 - In the **Severity** field, select the severity levels of the alerts that you want to view, such as **Error** or **Warning**.

Viewing the job log from an alert

The job log provides detailed job information, storage and media information, job options, file statistics, and job completion status for completed jobs. You can access the job log from the alerts that were generated for jobs.

To view the job log from an alert

- 1 Access the **Active Alerts** pane on the **Home** tab, the **Backup and Restore** tab, or the **Storage** tab.
- 2 Right-click the alert for which you want to view the job log, and then select **View Job Log**.
- 3 Do any of the following:
 - To search for a specific word or phrase, click **Find**. Type the text you want to find, and then click **Next**.
Be sure to expand all sections of the job log. The Find feature searches only the expanded sections of the job log.
 - To print the job log, click **Print**. To print the log, you must have a printer attached to your system and configured.
 - To save the job log as an .html file or a .txt file, click **Save As** and then select the file name, file location, and file type.

Responding to active alerts

You can respond to active alerts and continue or cancel the operation, depending on the alert condition. By default, Backup Exec displays all enabled alerts, and all

alerts that require a response. If you have set filters, only those alerts that are selected appear in addition to any alerts that require a response.

If you click **Close** on the alert response dialog box, the dialog box closes, but the alert remains active. To clear the alert, you must select a response such as **OK**, **Yes**, **No**, or **Cancel**. You can configure automatic responses for some alert categories.

See [“Configuring alert categories”](#) on page 302.

Some alerts provide a Unique Message Identifier (UMI) code. This code is a hyperlink to the Symantec Technical Support website. You can access the technical notes that are related to the alert.

To respond to an active alert

- 1 Access the **Active Alerts** pane on the **Home** tab, the **Backup and Restore** tab, or the **Storage** tab.
- 2 Right-click the alert that you want to respond to, and then click **Respond** or **Respond OK**.
- 3 Click a response for the alert, such as **Respond OK** or **Respond**.

Clearing all informational alerts manually

You can configure individual alert categories to be cleared automatically after a certain period of time. Informational alerts may be generated often, so you may want to clear all informational alerts manually before the system moves them automatically.

To clear all informational alerts manually

- 1 Access the **Active Alerts** pane on the **Home** tab, the **Backup and Restore** tab, or the **Storage** tab.
- 2 Right-click an informational alert, and then select **Clear All Informational Alerts**.

See [“Configuring alert categories”](#) on page 302.

Setting up notification for alerts

You can configure Backup Exec to notify recipients when alerts occur. Setting up notifications for alerts is a three-step process.

Table 7-4 How to set up notification for alerts

Step	Action
Step 1	Configure the method you want to use to notify the recipient. The notification methods are text message or email. See “Configuring email or text message notification for alerts” on page 295.
Step 2	Add the contact information for the people or groups that you want to receive notifications. See “Adding recipients for notification” on page 297. See “Adding a recipient group for alert notifications” on page 299.
Step 3	Assign each recipient to the receive notifications for specific alert categories. See “Assigning recipients to receive notifications for specific alert categories” on page 304. See “Configuring alert categories” on page 302. See “Sending a notification when a job completes” on page 304.

Configuring email or text message notification for alerts

You can set up Backup Exec to send email or text messages to specified recipients when an alert occurs. Email notification requires an email account to be used as the sender. For example, you might want to use an email account for the backup administrator or the IT administrator. To configure email notifications, enter the name of the sender's mail server, the port number that the server uses, and the sender's name and email address. You can also set up Backup Exec to authenticate the emails that are sent for alerts.

Table 7-5 Types of notification

Notification type	Description
Email notification	Backup Exec uses SMTP for email notifications and supports authentication and Transport Layer Security (TLS). Notification email messages can be sent to Microsoft Outlook, Lotus Notes, and web-based email applications, such as Gmail or Yahoo mail.

Types of notification (*continued*)

Notification type	Description
Text message notification	<p>For a text message notification, Backup Exec attempts to format the message to contain fewer than 144 characters to meet text messaging protocol restrictions. By limiting a notification to fewer than 144 characters, the notification is more likely to be sent in a single text message instead of broken up into multiple messages. However, the text messaging service provider determines how the notifications are delivered.</p> <p>Text message notifications are sent in the following formats:</p> <ul style="list-style-type: none"> ■ Job-related notification: Backup Exec: <Server Name> : <Job Name> : <Status> ■ Alert-related notification: Backup Exec: <Server Name> : <Alert Type>

After the sender's email information has been entered, then information about recipients can be set up.

Note: An SMTP-compliant email system, such as a POP3 mail server, is required for email notifications.

Note: After you configure email or text message notification, you cannot remove the configuration to disable notifications. However, you can disable notification for individual recipients.

To configure Backup Exec to send email or text notification for alerts

- 1 Click the Symantec Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Email and Text Notification**.
- 3 Set up the sender's email and/or text messaging information.

If you want to send alert notifications by email

Do the following:

- Under **Email configuration**, enter the name of the mail server, the port number that the mail server uses, and the sender's name and email address.
- If you want to authenticate the email information that you entered, check **Enable email authentication**, and then enter the user name and password for the sender's email account.

If you want to send alert notifications by text message

In **Text message service provider address**, enter the fully-qualified domain name of the sender's text messaging service provider.

Example: If a company called "MyPhone" provides text messaging services, then enter "MyPhone.com" in the **Text message service provider address** field. You can override this default address for the individuals who do not use this provider.

Note: Text message notification is sent as SMTP mail to an email address that is provided by a text messaging service provider. To enable notification by text message, you must enter the information about the sender's email account in the **Email configuration** section in addition to the fully-qualified domain name of the default text messaging service provider.

4 Click **OK**.

You can now add information about the recipients who need to receive alert notifications.

See ["Adding recipients for notification"](#) on page 297.

Adding recipients for notification

Individuals or groups can be set up to receive notifications when alerts occur in Backup Exec. When you set up an individual recipient, you indicate whether the person wants to receive notifications by email, text message, or both. A group recipient contains the individual recipients that you select. Each individual within a

group receives notifications by the method that is indicated for the individual; email, text message, or both.

Note: Information about the notification sender must be configured before recipients can be configured.

See [“Configuring email or text message notification for alerts”](#) on page 295.

To add recipients for email or text message notification

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Notification Recipients**.
- 3 On the **Manage Recipients** dialog box, click **Add a recipient**.
- 4 In the **Name** field, type the name of the recipient.
- 5 Select the method of notification for this recipient:

To send emails to this recipient

Check **Send notifications by email**, and then enter the person's email address.

To limit the number of emails that are sent within a specific amount of time, check **Send no more than x emails within x minutes/hours**, and then enter the maximum number of emails to send in a specific time period.

To send text messages to this recipient

Check **Send notifications by text message**, and then enter the person's cell phone number and text message service provider's address.

Note: If the recipient is located in a different country, you must include the exit code of the country from which the message is generated and the country code for the country in which the recipient is located. For example, the exit code for the United States is 011. The country code for Italy is 39. To send a message from the United States to a recipient that is located in Italy, enter **011 39** and the recipient's phone number.

The phone number can include spaces and the following characters:

- Opening and closing quotes
- Period
- Plus sign
- Dash
- Opening and closing parentheses
- Forward slash

To limit the number of text messages that are sent within a specific amount of time, check **Send no more than x text messages within x minutes/hours**, and then enter the maximum number of text messages to send in a specific time period.

6 Click **OK**.

You can now assign the recipients to the alert categories for which they should receive notifications.

See [“Assigning recipients to receive notifications for specific alert categories”](#) on page 304.

Adding a recipient group for alert notifications

Groups are configured by adding recipients as group members. A group contains one or more recipients and each recipient receives the notification message. A group can only include individuals. A group cannot contain other groups.

To add a recipient group for alert notifications

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Notification Recipients**.
- 3 On the **Manage Recipients** dialog box, click **Add a group**.
- 4 In the **Name** field, type a unique name for this notification group.
- 5 To add members to the group, select recipients from the **All recipients** list, and then click **Add** to move them to the **Selected recipients** list.

To remove members from the group, select recipients from the **Selected recipients** list, and then click **Remove** to move them to the **All recipients** list.

- 6 When you have completed the group, click **OK**.

The group is added to the list of recipients on the **Manage Recipients** dialog box.

See [“Adding recipients for notification”](#) on page 297.

Removing a recipient from a group

When you remove a recipient from a group, the recipient no longer receives the notifications that the group is configured to receive. The recipient continues to receive notifications for which the recipient is configured to receive as an individual recipient.

To remove a recipient from a group

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Notification Recipients**.
- 3 On the **Manage Recipients** dialog box, double-click the group that contains the recipient.
- 4 Under **Selected recipients**, select the recipient that you want to remove, and then click **Remove**.

Disabling email or text message alert notification for a recipient

If a person no longer wants to receive alert notifications, you can disable notifications for that person.

Note: Instead of disabling notifications completely, you can also change the alert categories for which a recipient receives notifications.

See [“Disabling notifications for a specific alert category”](#) on page 305.

To disable email or text message alert notification for a recipient

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Notification Recipients**.
- 3 On the **Manage Recipients** dialog box, select the name of the recipient, and then click **Edit**.
- 4 Do any of the following:
 - To disable email notifications, clear the **Send notifications by email** check box.
 - To disable text message notifications, clear the **Send notifications by text message** check box.
- 5 Click **OK**.

Editing recipient notification properties

You can edit the recipient notification properties at any time and change the recipient information, such as an email address or cell phone number. For a group, you can add recipients to the group or remove recipients from the group.

To edit the recipient notification properties

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Notification Recipients**.
- 3 On the **Manage Recipients** dialog box, select the recipient that you want to edit.
- 4 Click **Edit**.
- 5 Edit the properties for the selected recipient.
- 6 Click **OK**.

Deleting recipients

You can delete the recipients that do not want to receive notification messages. The recipient is permanently removed upon deletion. Alternatively, you can disable notification for recipients.

See [“Disabling notifications for a specific alert category”](#) on page 305.

See [“Disabling email or text message alert notification for a recipient”](#) on page 300.

To delete a recipient

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Notification Recipients**.
- 3 On the **Manage Recipients** dialog box, select the recipient that you want to delete.
- 4 Click **Delete**.
- 5 Click **Yes** to confirm that you want to delete this recipient.
- 6 Click **OK**.

Configuring alert categories

Alert categories are the events or the conditions that cause alerts. Alert categories encompass many circumstances or problems that affect the system, jobs, media, or storage sources. Each alert category can include one or more events that generate an alert. For example, a Job Failed error may occur for many reasons. The alert types can help you to determine which alerts need immediate attention and which alerts require a response. You can set up alert categories to enable or disable alerts and to determine what actions should take place when an alert occurs.

Most alerts are enabled by default, however the following alert categories are initially disabled:

- Backup job contains no data
- Job Start
- Job Success

You can disable alert categories if they are informational or warning alerts. You cannot disable alert categories if they are error or attention required alerts.

Each time you change the alert configuration, it is recorded in the audit log. You can view the audit log at any time to view the changes that were made to the alert category.

To configure alert category properties

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Alert Categories**.
- 3 Under **Alert category**, select the alert that you want to configure.

4 Under **Category Properties**, select the appropriate options.

Enable alerts for this category	Enables or disables the alert. You cannot disable error and attention required alerts.
Include the job log with email notifications	Sends the job log to the recipient that is configured for notification. If you select this option, be sure to select at least one recipient from the Send notification to the following recipients area at the bottom of the dialog box.
Record event in the Windows Event Log	<p>Enters the alert into the Windows Event Viewer. The Windows Event log displays all the property information for the alert.</p> <p>If a link appears in the Windows Event log you can search the Symantec Technical Support website for information about the Event ID.</p>
Send SNMP notifications	Indicates whether SNMP notifications are enabled or cleared for the alert. SNMP must be installed to use this option.
Automatically clear after X hours/minutes	<p>Lets you enter the number of minutes, hours, or days you want the alert to remain active before it is cleared.</p> <p>Note: Error alerts cannot be cleared automatically, so this option is disabled for error alerts.</p>
Respond with	<p>Indicates the response that you want Backup Exec to send automatically when the alert is cleared. This option is available only for the Media Overwrite and Media Insert alert categories and only when the Automatically clear after x days/hours/minutes option is selected. The choices are Cancel, No, Yes, or OK.</p>

Send notification to the following recipients

Lets you select the name of a recipient to notify when this type of alert occurs. You must have recipients configured to use this option.

If the recipient to which you want to send notifications is not in the list, click **Manage Recipients** to add the recipient.

- 5 Repeat steps 2 - 4 to configure additional alert categories.
- 6 Click **OK** to save the properties that you selected.

Assigning recipients to receive notifications for specific alert categories

After you have set up notification recipients, you should determine the alert categories for which they should receive notifications. For example, some recipients may only want to receive notifications about job failures and tape errors while other recipients may want to receive notifications for all error alert categories.

To assign recipients to receive notifications for specific alert categories

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Alert Categories**.
- 3 Select an alert category from the list.
- 4 If the category is not enabled, click **Enable alerts for this category**.
- 5 Under **Send notification to the following recipients**, check the check box next to the name of every recipient that needs to receive notifications for the selected alert category.
- 6 Click **OK**.

Sending a notification when a job completes

You can assign recipients to be notified when a job completes. Recipients must be set up before you can set up notification.

To send a notification when a job completes

- 1 Create a new job or edit an existing job.
- 2 On the **Backup Options** dialog box, in the left pane, click **Notification**.

- 3 Select the check box for each recipient that you want to notify when each type of job completes.
- 4 To send the job log with the notification to an email address, check **Include job log in email notifications**.
- 5 You can continue selecting other options from the **Options** dialog box or click **OK**.

Notification options for jobs

When you set up or edit a job, you can select recipients to receive notification when the job completes.

See [“Sending a notification when a job completes”](#) on page 304.

Table 7-6 Notification options for jobs

Item	Description
Recipient name	Shows the names of the individual and group recipients.
Recipient type	Indicates Recipient for an individual recipient or Group for a group recipient.
Include job log in email notifications	<p>Enables Backup Exec to include a copy of the job log with the notification. This option applies only to email recipients. The maximum attachment size, in kilobytes, can be configured in the following registry key:</p> <p>HKLM\Software\Symantec\Backup Exec for Windows\Backup Exec\Server\Max Notification Attachment Size</p> <p>The attachment size can also be determined by the settings on your mail server.</p>
Manage Recipients	Lets you add, edit, or delete recipients.
Properties	Lets you view or change the properties of a selected recipient.

Disabling notifications for a specific alert category

When a recipient no longer needs to receive notifications for an alert category, you can stop the notification.

To disable notifications for a specific alert category

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Alerts and Notifications**, and then select **Alert Categories**.
- 3 Under **Alert category**, select the category for which a recipient no longer needs to receive notification.
- 4 Under **Send notifications to the following recipients**, clear the check box next to the recipient for whom you want to stop notification.
- 5 Click **OK**.

See [“Disabling email or text message alert notification for a recipient”](#) on page 300.

Configuring default alert settings

The default alert settings let you enable or disable the pop-up alerts for the four types of alerts and also to determine how long those pop-up alerts remain on the screen. If you disable the pop-up alerts for a particular alert type, that alert type still appears in the list of active alerts in other parts of the Backup Exec Administration Console, unless you have filtered that same alert type from the list of active alerts.

To configure default alert settings

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Backup Exec Settings**, and then select **Alerts**.
- 3 If you want to receive a reminder alert to renew your maintenance contracts, check **Send an alert reminder to renew your maintenance contracts on**, and then enter a date.

- 4 In the **Display pop-up alerts for the following alert types** group box, check the check boxes for the types of alerts that you want to see in a pop-up alert. Clear the check boxes for the types of alerts that you do not want to see in a pop-up alert.

Informational	Informational alerts provide status messages for the conditions that you might want to know about. They do not require a response and are not critical. The pop-up informational alerts have a blue color.
Error	Error alerts indicate the issues that affect job processing or the integrity of your backup. You must respond to them manually. The pop-up error alerts have a red color.
Warning	Warning alerts indicate the conditions that may or may not cause jobs to fail. You should monitor the conditions and take actions to resolve them. The pop-up warning alerts have a yellow color.
Attention required	Attention required alerts indicate the issues that require a response before a job or an operation can continue. You can respond to this type of alert directly from the pop-up. The pop-up attention-required alerts have a purple color.

- 5 In the **Number of seconds to display pop-up alerts** field, enter the amount of time you want pop-up alerts to remain on the screen.
- 6 Click **OK**.

Enabling or disabling pop-up alerts

By default, Backup Exec displays informational, error, warning, and attention required alert types in pop-up alert windows when an error occurs. You can change the default settings so that you see only the types of alerts that are important to you.

To enable or disable pop-up alerts

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Backup Exec Settings**, and then select **Alerts**.

- 3 In the **Display pop-up alerts for the following alert types** group box, check the check boxes for the types of alerts that you want to see in a pop-up alert. Clear the check boxes for the types of alerts that you do not want to see in a pop-up alert.
- 4 Click **OK**.

SNMP traps for Backup Exec alerts

SNMP (Simple Network Management Protocol) is a method by which a network can be monitored from a central location. SNMP-enabled network applications like Backup Exec report to an SNMP console (a management workstation). The console receives messages (traps) from Backup Exec regarding status and error conditions. An MIB is available in the WINNT\SNMP\language directory on the Backup Exec installation media that you can load into your SNMP console.

The Object Identifier prefix for Symantec is:

1.3.6.1.4.1.1302

Backup Exec SNMP traps (messages) have unique object IDs and may include up to four strings.

The following SNMP trap types are supported:

Table 7-7 SNMP traps

Trap Type	Object ID	String 1	String 2	String 3	String 4
Product Start	1302.3.1.1.9.1	Backup Exec: Application initializing	machine name	product, version, revision	
Product Stop	1302.3.1.1.9.2	Backup Exec: Application terminating	machine name	product, version, revision	
Job Canceled	1302.3.1.2.8.2	Backup Exec: Job canceled by Operator	machine name	job name	local or remote Operator name
Job Failed	1302.3.1.2.8.1	Backup Exec: Job failed	machine name	job name	detail message
Storage device requires human intervention	1302.3.2.5.3.3	Backup Exec: Storage device requires attention	machine name	job name	detail message

Table 7-7 SNMP traps (continued)

Trap Type	Object ID	String 1	String 2	String 3	String 4
Robotic library requires human intervention	1302.3.2.4.3.3	Backup Exec: robotic library device requires attention	machine name	job name	detail message
Simplified Disaster Recovery Message	1302.3.1.4.2.1.2	SDR full backup success	machine name	job name	detail message
Backup Exec system error	1302.3.1.1.9.3	The application has encountered an error	machine name	job name	detail message
Backup Exec general information	1302.3.1.1.9.4	Information on normal events	machine name	job name	detail message
Job Success	1302.3.1.2.8.3	The job succeeded	machine name	job name	detail message
Job Success with exceptions	1302.3.1.2.8.4	The job succeeded, but there was a problem	machine name	job name	detail message
Job Started	1302.3.1.2.8.5	The job has started	machine name	job name	detail message
Job Completed with no data	1302.3.1.2.8.6	The job succeeded, but there was no data	machine name	job name	detail message
Job Warning	1302.3.1.2.8.7	The job has a warning	machine name	job name	detail message
PVL Device Error	1302.3.1.5.1.1.1	The device has encountered an error	machine name	job name	detail message
PVL Device Warning	1302.3.1.5.1.1.2	The device has encountered a warning	machine name	job name	detail message
PVL Device Information	1302.3.1.5.1.1.3	Normal device information	machine name	job name	detail message
PVL Device Intervention	1302.3.1.5.1.1.4	Device requires attention	machine name	job name	detail message

Table 7-7 SNMP traps (*continued*)

Trap Type	Object ID	String 1	String 2	String 3	String 4
PVL Media Error	1302.3.1.5.2.1.1	There is an error with the media	machine name	job name	detail message
PVL Media Warning	1302.3.1.5.2.1.2	There may be a problem with the media	machine name	job name	detail message
PVL Media Information	1302.3.1.5.2.1.3	Normal media information	machine name	job name	detail message
PVL Media Intervention	1302.3.1.5.2.1.4	Media requires attention	machine name	job name	detail message
Catalog Error	1302.3.1.5.3.1.1	There is an error with the catalog	machine name	job name	detail message
Tape Alert Error	1302.3.1.5.4.1.1	There is a TapeAlert error	machine name	job name	detail message
Tape Alert Warning	1302.3.1.5.4.1.2	There is a TapeAlert warning	machine name	job name	detail message
Tape Alert Information	1302.3.1.5.4.1.3	Normal TapeAlert information	machine name	job name	detail message
Database Maintenance Error	1302.3.1.5.5.1.1	There is a database maintenance error	machine name	job name	detail message
Database Maintenance Information	1302.3.1.5.5.1.2	Normal database maintenance information	machine name	job name	detail message
LiveUpdate Error	1302.3.1.5.6.1.1	There is a software update error	machine name	job name	detail message
LiveUpdate Warning	1302.3.1.5.6.1.2	There is a software update warning	machine name	job name	detail message
LiveUpdate Information	1302.3.1.5.6.1.3	Normal software update information	machine name	job name	detail message
Install Update Warning	1302.3.1.5.7.1.1	There is an install warning	machine name	job name	detail message

Table 7-7 SNMP traps (continued)

Trap Type	Object ID	String 1	String 2	String 3	String 4
Install Update Information	1302.3.1.5.7.1.2	Normal Install information	machine name	job name	detail message

See [“Installing and configuring the SNMP system service”](#) on page 311.

Installing and configuring the SNMP system service

To receive Backup Exec traps at the SNMP console, you must configure the SNMP system service with the SNMP console's IP address.

SNMP starts automatically after installation. You must be logged on as an administrator or a member of the Administrators group to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure.

To install the SNMP system service and configure it to send traps to the SNMP console

- 1 From the Windows Control Panel, select **Add/Remove Programs**.
- 2 Click **Add/Remove Windows Components**.
- 3 In Add/Remove Windows Components, select **Management and Monitoring Tools**, and then click **Details**.

When selecting the component, do not select or clear its check box.

- 4 Select **Simple Network Management Protocol**, and then click **OK**.
- 5 Click **Next**.

Installing the Windows Management Instrumentation performance counter provider

Windows Management Instrumentation (WMI) is an infrastructure through which you can monitor and control system resources. Backup Exec includes performance counter and SNMP providers that can be manually installed and used with WMI.

To install the WMI performance counter provider

- 1 Insert the Backup Exec Installation media.
- 2 At the command prompt, type the following:

```
mofcomp <CD Drive Letter>:\winnt\wmi\backupexecperfmon.mof
```

Installing the Windows Management Instrumentation provider for SNMP

Windows Management Instrumentation (WMI) is an infrastructure through which you can monitor and control system resources. Backup Exec includes performance counter and SNMP providers that can be manually installed and used with WMI.

To use the WMI SNMP provider you must set up SNMP notification.

To install the WMI SNMP provider

- 1 Before you install the SNMP provider that is included with Backup Exec, you must have the Microsoft SNMP provider installed on your system.

For more information, refer to your Microsoft documentation.

- 2 Insert the Backup Exec Installation media.
- 3 At the command prompt, type the following:

```
mofcomp <CD Drive Letter>:\winnt\wmi\snmp\eng\bkupexecmib.mof
```

Uninstalling the Windows Management Instrumentation performance counter provider

You must uninstall the Windows Management Instrumentation (WMI) performance counter provider and the WMI SNMP provider separately.

To uninstall the WMI performance counter provider

- ◆ At the command line, type:

```
mofcomp <CD Drive  
Letter>:\winnt\wmi\deletebackupexecperfmon.mof
```

Uninstalling the Windows Management Instrumentation provider for SNMP

You must uninstall the Windows Management Instrumentation (WMI) performance counter provider and the WMI SNMP provider separately.

To uninstall the WMI SNMP provider

- ◆ At the command line, type:

```
Smi2smir /d Backup_Exec_MIB
```

Disk-based and network-based storage

This chapter includes the following topics:

- [Features and types of disk-based storage and network-based storage](#)
- [Storage trending statuses for disk storage and virtual disks](#)
- [Setting low disk space thresholds on disk-based storage](#)
- [Configuring disk storage](#)
- [How to restore data from a reattached or reinserted disk-based storage device](#)
- [Configuring disk cartridge storage](#)
- [How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage](#)
- [Backup sets](#)

Features and types of disk-based storage and network-based storage

Features of disk-based storage include the following:

- Automatic discovery of locally accessible disk volumes.
- Disk space monitoring. Alerts are sent when the disk space thresholds that you set are reached.
- Storage trending analysis that provides predictions of low disk space for disk storage and virtual disks.

- Data lifecycle management, which automatically deletes expired backup sets and reclaims the disk space for use by new backup sets.

Disk-based storage includes the following types of storage:

Table 8-1 Types of disk-based storage

Types of disk-based storage	Description
Disk storage	<p>Disk storage is a location on a locally attached internal hard drive, a USB device, a FireWire device, or a network-attached storage device to which you can back up data.</p> <p>See “Configuring disk storage” on page 318.</p>
Disk cartridge devices	<p>Disk cartridges are a type of storage that usually remains attached to the Backup Exec server while you remove the media, such as RDX. If you are not sure if the storage has removable media, you can open the Computer folder on your Windows computer. The devices that contain removable media are listed.</p> <p>See “Configuring disk cartridge storage” on page 328.</p>
Deduplication disk storage	<p>Deduplication disk storage is a disk-based backup folder that is located on the Backup Exec server and which provides integrated deduplication. You must install the Symantec Backup Exec Deduplication Option to use this data-reduction strategy that optimizes storage and network bandwidth.</p> <p>See “About the Deduplication Option” on page 795.</p>
Storage arrays and virtual disks	<p>Storage arrays and virtual disks are available with the Symantec Backup Exec Storage Provisioning Option, which lets you configure, manage, and monitor storage arrays that are attached to the Backup Exec server.</p> <p>See “Features of the Storage Provisioning Option” on page 1268.</p>

Table 8-1 Types of disk-based storage (continued)

Types of disk-based storage	Description
Vault stores	<p>Vault stores are disk-based containers for the archived data that Backup Exec archives from one server. You must install the Symantec Backup Exec Archiving Option before you can use vault stores.</p> <p>See “About the Archiving Option” on page 1313.</p>

Network-based storage includes NDMP servers, OpenStorage devices, and the Remote Media Agent for Linux.

Table 8-2 Types of network storage

Type of storage	Description
NDMP servers	<p>NDMP servers are network-attached storage (NAS) that supports the Network Data Management Protocol (NDMP) to allow the use of devices that are attached to the servers.</p> <p>See “Features of the NDMP Option” on page 1170.</p>
OpenStorage devices	<p>OpenStorage devices are network-attached storage that supports Symantec's OpenStorage technology.</p> <p>See “Configuring an OpenStorage device” on page 804.</p>
Remote Media Agent for Linux	<p>The Remote Media Agent for Linux lets you back up data from remote computers to the storage devices that are directly attached to a Linux server. You can also back up to a simulated tape library on a Linux server.</p> <p>See “About the Remote Media Agent for Linux” on page 1241.</p>

See [“Storage trending statuses for disk storage and virtual disks”](#) on page 316.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 335.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 451.

See [“Backup sets ”](#) on page 342.

Storage trending statuses for disk storage and virtual disks

Backup Exec gathers disk usage information for disk storage and virtual disks. Backup Exec then performs statistical analysis of used disk space and free disk space. The analysis provides an estimate of how many days remain before the disk storage or virtual disk is full.

Alerts provide information about whether the current disk space resources are sufficient, and can help you plan when to increase disk space.

Table 8-3 Storage trending statuses

Storage trending status	Description
Remaining storage: x days	An estimate of the remaining number of days of storage space, based on the current usage of disk space.
History of used space is still being gathered	<p>This status may appear for any of the following reasons:</p> <ul style="list-style-type: none">■ The disk storage device has not been configured long enough to get a statistical estimate. <p>Note: After you create disk storage, Backup Exec may take approximately one month to gather enough information to provide a storage estimate.</p> <ul style="list-style-type: none">■ This storage may be on a managed Backup Exec server that is currently in a rolling upgrade.
Current storage is sufficient	The environment contains enough disk space to meet storage requirements for the next 30 days.
No estimate due to an inconclusive history of used space	A storage trend cannot be obtained. Unusual increases or decreases in the amount of free disk space in the last 30 days can cause this status.

Table 8-3 Storage trending statuses (*continued*)

Storage trending status	Description
Not enough statistical information is available	Backup Exec has not collected enough sample data for statistical analysis. Note: After you create disk storage, Backup Exec may take approximately one month to gather enough information to provide a storage estimate.

See [“Configuring disk storage”](#) on page 318.

See [“Features of the Storage Provisioning Option”](#) on page 1268.

Setting low disk space thresholds on disk-based storage

You can change the default values for three levels of low disk space conditions. When the storage device's used capacity reaches this threshold, Backup Exec sends an alert, and the data lifecycle management feature immediately searches the device for expired backup sets that it can delete.

Data lifecycle management also runs on a disk cartridge if the cartridge reaches capacity during a backup job. The backup may not need to span to another cartridge if data lifecycle management deletes enough expired backup sets. If a job does span and you insert a new disk cartridge that is full, data lifecycle management deletes expired backup sets on the new cartridge.

To set low disk space thresholds on disk-based storage

- 1 On the **Storage** tab, double-click the storage on which you want to change the low disk space thresholds.
- 2 In the left pane, click **Properties**.
- 3 Change the value on any or all of the following properties:
 - **Low disk space.**
The first low disk space threshold at which you want Backup Exec to send an alert. The default value is 25%.
 - **Low disk space - Warning.**
The second low disk space threshold at which you want Backup Exec to send an alert. The default value is 15%. This threshold must be less than the Low disk space threshold.

- **Low disk space - Critical.**

The third low disk space threshold at which you want Backup Exec to send an alert. The default value is 5%. This threshold must be less than the Warning threshold.

4 Click **Apply**.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 335.

Configuring disk storage

Disk storage is a location on a locally attached internal hard drive, a USB device, a FireWire device, or a network-attached storage device to which you can back up data. You specify how long you want to keep the data that you back up to disk storage when you create a backup job. Backup Exec's data lifecycle management feature automatically deletes expired backup sets and reclaims the disk space. If you want to keep the backup data longer than the period that you specify when you create the backup job, you should create a duplicate backup job. A duplicate backup job can copy the backup data from the original storage device to tape media or to disk cartridge media, which you can then send for long-term or off-site storage. You can also keep the backup sets from automatically expiring by retaining the backup sets. Backup Exec then retains all dependent backup sets as well.

To be eligible for configuration as disk storage, a disk must have at least 1 GB of disk space and cannot be configured as deduplication disk storage. Although you can configure disk storage and deduplication disk storage on the same disk, it is not recommended.

When you create disk storage on a disk that is attached to the network, you must specify the path to an existing share. You should use the server name in the UNC path rather than an IP address.

Note: Before you create the disk storage on a network share, you must give read and write permissions to the Backup Exec service account. The Backup Exec service account is on the Backup Exec server that you want to access the network share.

When you create disk storage on a local disk, Backup Exec lets you specify any of the following locations:

- Volumes with or without drive letters.
You can create only one disk storage on a volume.
- Unformatted partitions.

Backup Exec formats and partitions the drive for you, if necessary.

- Drives that do not have partitions.

Backup Exec creates a folder named BEControl on the root of the volume. Do not delete or edit the contents of the BEControl folder, and do not copy it to other volumes or drive letters.

In Windows Explorer, the backup files that the disk storage device contains display with a .bkf file extension. Each disk storage device also contains a file named changer.cfg and a file named folder.cfg, which store information about the backup files. Do not delete or edit the changer.cfg or folder.cfg files.

A subfolder with a prefix of IMG in the name may display in a disk storage device. This subfolder appears if the option to enable Granular Recovery Technology (GRT) was selected for backup, or if you select the disk storage device as storage for backup data.

You must use the **Configure Storage** wizard to create disk storage. In the **Configure Storage** wizard, Backup Exec provides a list of disks on which you can create disk storage. The disks do not appear in the list in the alphabetical order of the drive letter. Instead, the disk that appears first in the list has the most amount of disk space. You can select any disk that you want, but the disk that Backup Exec recommends for use appears at the top of the list. The disk that you use as the system drive always appears last in the list. Symantec recommends that you do not configure disk storage on the system drive.

Note: If Windows data deduplication is enabled on the disk storage volume, Backup Exec excludes the backup data in the folder \BEData from deduplication, unless the \BEData folder already exists. Backup Exec must exclude backup data from deduplication for you to use Simplified Disaster Recovery (SDR) to perform a local recovery of the Backup Exec server.

If Windows data deduplication is enabled on the disk storage volume, local disaster recovery using SDR fails. The Windows Preinstallation Environment (Windows PE) that SDR uses cannot read the files that Windows data deduplication processes.

To configure disk storage

- 1 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.
- 2 Do one of the following:

If the Central Admin Server Option is not installed	Click Disk-based storage , and then click Next .
---	--

If the Central Admin Server Option is installed

Do the following in the order listed:

- Select the Backup Exec server on which you want to configure storage, and then click **Next**.
- Click **Disk-based storage**, and then click **Next**.

- 3 Click **Disk storage**, and then click **Next**.
- 4 Enter a name and description for the disk storage device, and then click **Next**.
- 5 Specify if you want to create the disk storage device on a local disk or on a network share, enter the location or path, and then click **Next**.
- 6 Specify how many write operations to let run at the same time on this disk storage device, and then click **Next**.
- 7 Review the summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To configure the disk storage device

Click **Finish**.

See [“Changing the location of a disk storage device”](#) on page 320.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 335.

Changing the location of a disk storage device

You can change the location of an existing disk storage device. You must have a different volume available to which you can move the files in the \BEData folder.

Note: When you copy files from the original disk storage device to the new location, do not copy .cfg files.

To change the location of a disk storage device

- 1 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.
- 2 Click **Disk storage**, and then click **Next**.
- 3 Enter a different name and description than the original disk storage device, and then click **Next**.
- 4 Specify a different drive letter than the original disk storage device, and then click **Next**.
- 5 Specify the number of write operations that you want to let run at the same time on this disk storage device, and then click **Next**.
- 6 Review the summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To configure the disk storage device

Click **Finish**.

- 7 In Windows Explorer, copy and paste the following files from the \BEData folder on the original volume to the \BEData folder on the new volume:
 - .Bkf files
 - Any subfolders with a prefix of IMG in the name
- 8 In Windows Explorer, delete all of the files from the original disk storage device.
- 9 On the Backup Exec Administration Console, on the **Storage** tab, right-click the original disk storage device, and then click **Delete**.
- 10 Rename the new disk storage device with the name of the original disk storage device.
- 11 Right-click the new disk storage device, and then click **Inventory and Catalog**.
See [“Inventorying and cataloging a storage device”](#) on page 454.
See [“Configuring disk storage”](#) on page 318.

Editing disk storage properties

You can edit disk space management settings for the disk storage device.

To edit disk storage properties

- 1
- On the **Storage** tab, double-click the storage for which you want to edit properties.
- 2
- In the left pane, click **Properties**.
- 3
- Edit any of the following options:

Name	Displays the name of the disk storage. You can edit this field.
Description	Displays a description of the disk storage. You can edit this field.
Limit Backup Exec to read-only operations	<p>Prevents Backup Exec from deleting expired backup sets on this disk storage when you reattach the disk storage to the Backup Exec server. Otherwise, Backup Exec's data lifecycle management feature deletes any backup sets that are expired and reclaims the disk space.</p> <p>The default value is No.</p> <p>This option applies only when the disk storage has been detached from the Backup Exec server for the number of days that you specify in the global setting. The default number of days is 14.</p> <p>See “How data lifecycle management (DLM) deletes expired backup sets on disk-based storage” on page 335.</p> <p>See “Backup sets ” on page 342.</p> <p>See “How to restore data from a reattached or reinserted disk-based storage device” on page 327.</p>
Maximum file size	<p>Displays the maximum file size on the disk storage. The data from the backup job is contained in a file on the disk.</p> <p>The default value is 50 GB or the capacity of the disk storage.</p>

Preallocate disk space incrementally up to the maximum file size	<p>Creates the file when the backup job starts by preallocating space incrementally, according to the size of the increment that you set in Preallocation increment. As the job uses the disk space, more disk space is preallocated up to the maximum file size. When the job completes, the file size is then reduced to the amount of disk space that the job used.</p> <p>For example, if you enable preallocation and set the preallocation increment to 4 GB, then 4 GB of disk space is preallocated when the job starts. After the job uses 4 GB, then Backup Exec allocates another 4 GB. Disk space continues to be preallocated by 4 GB until the job completes. If the job only uses 13 GB of the 16 GB that was allocated, then the file size is reduced to 13 GB.</p> <p>The default value is Disabled.</p>
Preallocation increment	<p>Displays the amount of disk space by which to increase the file size. The file size increases by this increment as the job requires disk space, up to the maximum file size.</p> <p>The default value is 1 GB.</p>
Auto detect block and buffer size	<p>Indicates if Backup Exec automatically detects the preferred settings for the block size and buffer size for the disk storage.</p> <p>The default value is Enabled.</p> <p>If you disable this setting, you can then choose the block size and buffer size to use.</p>

Block size

Displays the size of the blocks of data that are written to new media in this disk storage device if the option **Auto detect block and buffer size** is disabled. The default is the preferred block size.

Some storage devices provide better performance when larger block sizes are used. The preferred block size can range from 512 bytes to 64 kilobytes or larger. If you use a storage device that supports larger block sizes, you can change the block size. However, if the option to change the block size is unavailable, you must configure the device to use a larger size.

See the manufacturer's documentation for help in configuring the device.

Backup Exec does not ensure that the storage device supports the requested block size. If the requested block size is not supported, it defaults to its standard block size.

If the device does not support block size configuration, this option is unavailable.

Buffer size

Displays the amount of the data that is sent to the disk storage device on each read or write request if the option **Auto detect block and buffer size** is disabled. The buffer size must be an even multiple of the block size.

Depending on the amount of memory in your system, increasing this value may improve storage performance. Each type of storage device requires a different buffer size to achieve maximum throughput.

If the preferred block size is greater than 64 KB, the default buffer size is the same as the default block size. If the preferred block size is less than 64 KB, then the default buffer size is 64 KB.

Low disk space - Critical

Displays the critically low disk space threshold at which you want Backup Exec to send an alert. The color of the capacity bar on the **Storage** tab turns red to indicate critically low available space. Backup Exec sends alerts when the amount of free disk space drops below the low disk space threshold, and again if it drops below the warning threshold. The amount of free disk space does not include the disk space that is reserved for non-Backup Exec operations.

You can change the value of the threshold. This threshold must be less than the warning threshold.

The default value is 5%.

Low disk space - Warning

Displays the low disk space threshold at which you want Backup Exec to send an alert. The color of the capacity bar on the **Storage** tab turns orange to indicate a low disk space condition. If free disk space drops below the warning threshold to the critical threshold, another alert is sent. The amount of free disk space does not include the disk space that is reserved for non-Backup Exec operations.

You can change the value of the threshold. This threshold must be less than the low disk space threshold.

The default value is 15%.

Low disk space

Displays the low disk space threshold at which you want Backup Exec to send an alert. The color of the capacity bar on the **Storage** tab turns yellow to indicate the first of three low disk space conditions. If free disk space drops below this threshold to the amount that is specified in the warning threshold, another alert is sent. If free disk space drops below the warning threshold to the critical threshold, another alert is sent. The amount of disk space does not include the disk space that is reserved for non-Backup Exec operations.

When low disk space reaches this threshold, data lifecycle management immediately searches the device for expired backup sets that it can delete.

You can change the value of the threshold.

The default value is 25%.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 335.

Disk space to reserve for non-Backup Exec operations

Displays the amount of disk space to set aside for applications other than Backup Exec.

The default value is 10 MB.

Auto detect settings

Indicates if Backup Exec automatically detects the preferred settings for read and write buffers for the disk storage.

Buffered read

Indicates the following when the setting is enabled:

- You do not want Backup Exec to automatically detect settings for this disk storage device.
- You want this disk storage to allow buffered read, which is the reading of large blocks of data.

Enabling buffered reads may provide increased performance.

Buffered write

Indicates the following when the setting is enabled:

- You do not want Backup Exec to automatically detect settings for this disk storage device.
- You want this disk storage to allow buffered write, which is the writing of large blocks of data.

Concurrent write sessions

Displays the number of concurrent write operations that you want to allow to this disk storage device.

4 Click **Apply**.

See [“Configuring disk storage”](#) on page 318.

How to restore data from a reattached or reinserted disk-based storage device

If the backup sets on a disk storage device or a disk cartridge expire while that device is detached, Backup Exec deletes the catalogs for those backup sets. To restore from those backup sets at some future time, you must run an inventory and catalog operation on the device when you reattach it. When you run the inventory and catalog operation, Backup Exec sets a new expiration date for each backup set by using the backup set's original retention setting, calculated from the current date. Backup Exec also resets the expiration date for any backup set on the storage device that expires within seven days of the current date.

If you want the backup sets to expire, you can disable the storage device property **Limit Backup Exec to read-only operations**. To find this option, on the **Storage** tab, right-click the storage device, click **Details**, and then click **Properties**. Do not run an inventory and catalog operation. Backup Exec reclaims the disk space on that storage device during data lifecycle management. You can also delete the backup sets.

See [“Inventorying and cataloging a storage device”](#) on page 454.

See [“Backup sets ”](#) on page 342.

Configuring disk cartridge storage

Disk cartridges are a type of storage that usually remains attached to the Backup Exec server while you remove the media, such as RDX. If you are not sure if the storage has removable media, you can open the Computer folder on your Windows computer. The devices that contain removable media are listed.

Backup Exec uses data lifecycle management to automatically expire the backup sets that are stored on disk cartridge media. The backup sets on disk cartridge media are kept for the amount of time that you specify in the backup job properties. Backup Exec automatically reclaims the space as the backup data expires. You can keep the backup sets from automatically expiring by retaining the backup sets. Backup Exec then retains all dependent backup sets as well.

You must use the **Configure Storage** wizard to configure a disk cartridge device.

Available storage operations for disk cartridges and media are dependent on the type of disk cartridge that you have. For example, fewer operations are available for USB memory sticks than are available for RDX devices.

To configure disk cartridge storage

- 1 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.
- 2 Do one of the following:

If the Central Admin Server Option is not installed	Click Disk-based storage , and then click Next .
---	--

If the Central Admin Server Option is installed	Do the following in the order listed: <ul style="list-style-type: none">■ Select the Backup Exec server that you want to configure storage for, and then click Next.■ Click Disk-based storage, and then click Next.
---	--

- 3 Click **Disk cartridge device**, and then click **Next**.
- 4 Enter a name and description for the disk cartridge device, and then click **Next**.
- 5 Specify where the disk cartridge device is located, and then click **Next**.
- 6 Review the summary, and then do one of the following:

To change the configuration	Do the following in the order listed: <ul style="list-style-type: none">■ Click the heading that contains the items that you want to change.■ Make any changes, and then click Next until the summary appears.■ Click Finish.
To configure the disk cartridge device	Click Finish .

See [“Editing disk cartridge properties”](#) on page 329.

Editing disk cartridge properties

You can edit settings for the disk cartridge storage.

See [“Configuring disk cartridge storage”](#) on page 328.

To edit disk cartridge properties

- 1 On the **Storage** tab, double-click the disk cartridge for which you want to edit properties.
- 2 In the left pane, click **Properties**.
- 3 Edit any of the following options:

Name	Displays the name of the disk cartridge. Disk cartridge names cannot exceed 128 characters. You can rename the disk cartridge.
Description	Displays a description of the disk cartridge. You can change the description.
Maximum file size	Displays the maximum file size on the disk cartridge. The data from the job is contained in a file on the disk cartridge. The default value is 50 GB or the capacity of the disk cartridge media.

Preallocate disk space incrementally up to the maximum file size	<p>Creates the file when the job starts by preallocating space incrementally, according to the size of the increment that you set in Preallocation increment. As the job uses the disk space, more disk space is preallocated up to the maximum file size. When the job completes, the file size is then reduced to the amount of disk space that the job used.</p> <p>For example, if you enable preallocation and set the preallocation increment to 4 GB, then 4 GB of disk space is preallocated when the job starts. After the job uses 4 GB, then Backup Exec allocates another 4 GB. Disk space continues to be preallocated by 4 GB until the job completes. If the job only uses 13 GB of the 16 GB that was allocated, then the file size is reduced to 13 GB.</p> <p>The default value is Disabled.</p>
Preallocation increment	<p>Displays the amount of disk space by which to increase the file size if the option to preallocate disk space is enabled. The file size increases by this increment as the job requires disk space, up to the maximum file size.</p> <p>The default value is 1 GB.</p>
Auto detect block and buffer size	<p>Indicates if Backup Exec automatically detects the preferred settings for the block size and buffer size for the disk storage.</p> <p>The default value is Enabled.</p> <p>If you disable this setting, you can then choose the block size and buffer size to use.</p>

Block size

Displays the size of the blocks of data that are written to new media in this disk cartridge if the option **Auto detect block and buffer size** is disabled. The default is the preferred block size.

Some storage devices provide better performance when larger block sizes are used. The preferred block size can range from 512 bytes to 64 kilobytes or larger. If you use the storage that supports larger block sizes, you can change the block size. However, if the option to change the block size is unavailable, you must configure the device to use a larger size.

See the manufacturer's documentation for help in configuring the storage.

Backup Exec does not ensure that the storage device supports the requested block size. If the requested block size is not supported, it defaults to its standard block size.

If the storage does not support block size configuration, this option is unavailable.

Buffer size

Displays the amount of the data that is sent to the disk cartridge on each read or write request if the option **Auto detect block and buffer size** is disabled. The buffer size must be an even multiple of the block size.

Depending on the amount of memory in your system, increasing this value may improve storage performance. Each type of storage requires a different buffer size to achieve maximum throughput.

If the preferred block size is greater than 64 KB, the default buffer size is the same as the default block size. If the preferred block size is less than 64 KB, then the default buffer size is 64 KB.

Low disk space - Critical

Displays the disk space threshold at which the color of the capacity bar on the **Storage** tab turns red to indicate critically low available space. Backup Exec does not send low disk space alerts for disk cartridge devices.

You can change the value of the threshold, but it must be less than the warning threshold.

The default value is 5%.

See [“Storage tab overview in Backup Exec”](#) on page 425.

Low disk space - Warning

Displays the disk space threshold at which the color of the capacity bar on the **Storage** tab turns orange to indicate a low disk space condition. Backup Exec does not send low disk space alerts for disk cartridge devices.

You can change the value of the threshold, but it must be less than the low disk space threshold.

The default value is 15%.

See [“Storage tab overview in Backup Exec”](#) on page 425.

Low disk space

Displays the disk space threshold at which the color of the capacity bar on the **Storage** tab turns yellow to indicate the first of three low disk space conditions. Backup Exec does not send low disk space alerts for disk cartridge devices. When the disk cartridge media reaches this threshold, the data lifecycle management feature immediately searches this disk cartridge media for expired backup sets that it can delete

You can change the value of the threshold.

The default value is 25%.

See [“Storage tab overview in Backup Exec”](#) on page 425.

Auto detect settings

Indicates if Backup Exec automatically detects the preferred settings for read and write buffers for the disk cartridge.

The default value is **Enabled**.

Buffered read

Indicates the following when the setting is enabled:

- You do not want Backup Exec to automatically detect settings for this disk cartridge.
- You want this disk cartridge to allow buffered read, which is the reading of large blocks of data.

Enabling buffered read operations may provide increased performance.

The default value is **Enabled**. If you disable **Auto detect settings**, this setting also changes to **Disabled**.

Buffered write

Indicates the following when the setting is enabled:

- You do not want Backup Exec to automatically detect settings for this disk cartridge.
- You want this disk cartridge to allow buffered write, which is the writing of large blocks of data.

The default value is **Enabled**. If you disable **Auto detect settings**, this setting also changes to **Disabled**.

- 4 Click **Apply**.

Editing disk cartridge media properties

You can edit the properties of the disk cartridge media.

See [“Configuring disk cartridge storage”](#) on page 328.

To edit disk cartridge media properties

- 1 On the **Storage** tab, double-click the disk cartridge device that contains the media.
- 2 In the left pane, click **Media**.

3 Edit any of the following options:

Media label

Displays the media label that Backup Exec assigns automatically or that the administrator assigns.

You can edit the media label, which is limited to 32 characters. Editing the label changes the name of the media in the Backup Exec user interface.

Media description

Displays the original media label if the media is imported media. You can edit the media description to make it a more descriptive label. The description is limited to 128 characters.

Preserve description

Keeps the media description when you select **Yes** in the drop-down box. The media description is kept until an overwrite backup job runs or an erase or label storage operation job runs.

By default, the media description is not retained. This option is set to **No** by default.

How data lifecycle management (DLM) deletes expired backup sets on disk-based storage

Limit Backup Exec to read-only operations

Prevents Backup Exec from deleting expired backup sets on this disk cartridge media when you reinsert the cartridge media into the storage device. If you select **Yes** in the drop-down menu, then Backup Exec's data lifecycle management feature deletes any backup sets that are expired and reclaims the disk space.

The default value is **No**.

This option applies only when the disk cartridge media is absent from the Backup Exec server for the number of days that you specify in the global setting. By default, the number of days that the disk cartridge media must be absent before this option takes effect is 30 days.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 335.

See [“Backup sets ”](#) on page 342.

See [“How to restore data from a reattached or reinserted disk-based storage device”](#) on page 327.

4 Click **Apply**.

How data lifecycle management (DLM) deletes expired backup sets on disk-based storage

Backup Exec uses data lifecycle management (DLM) to automatically delete expired backup sets on disk storage, disk cartridge media, deduplication storage, storage arrays, and virtual disks. You specify how long to keep backup data when you create a backup job that is sent to a disk-based storage device. When the amount of time to keep the backup data expires, the data lifecycle management feature deletes the backup sets and reclaims the disk space unless there are dependent backup sets such as incrementals.

By default, Backup Exec keeps the most recent backup sets that are necessary to restore any backed-up component of a server, even if the backup sets expire. If backup sets are dependent on other backup sets, then Backup Exec does not delete the backup set until all expiration dates on the backup sets are reached. Even if

How data lifecycle management (DLM) deletes expired backup sets on disk-based storage

the backup set is displayed as expired, the data is available until all dependent backup sets expire as well.

For example, you create a backup definition that contains a full backup and an incremental backup for the C: volume on a server. The first full backup runs, followed by the first incremental backup, and then the second incremental backup. The second full backup runs, followed by the third incremental backup, and then the fourth incremental backup. No more backups are run. All of the backup sets created by these backups eventually expire and are deleted by DLM. However, the backup sets that resulted from the second full backup and the third and fourth incremental backups are kept.

Backup Exec keeps these related backup sets because they are the most recent backup sets that you need to restore the C: volume. By keeping the last related backup sets, you have the data to restore the volume.

Warning: DLM deletes all expired backup sets that are created by a one-time backup job. DLM does not keep the last backup set after the retention date expires if the backup set is from a one-time backup.

To prevent the backup sets from being automatically deleted, you can manually retain specific backup sets or you can change the expiration date of the backup set. If you retain a backup set, Backup Exec then retains all dependent backup sets as well.

See [“Retaining backup sets on disk-based storage to prevent them from expiring”](#) on page 346.

See [“Changing the expiration date of backup sets on disk-based storage”](#) on page 345.

DLM searches for expired backup sets to delete from disk-based storage at the following times:

- Every 4 hours.
DLM starts for the first time 4 hours after you install Backup Exec and the Backup Exec service starts, and then runs every 4 hours. If you restart the Backup Exec service, the four-hour DLM cycle also restarts.
- When the low disk space threshold for disk-based storage is reached.
The low disk space threshold is a storage device property. When the storage device's used capacity reaches this threshold, DLM immediately searches the device for the expired backup sets that it can delete.
- When you manually expire a backup set.
When you manually expire a backup set on a standalone Backup Exec server, DLM runs immediately on the storage device on which the backup set is located.

How data lifecycle management (DLM) deletes expired backup sets on disk-based storage

In a Central Admin Server Option (CASO) environment, if you manually expire a backup set from the central administration server, DLM immediately runs on the server on which the backup set was created. The server can be either the central administration server or the managed Backup Exec server. DLM runs only on the storage device from which the backup set was manually expired. If you manually expire a backup set from a managed Backup Exec server, DLM runs immediately on the storage device from which the backup set was manually expired.

To monitor the backup sets that data lifecycle management deletes, you can view the **Backup Set Retention** category in the audit log. You can also run the audit log report to view the backup sets that data lifecycle management deletes.

See [“Configuring audit logs”](#) on page 624.

See [“Audit Log report”](#) on page 655.

Storage options that can affect how data lifecycle management deletes backup sets are described in the following table:

Table 8-4 Storage options in Backup Exec that affect data lifecycle management

Storage option	Description
Allow Backup Exec to delete all expired backup sets	<p>This global setting lets Backup Exec delete expired backup sets, even if they are the last remaining backup sets that you need to restore a server. Use this option when you want to delete backup data after a period of time. Alternatively, you can manually delete backup sets.</p> <p>Warning: If you enable this option, the data that you need to restore a server may not be available.</p> <p>The following conditions may cause you to lose backup sets when you enable this option:</p> <ul style="list-style-type: none"> ■ If the backup sets from the last full backup job expire before the next full backup runs. Ensure that when you create jobs, the backup sets are kept longer than the amount of time between full backups. ■ If the backup job fails or is missed, and is not rerun before the backup sets expire. Monitor any failed or missed jobs, and ensure that you rerun them before the backup sets from the previous full backup expire. <p>Note: In a Central Admin Server Option (CASO) environment, this option is only available on the central administration server. If you enable this option on the central administration server, DLM deletes all expired backup sets on the central administration server as well as on all of the managed Backup Exec servers. This option deletes all expired backup sets on both centrally managed and locally managed Backup Exec servers in a CASO environment.</p> <p>To access this option, click the Backup Exec button > Configuration and Settings > Backup Exec Settings > Storage.</p> <p>See “Editing global settings for storage” on page 439.</p>

Table 8-4 Storage options in Backup Exec that affect data lifecycle management (*continued*)

Storage option	Description
Limit Backup Exec to read-only operations on a disk-based storage device if it has been detached for x number of days	<p>This global setting prevents Backup Exec from deleting expired backup sets on any disk-based storage device as soon as you reattach it. When you limit Backup Exec to read-only operations, you have time to review any expired backup sets and determine if you want to keep them before data lifecycle management deletes them. To keep expired backup sets, you can retain them or change the expiration date.</p> <p>See “Retaining backup sets on disk-based storage to prevent them from expiring” on page 346.</p> <p>See “Changing the expiration date of backup sets on disk-based storage” on page 345.</p> <p>Backup Exec is limited to read-only operations only on the disk-based storage devices that are detached from the Backup Exec server for the specified number of days.</p> <p>To access this option, click the Backup Exec button > Configuration and Settings > Backup Exec Settings > Storage.</p> <p>See “Editing global settings for storage” on page 439.</p>

Table 8-4 Storage options in Backup Exec that affect data lifecycle management (*continued*)

Storage option	Description
Limit Backup Exec to read-only operations on a disk cartridge if it has not been inserted for x number of days	<p>This global setting prevents Backup Exec from deleting expired backup sets on any disk cartridge as soon as you insert it into a disk cartridge device. When you limit Backup Exec to read-only operations, you have time to review any expired backup sets and determine if you want to keep them before data lifecycle management deletes them. To keep expired backup sets, you can retain them or change the expiration date.</p> <p>See “Backup sets ” on page 342.</p> <p>Backup Exec is limited to read-only operations only on the disk cartridges that are detached from the Backup Exec server for the specified number of days.</p> <p>To access this option, click the Backup Exec button > Configuration and Settings > Backup Exec Settings > Storage.</p> <p>See “Editing global settings for storage” on page 439.</p>

Table 8-4 Storage options in Backup Exec that affect data lifecycle management (*continued*)

Storage option	Description
Limit Backup Exec to read-only operations	<p>This device property prevents data lifecycle management from running on a specific disk-based or disk cartridge storage device that you reattach or reinsert it. This option applies only when the storage device has been absent from the Backup Exec server for the number of days that you specify in either of the two previous global settings described in this table.</p> <p>To access this option, on the Storage tab, right-click the device, click Details, and then click Properties.</p> <p>See “Editing disk storage properties” on page 321.</p> <p>See “Editing disk cartridge properties” on page 329.</p> <p>See “Editing the properties of a deduplication disk storage device” on page 811.</p> <p>See “Editing properties of virtual disks on storage arrays” on page 1292.</p>

Table 8-4 Storage options in Backup Exec that affect data lifecycle management (*continued*)

Storage option	Description
Low disk space	<p>This device property is the first of three low disk space conditions. When the storage device's used capacity reaches this threshold, the data lifecycle management feature immediately searches the device for expired backup sets that it can delete.</p> <p>Data lifecycle management also runs on a disk cartridge if the cartridge reaches capacity during a backup job. The backup may not need to span to another cartridge if data lifecycle management deletes enough expired backup sets. If a job does span and you insert a new disk cartridge that is full, data lifecycle management deletes expired backup sets on the new cartridge.</p> <p>To access this option, on the Storage tab, right-click the device, click Details, and then click Properties.</p> <p>See "Setting low disk space thresholds on disk-based storage" on page 317.</p>

You should review the best practices for data lifecycle management at the following URL:

<http://entsupport.symantec.com/umi/V-269-43>

Backup sets

A backup set is a collection of the data that you back up from a single source of content. A single source of content can be a server or a Microsoft Exchange data set, for example. If you select multiple sources of content, Backup Exec creates multiple backup sets. When you run a backup job, Backup Exec creates the backup sets and writes them on storage. To restore data, you select the backup sets that contain the data that you want to restore.

Backup Exec keeps the backup sets that are stored on disk storage and disk cartridge media for as long as you specify in the backup job properties. By default, the amount of time that backup sets are stored is based on the type of backup job and its schedule.

For example, you can specify to keep the backup sets from a full backup for two weeks on a disk-based storage device. After two weeks, the backup sets expire and Backup Exec uses the data lifecycle management feature to delete the backup sets and reclaim that disk space. If you later create an incremental backup job, Backup Exec keeps the full backup sets for two weeks, plus the amount of time that it keeps the incremental backup sets. If you keep the incremental backup sets for four weeks, then Backup Exec keeps the full backup sets for six weeks. The data from a full backup job is kept as long as the data from its associated incremental backup jobs. Backup Exec does not reclaim the disk space for backup sets from a job that depends on another job until the data retention expires for all of the associated jobs. Even if the backup set is displayed as expired, the data is available until all dependent backup sets expire as well.

Backup Exec manages the retention of backup sets differently depending on the type of storage to which you back up the data.

Table 8-5 Storage types and backup set retention

Type of storage	Backup data retention
Disk storage, disk cartridge devices, deduplication disk storage, storage arrays and virtual disks	<p>Backup Exec uses data lifecycle management to automatically delete expired backup sets from disk-based storage. By default, Backup Exec keeps the most recent backup sets that are necessary to restore any backed-up component of a server, even if the backup sets expire. Even if the backup set is displayed as expired, the data is available until all dependent backup sets expire as well.</p> <p>For disk-based storage and any disk cartridges that you reattach to the Backup Exec server after a specified number of days, you can prevent Backup Exec from reclaiming that disk space. A global setting limits Backup Exec to read-only operations on disk-based storage or a disk cartridge if it has been detached for a specified number of days. You can also limit Backup Exec to read-only operations per disk storage or disk cartridge by enabling the setting on the device properties.</p> <p>See “How data lifecycle management (DLM) deletes expired backup sets on disk-based storage” on page 335.</p> <p>See “How to restore data from a reattached or reinserted disk-based storage device” on page 327.</p>

Table 8-5 Storage types and backup set retention *(continued)*

Type of storage	Backup data retention
Tape cartridge media	<p>Backup Exec uses the Advanced Device and Media Management (ADAMM) feature to manage data retention on tape cartridge media. ADAMM expires the backup sets that are stored on media according to a set of rules that you apply to the media. Backup sets are not automatically deleted from tape cartridge media, but they can be overwritten, depending on the rules that you specify. The set of rules that manages tape cartridge media is called a media set. You create media sets that specify append periods, overwrite protection periods, and vaulting periods.</p> <p>See “Default media sets” on page 379.</p>

You can perform the following actions on backup sets that are on disk-based storage:

- Change the expiration date of backup sets to keep them longer or expire them immediately.
- Extend the amount of time that you keep backup sets by retaining them.
- Release any backup sets that are retained to let them expire automatically.

For all backup sets, including those on tape cartridge media, you can perform the following actions:

- Catalog backup sets so that you can view the data that is contained in the backup sets and search for files to restore.
- View the contents of backup sets and browse the backed up data that is contained in them.
- View the system properties and job properties of backup sets.

See [“Changing the expiration date of backup sets on disk-based storage”](#) on page 345.

See [“Retaining backup sets on disk-based storage to prevent them from expiring”](#) on page 346.

See [“Releasing retained backup sets on disk-based storage”](#) on page 347.

See [“Cataloging backup sets”](#) on page 251.

See [“Viewing the contents or properties of backup sets”](#) on page 348.

Changing the expiration date of backup sets on disk-based storage

You can change the expiration date of backup sets that are on disk-based storage to keep them for a longer or shorter amount of time. You can also expire backup sets immediately, if you no longer want to keep them. Data lifecycle management (DLM) automatically deletes expired backup sets from disk-based storage.

When you manually expire backup sets, Backup Exec checks those backup sets to ensure that no other backup sets are dependent upon them. Backup sets from incremental and differential jobs are dependent upon the backup sets that come from the full backup job in the same backup definition. You cannot expire only the backup sets that come from the full backup job because the dependent backup sets would not function without them. When Backup Exec detects dependent backup sets, it gives you the option to expire the backup sets and any dependent backup sets.

When you manually expire a backup set on a standalone Backup Exec server, data lifecycle management (DLM) runs immediately on the storage device on which the backup set is located and deletes the expired backup set. In a Central Admin Server Option (CASO) environment, if you manually expire a backup set from the central administration server, DLM immediately runs on the server on which the backup set was created and deletes the expired backup set. The server can be either the central administration server or the managed Backup Exec server. DLM runs only on the storage device from which the backup set was manually expired to delete the expired backup set. If you manually expire a backup set from a managed Backup Exec server, DLM runs immediately on the storage device from which the backup set was manually expired and deletes the expired backup set.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 335.

To change the expiration date of backup sets on disk-based storage

- 1 On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the backup sets.
- 2 In the left pane, click **Backup Sets**.
- 3 Do one of the following:
 - To change the expiration date for a single backup set, right-click the backup set.
 - To change the expiration date for multiple backup sets, Shift + click or Ctrl + click the backup sets, and then right-click one of the selected backup sets.
- 4 Do either of the following:

To expire backup sets immediately

- Click **Expire**.
Backup Exec displays the backup set and any dependent backup sets.
- To expire the backup sets and any dependent backup sets, click **Yes** or **Yes To All**.

To change the expiration date of backup sets

- Click **Expiration**.
- Enter the new expiration date in the **Expiration** field.
- Click **OK**.

See [“Backup sets”](#) on page 342.

Retaining backup sets on disk-based storage to prevent them from expiring

You can prevent backup sets on disk-based storage from automatically expiring by retaining the backup sets. Backup Exec retains all dependent backup sets as well. For example, if you choose to retain an incremental backup set, Backup Exec retains all backup sets dating back to, and including, the last full backup job. You may need to retain backup sets for legal purposes, such as compliance with data retention laws.

After you retain a backup set, Backup Exec prevents the backup set from expiring indefinitely. If you decide that you no longer need to retain a backup set, you must release it so that it can expire automatically. Data lifecycle management (DLM) automatically deletes expired backup sets from disk-based storage.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 335.

To retain backup sets on disk-based storage to prevent them from expiring

- 1 On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the backup sets that you want to retain.
- 2 In the left pane, click **Backup Sets**.
- 3 Do one of the following:
 - To retain a single backup set, right-click the backup set.
 - To retain multiple backup sets, Shift + click or Ctrl + click the backup sets, and then right-click one of the selected backup sets.
- 4 Click **Retain**.

- 5 In the **Reason to retain backup sets** field, select the reason that you want to retain the backup sets. You can choose from the following options:

Legal

Select this option if the reason for retaining the backup sets is a legal one. You may have to retain backup sets to comply with corporate or regulatory data retention policies.

User defined

Select this option if the reason for retaining the backup sets is something other than a legal one.

- 6 In the **Explanation** field, type any additional information about why you retained the backup sets. Entering an explanation in this field can help remind you why you retained the backup sets or for how long they should be retained.

- 7 Click **OK**.

See [“Backup sets”](#) on page 342.

See [“Releasing retained backup sets on disk-based storage”](#) on page 347.

Releasing retained backup sets on disk-based storage

You can override the retention period for backup sets that are on disk-based storage by manually retaining them. When you choose to retain backup sets on disk-based storage, Backup Exec prevents the backup sets from automatically expiring when their retention period is over. You can manually retain backup sets indefinitely.

If you no longer need the retained backup sets, you can allow them to expire. First, you need to remove the backup sets' retained status. Then Backup Exec expires the backup sets automatically according to the backup sets' storage settings. Data lifecycle management (DLM) automatically deletes expired backup sets from disk-based storage.

To release retained backup sets on disk-based storage

- 1 On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the backup sets that you want to release.
- 2 In the left pane, click **Backup Sets**.
- 3 Do one of the following:
 - To release a single backup set, right-click the backup set.
 - To release multiple backup sets, Shift + click or Ctrl + click the backup sets, and then right-click one of the selected backup sets.
- 4 Click **Retain**.

5 Select **Do not retain**.

6 Click **OK**.

See [“Backup sets ”](#) on page 342.

See [“Retaining backup sets on disk-based storage to prevent them from expiring”](#) on page 346.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 335.

Viewing the contents or properties of backup sets

After you complete a backup job, you can view the data that is contained in the backup sets that are created. Viewing the contents of backup sets can help you to confirm what data was backed up. You may also want to view the contents of backup sets before you run a restore job to verify the data that they contain.

You can also view the following backup set properties:

- Backup source
- Backup date
- Expiration date
- Backup method
- Size
- Location
- Backup set description
- Data encryption
- True image
- Server name
- Catalog file name
- Snapshot

To view the contents or properties of backup sets

- 1 On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device that is related to the backup sets that you want to view.
- 2 In the left pane, click **Backup Sets**.

- 3 Double-click the backup set that you want to view.

Note: On the **Backup and Restore** tab, you must expand the backup source to see the backup set.

- 4 Do either of the following:

To view the contents of the backup sets

In the left pane, click **Contents**.

The contents of the backup set display in the left pane in a tree view. You can expand folders and drives to view their contents in the right pane.

To view the properties of the backup sets

In the left pane, click **Properties**.

See [“Backup sets ”](#) on page 342.

Legacy backup-to-disk folders

This chapter includes the following topics:

- [Legacy backup-to-disk folders](#)

Legacy backup-to-disk folders

In previous versions of Backup Exec, the backup-to-disk feature let you back up data to a folder on a hard disk. These legacy backup-to-disk folders are now read-only. You can continue to inventory, catalog, and restore data from a backup-to-disk folder. You can use remote Simplified Disaster Recovery to perform a disaster recovery from backup-to-disk folders.

In Windows Explorer, the backup-to-disk folders display in the path you specified when you added the folders. The backup-to-disk files display with a .bkf file extension. Each backup-to-disk folder also contains a file named changer.cfg and a file named folder.cfg, which store information about the backup-to-disk files.

Note: Do not delete or edit the changer.cfg or folder.cfg files.

A subfolder with a prefix of IMG in the name may display under a backup-to-disk folder.

In previous versions, Backup Exec created this subfolder when the following conditions were met in a backup job:

- The option to enable Granular Recovery Technology (GRT) was selected.
- A backup-to-disk folder was selected as storage for the backup data.

You cannot send backup data to a backup-to-disk folder. Symantec recommends that you back up data to a disk-based storage device.

See [“Features and types of disk-based storage and network-based storage”](#) on page 313.

See [“Restoring data from a legacy backup-to-disk folder”](#) on page 354.

See [“Editing legacy backup-to-disk folder properties”](#) on page 351.

See [“Changing the location of a legacy backup-to-disk folder”](#) on page 352.

See [“Recreating a legacy backup-to-disk folder and its contents”](#) on page 353.

Editing legacy backup-to-disk folder properties

You can edit the name and description of a legacy backup-to-disk folder. You can also change the settings for buffered reads.

To edit legacy backup-to-disk folder properties

- 1 On the **Storage** tab, double-click the backup-to-disk folder for which you want to edit the properties.
- 2 In the left pane, click **Properties**.
- 3 Edit the appropriate options.
- 4 Click **Apply**.

See [“Legacy backup-to-disk folders ”](#) on page 350.

Changing the name or description of a legacy backup-to-disk folder

You can edit the name and description of a backup-to-disk folder.

See [“Legacy backup-to-disk folders ”](#) on page 350.

To change the name or description of a legacy backup-to-disk folder

- 1 On the **Storage** tab, double-click the backup-to-disk folder for which you want to edit the name or description.
- 2 In the left pane, click **Properties**.

3 Change either of the following properties:

Name	Displays the name of the backup-to-disk folder. Backup-to-disk folder names must not exceed 128 characters.
Description	Displays the description of the backup-to-disk folder.

4 Click **Apply**.

Changing the buffered read setting for a legacy backup-to-disk folder

You can change the buffered read setting for a legacy backup-to-disk folder so that it is enabled, or disabled.

To change the buffered read setting for a legacy backup-to-disk folder

- 1 On the **Storage** tab, double-click the backup-to-disk folder for which you want to change the buffered read and write settings.
- 2 In the left pane, click **Properties**.
- 3 Change any of the following properties:

Auto detect buffers	Indicates if Backup Exec automatically detects the preferred settings for read buffers for the backup-to-disk folder.
Buffered reads	<div>Indicates the following when the setting is Enabled:</div> <div><ul style="list-style-type: none">■ You do not want Backup Exec to automatically detect settings for this backup-to-disk folder.■ You want this backup-to-disk folder to allow buffered reads, which is the reading of large blocks of data.</div> <div>Enabling the buffered reads setting may provide increased performance.</div>

4 Click **Apply**.

Changing the location of a legacy backup-to-disk folder

You can change the location of a legacy backup-to-disk folder.

See [“Legacy backup-to-disk folders”](#) on page 350.

To change the location of a legacy backup-to-disk folder

- 1 On the **Storage** tab, right-click the original legacy backup-to-disk folder, and then click **Delete**.
- 2 Create a folder that has a different name and location than the existing legacy backup-to-disk folder.
- 3 In Windows Explorer, copy and paste all of the following files and folders to the new folder:
 - .Bkf files
 - .Cfg files
 - Any subfolders with a prefix of IMG in the name
- 4 In Windows Explorer, delete all of the files from the original backup-to-disk folder.
- 5 On the Backup Exec Administration Console, on the **Storage** tab, in the **Configure** operation group, click **Configure Storage**.
- 6 Follow the prompts to import the legacy backup-to-disk folder from the new location.

See [“Restoring data from a legacy backup-to-disk folder”](#) on page 354.
- 7 Right-click the backup-to-disk folder, and then click **Inventory and Catalog**.

Recreating a legacy backup-to-disk folder and its contents

If you have deleted a backup-to-disk folder from Backup Exec, but not from the disk, you can recreate the backup-to-disk folder and the files in it. You must know the name and path of the original backup-to-disk folder to retrieve it. If you deleted a backup-to-disk folder from the disk, you cannot recreate it.

See [“Legacy backup-to-disk folders”](#) on page 350.

To recreate a legacy backup-to-disk folder and its contents

- 1 On the **Storage** tab, in the **Configure** operation group, click **Configure Storage**.
- 2 Follow the prompts to import the legacy backup-to-disk folder that you want to recreate.

See [“Restoring data from a legacy backup-to-disk folder”](#) on page 354.
- 3 Right-click the backup-to-disk folder, and then click **Inventory and Catalog**.

See [“Inventorying and cataloging a storage device”](#) on page 454.

Restoring data from a legacy backup-to-disk folder

You must import an existing legacy backup-to-disk folder into Backup Exec to restore data from it. When you import a backup-to-disk folder, you cannot use the root of a volume or an administrative UNC share as the path. The administrative shares are ADMIN\$, IPC\$, and one for each local disk drive letter, such as C\$, D\$, and so on.

Note: When you import a legacy backup-to-disk folder, all of the backup data on that folder is set to expire one year from the data that the folder is imported. You can keep the data longer, or expire the data earlier.

See [“Changing the expiration date of backup sets on disk-based storage”](#) on page 345.

See [“Retaining backup sets on disk-based storage to prevent them from expiring”](#) on page 346.

To restore data from a legacy backup-to-disk folder

1 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.

2 Do one of the following:

If the Central Admin Server Option is not installed

Select **Disk-based Storage**, and then click **Next**.

If the Central Admin Server Option is installed

Do the following in the order listed:

- Select the Backup Exec server on which you want to import a legacy backup-to-disk folder, and then click **Next**.
- Select **Disk-based Storage**, and then click **Next**.

3 Select **Import a legacy backup-to-disk folder**, and then click **Next**.

4 Enter a name and description for the legacy backup-to-disk folder, and then click **Next**.

5 Enter the path to the existing legacy backup-to-disk folder, and then click **Next**.

6 Review the summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To import the legacy backup-to-disk folder Click **Finish**.

- 7 When the legacy backup-to-disk folder appears on the Storage view, right-click the folder, and then click **Inventory and Catalog**.

Note: Depending on the number of backup sets involved, this operation may take a long time to complete. Consider running this operation during off-peak hours.

- 8 On the **Storage** tab, in the **All Storage** view, navigate to the legacy backup-to-disk folder from which you want to restore.
- 9 Right-click the legacy backup-to-disk folder, and then click **Restore**.

Tape storage

This chapter includes the following topics:

- Support for tape drives and robotic libraries
- About the Virtual Tape Library Unlimited Drive Option
- About the Library Expansion Option
- Adding or replacing devices with the Hot-swappable Device Wizard
- Installing Symantec tape device drivers
- Editing tape drive properties
- Viewing tape drive statistics
- Robotic libraries in Backup Exec
- Creating robotic library partitions
- Managing tapes
- Labeling tape media
- How WORM media is used in Backup Exec
- Default media vaults
- Retiring damaged tape media
- Deleting tape media
- Erasing tape or disk cartridge media
- About cataloging tape or disk cartridge media that contains encrypted backup sets
- Associating tape media with a media set

- [Editing tape media properties](#)
- [Tape media rotation strategies](#)

Support for tape drives and robotic libraries

When you install Backup Exec, it automatically recognizes all tape storage that is attached to the Backup Exec server. Tape storage includes tape drives, robotic libraries, virtual tape libraries, and simulated tape libraries.

When you install Backup Exec, support for the following items is included:

- The first robotic library drive per robotic library.
- Every single-drive virtual tape library.

Support for additional drives is available with the Library Expansion Option and the Virtual Tape Library Unlimited Drive Option.

You can use the **Configure Storage** wizard to perform the following actions for tape storage:

- Partition robotic library slots.
- Install Symantec tape device drivers.

Note: You do not need to install Symantec tape device drivers if Backup Exec runs on Windows Server 2012 and later. Symantec kernel-mode drivers and tapeinste.exe are no longer installed if Backup Exec runs on Windows Server 2012 and later.

See [“Installing Symantec tape device drivers”](#) on page 360.

- Replace or add hot-swappable storage on a Backup Exec server without having to restart the server.
- Create media sets to manage the backup data on tapes.

See [“About the Virtual Tape Library Unlimited Drive Option”](#) on page 358.

See [“About the Library Expansion Option”](#) on page 358.

See [“Default media sets”](#) on page 379.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 451.

About the Virtual Tape Library Unlimited Drive Option

When you install Backup Exec, support for every single-drive virtual tape library is included. The Virtual Tape Library Unlimited Drive Option enables support for all additional drives in each virtual tape library. To install the Virtual Tape Library Unlimited Drive Option, add a license.

You can find a list of compatible types of storage in the Backup Exec Hardware Compatibility List:

<http://entsupport.symantec.com/umi/v-269-2>

You can find license information for the Virtual Tape Library Unlimited Drive Option in the Backup Exec Customer Pricing and Licensing Guide:

<http://entsupport.symantec.com/umi/V-269-21>

See “Installing additional Backup Exec options to the local Backup Exec server” on page 76.

About the Library Expansion Option

When you install Backup Exec, support for the first drive in every robotic library is included. The Library Expansion Option enables support for each additional drive in a robotic library.

To install the Library Expansion Option, add a license.

You can find a list of compatible types of storage in the Backup Exec Hardware Compatibility List:

<http://entsupport.symantec.com/umi/v-269-2>

You can find license information for the Library Expansion Option in the Backup Exec Customer Pricing and Licensing Guide:

<http://entsupport.symantec.com/umi/V-269-21>

See “Installing additional Backup Exec options to the local Backup Exec server” on page 76.

Adding or replacing devices with the Hot-swappable Device Wizard

Use the **Hot-swappable Device Wizard** to replace or add hot-swappable storage on a Backup Exec server without having to restart the server.

If you remove and then reconnect Universal Serial Bus (USB) tape devices to the USB port, you must run the **Hot-swappable Device Wizard** to allow Backup Exec to rediscover the devices.

For iSCSI-attached devices, you must list the device as a **Persistent Target** in the iSCSI control panel applet, and then run the **Hot-swappable Device Wizard**. Listing the device as a **Persistent Target** lets Backup Exec rediscover the device whenever you restart the Backup Exec server.

After you start the **Hot-swappable Device Wizard**, you are prompted to close the Backup Exec Administration Console. The **Hot-swappable Device Wizard** waits until any jobs that were processing are completed. The wizard pauses the Backup Exec server and stops the Backup Exec services. You can then add or replace any storage devices. The wizard detects the new device or replaced device, and adds information about the device to the Backup Exec database. The wizard is then completed, and you can reopen the Backup Exec Administration Console.

Any new storage appears in the **Storage** tab, and usage statistics for the storage begin accumulating. You can enable the new storage in a storage device pool.

Any replaced storage appears in the **Storage** tab, in the **All Storage** view with a status of Offline.

Note: Start the **Symantec Hot-swappable Device Wizard** before you add or replace storage.

To add or replace devices with the Hot-swappable Device Wizard

1 Do one of the following:

For iSCSI-attached storage:

In the iSCSI control panel applet, add the storage to the **Persistent Targets** list.

Continue with the next step.

For any other hot-swappable storage:

Continue with the next step.

2 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.

3 When you are prompted for the type of storage that you want to configure, select **Tape Storage**, and then click **Next**.

4 Select **Run the Hot-swappable Device Wizard**, click **Next**, and then follow the on-screen prompts.

Installing Symantec tape device drivers

You can use the **Configure Storage** wizard to install tape device drivers and to delete entries for tape devices that are not available, have been removed, or are turned off.

This topic includes the following information:

[Symantec kernel-mode tape drivers and tapeinst.exe](#)

Before you install Symantec tape device drivers, do the following:

- Ensure that Backup Exec supports the tape device.
You can find a list of compatible types of storage in the Backup Exec hardware compatibility list:
<http://entsupport.symantec.com/umi/V-269-2>
- Run the Windows Device Manager to ensure that it lists the tape device.
- Get the latest Backup Exec device drivers.

To get the latest device drivers installed to the Backup Exec installation folder

- 1 Click the Backup Exec button, select **Technical Support**, and then click **Latest Backup Exec Device Drivers**.
- 2 Run the **Symantec Device Driver Installation Wizard**, in the following procedure.

To install Symantec tape device drivers by using the Configure Storage wizard

- 1 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.
- 2 On the wizard panel, select **Tape Storage**, and then click **Next**.
- 3 Select **Install tape device drivers**, and then click **Next**.
- 4 On the **Symantec Device Driver Installation Wizard**, follow the on-screen prompts.

Symantec kernel-mode tape drivers and tapeinst.exe

Symantec kernel-mode drivers and tapeinst.exe are no longer installed if Backup Exec runs on Windows Server 2012 or later. If the Backup Exec server runs on Windows Server 2012 or later, the **Configure Storage** wizard option **Install tape device drivers** is not available, and tapeinst.exe is not available from the default Backup Exec installation location. Tapeinst.exe is the program that installs Symantec kernel-mode tape class drivers. If the Backup Exec server runs on Windows Server 2008 R2 or earlier, you can run tapeinst.exe manually, or by using the Configure Storage wizard.

The Backup Exec 2014 hardware compatibility list does not list any tape devices that require kernel-mode drivers. If you use older or unsupported tape devices, test the devices before you use Windows Server 2012 and later to determine the functionality in case kernel-mode drivers are required.

Symantec user-mode drivers now access all supported tape devices. User-mode drivers have accessed most tape devices for multiple previous versions of Backup Exec.

User-mode drivers are not visible in the Windows Device Manager. When no kernel-mode drivers are installed, the Windows Device Manager may show the tape drive with a yellow exclamation mark. This display is acceptable and does not prevent the use of the device with Backup Exec and the user-mode drivers. Kernel-mode drivers from Microsoft or from the tape hardware vendor are generally acceptable for installation. Backup Exec will continue to access the tape devices through the user-mode drivers.

To install Symantec tape device drivers on a Backup Exec server that run on Windows Server 2008 R2 or earlier by running tapeinst.exe

- 1 From the Backup Exec installation directory, double-click the tapeinst.exe file.
The default installation directory is C:\Program Files\Symantec\Backup Exec
- 2 On the **Symantec Device Driver Installation Wizard**, follow the on-screen prompts.

Editing tape drive properties

You can edit the following tape drive properties.

See [“Support for tape drives and robotic libraries”](#) on page 357.

To edit tape drive properties

- 1 On the **Storage** tab, double-click the tape drive for which you want to edit properties.
- 2 In the left pane, click **Properties**.

3 Edit any of the following options:

Name

Displays the name of the tape drive. You can edit this field.

Description

Displays the description of the tape drive. You can edit this field.

Hardware compression

Indicates if hardware compression is enabled.

If this option is available, this drive is capable of supporting hardware compression.

If you configure a job to use hardware compression, and hardware compression is disabled on the device, then hardware compression is unavailable and is not used.

Block size

Displays the size of the blocks of data that are written to new media in this tape drive. The default is the preferred block size.

Some devices (for example, LTO devices) provide better performance when larger block sizes are used. The preferred block size can range from 512 bytes to 64 kilobytes or larger. If you use a tape drive that supports larger block sizes, you can change the tape drive's block size. However, if the tape drive does not allow a block size as large as you want, reconfigure the host bus adapter or the tape drive. After you reconfigure the hardware and restart the Backup Exec services, check if the block size that you want to use is available.

See the tape drive manufacturer's documentation for help in configuring the device.

Backup Exec does not ensure that the requested block size is supported by that tape drive. You should check the tape drive specifications to make sure that the block size is supported. If the tape drive does not support a block size, it defaults to its standard block size.

If the tape drive does not support block size configuration, this option is unavailable.

Buffer size

Displays the amount of data that is sent to the tape drive on each read or write request. The buffer size must be equal to the block size or an even multiple of the block size.

Depending on the amount of memory in your system, increasing this value may improve tape drive performance. Each type of tape drive requires a different buffer size to achieve maximum throughput.

Buffer count

Displays the number of buffers that are allocated for this tape drive.

Depending on the amount of memory in your system, increasing this value may improve device performance. Each type of tape drive requires a different number of buffers to achieve maximum throughput.

If you change the buffer count, you may need to adjust the high water count accordingly.

High water count

Displays the number of buffers to be filled before data is first sent to the tape drive, and any time that the tape drive underruns.

The high water count cannot exceed the buffer count. A value of 0 disables the use of high water logic; that is, each buffer is sent to the device as it is filled.

The default setting provides satisfactory performance in most instances; in some configurations, throughput performance may be increased when other values are specified in this field. If you increase or decrease the buffer count, the high water count should be adjusted accordingly. If a tape drive has a high water count default of 0, it should be left at 0.

Reset to default settings

Returns all of the preferred configuration settings to their defaults.

Read single block mode

Indicates if this tape drive reads only one block of data at a time, regardless of the size of the buffer block.

This option is disabled by default.

Write single block mode

Indicates if this tape drive writes only one block of data at a time. This option provides greater control over the handling of data write errors.

Symantec recommends that you select this option if the tape drive is shared.

This option is enabled by default.

Read SCSI pass-through mode

Indicates if this tape drive reads data without going through a Microsoft tape device API. This option allows the data to pass directly through the tape drive and allows more detailed information if device errors occur.

This option is disabled by default.

Write SCSI pass-through mode

Indicates if this tape drive writes data without going through the Microsoft tape device API. This option allows data to pass directly through the device driver and allows more detailed information if device errors occur.

Symantec recommends that you select this option if the tape drive is shared.

This option is enabled by default.

Servers that share this device

Displays the servers that can also use this device.

See [“Sharing storage devices”](#) on page 448.

Media Type

Indicates if barcode rules are enabled for the robotic library that this tape drive is attached to. If barcode rules are enabled, the media types are listed that this tape drive can read from and write to. Backup Exec uses the barcode rules to identify which type of media to use in a drive.

You can configure barcode rules, and enable or disable barcode rules for a robotic library.

See [“Configuring barcode rules for a robotic library”](#) on page 368.

Can Read From

Indicates if this tape drive can read from the media type.

See [“Configuring barcode rules for a robotic library”](#) on page 368.

Can Write To

Indicates if this tape drive can write to this media type.

See [“Configuring barcode rules for a robotic library”](#) on page 368.

- 4 Click **Apply**.

Viewing tape drive statistics

You can view statistics about tape drives.

See [“Support for tape drives and robotic libraries”](#) on page 357.

To view tape drive statistics

- 1 On the **Storage** tab, double-click the tape drive for which you want to view statistics.
- 2 In the left pane, click **Statistics**.

Robotic libraries in Backup Exec

Backup Exec's Advanced Device and Media Management (ADAMM) feature solves the problems that are associated with typical robotic library modules. Backup Exec accesses all of the media in the robotic library and uses the media that belongs to the specified media set. If the backup job exceeds the capacity of a media, Backup Exec searches all of the media that is contained in the robotic library and finds a suitable media to use.

For example, an operator has a robotic library with six slots. The operator inserts six blank tapes and targets backup jobs to various media sets within the robotic library. Backup Exec automatically allocates available tapes in the robotic library. If a job exceeds the capacity of one tape and another overwritable tape is available in the robotic library, the job automatically continues on that tape. When Backup Exec runs out of tapes, it prompts the operator to import overwritable media.

In a robotic library, Backup Exec selects the oldest recyclable media in the library to use first. If more than one media that meets the requirements is found, Backup Exec selects the media in the lowest-numbered slot. For example, Backup Exec selects media in slot 2 before it selects equivalent media in slot 4.

For restore jobs that use robotic libraries, Backup Exec accesses the source media regardless of its sequential placement in the magazine. For example, if the data for a restore job resides on two media in the magazine, the media do not have to be placed in adjacent slots for Backup Exec to restore the data. If Backup Exec

does not find the media that is required for the restore job in the robotic library, an alert is generated that requests the media that is necessary to complete the job.

See “Requirements for setting up robotic library hardware” on page 367.

See “Creating robotic library partitions” on page 373.

Requirements for setting up robotic library hardware

You can configure Backup Exec to work with robotic library drives by making associations between the robotic library’s drives, robotic arm, and Backup Exec. Backup Exec supports serialized drives. Manual configuration of serialized drives is not required.

You can find a list of supported types of storage at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

Ensure that the robotic library hardware is configured as follows:

- Ensure that the robotic arm is set to Random mode. Refer to your robotic library documentation for more information.
- Ensure the following for a multi-LUN robotic library:
 - The controller card is set to support multiple LUNs (if supported).
 - The target LUN for the tape drive is lower than the target LUN for the robotic library.
- Determine which drive is the first drive in the robotic library, and then arrange the SCSI IDs to match the sequence of the drive element addresses. Refer to your robotic library documentation to determine the drive element address for each storage device.
- Ensure that the SCSI ID of the robotic arm precedes the SCSI IDs of the drives in the robotic library. Do not use 0 or 1 because these SCSI IDs are typically reserved for boot devices.

In the following example, if your robotic library has two drives, the drive with the lowest drive element address should be assigned the lower SCSI ID.

Table 10-1 Example configuration for a multi-drive robotic library

Data transfer element (storage devices)	SCSI ID	Drive element address
Robotic arm	4	N/A
Storage device 0	5	00008000

Table 10-1 Example configuration for a multi-drive robotic library (*continued*)

Data transfer element (storage devices)	SCSI ID	Drive element address
Storage device 1	6	00008001

See [“Robotic libraries in Backup Exec”](#) on page 366.

Inventorying robotic libraries when Backup Exec services start

You can set a default so that all robotic libraries are included in the inventory job whenever Backup Exec services are started. Symantec recommends that you enable this default if media is often moved between robotic libraries. Backup Exec may take longer to start.

To inventory robotic libraries when Backup Exec services start

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, click **Storage**.
- 3 Click **Inventory robotic libraries when Backup Exec services start**.
- 4 Click **OK**.

See [“Inventorying a storage device”](#) on page 453.

Configuring barcode rules for a robotic library

If you have barcode support for a robotic library that uses different types of drives, you can create a barcode rule. Backup Exec uses the barcode rules to identify which type of media to use in a drive. When Backup Exec reads the barcode rule, it locates the type of media that corresponds to the prefix or suffix. Backup Exec then mounts the media into a drive that accepts that type of media.

The barcode rules apply to all of the robotic libraries that you enable for barcode rules. By default, barcode rules are disabled.

You can use the **Configure Storage** wizard to configure barcode rules, or you can add, edit, and delete barcode rules in the Backup Exec global settings. You can also enable or disable barcode rules in the robotic library's properties.

See [the section called “Configuring barcode rules for a robotic library by using the Configure Storage wizard”](#) on page 369.

See [the section called “Adding, editing, and deleting barcode rules by using the Backup Exec global settings”](#) on page 370.

See [the section called “Enabling or disabling barcode rules by using robotic library properties”](#) on page 370.

Configuring barcode rules for a robotic library by using the Configure Storage wizard

To configure barcode rules for a robotic library

1 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.

2 Do one of the following:

If the Central Admin Server Option is not installed Click **Tape storage**, and then click **Next**.

If the Central Admin Server Option is installed

Do the following in the order listed:

- Select the Backup Exec server on which you want to configure storage, and then click **Next**.
- Click **Tape storage**, and then click **Next**.

3 Click **Configure barcode rules for a robotic library**, and then click **Next**.

4 Select the robotic library that you want to configure barcode rules for, and then click **Next**.

5 To change the available barcode rules in the list, click any of the following:

To add a new barcode rule

Do the following in the order listed:

- Click **New**.
- Click the drop-down menu, and then select a media type
- Type the vendor name, and the barcode prefix and/or suffix.
- Click **OK**.

To edit an existing barcode rule

Do the following in the order listed:

- Select the barcode rule that you want to edit, and then click **Edit**.
- Make any changes, and then click **OK**.

To delete a barcode rule

Select a barcode rule, click **Delete**, and then click **Yes** to confirm the deletion.

6 Click **Next**.

- 7 Select a tape drive, and then for any of the listed media types, select a check box to specify if the tape drive can read from and write to that media type.

Click **Next**.

- 8 Review the summary, and then do one of the following:

To change the barcode configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next**.
- Click **Finish**.

To complete the barcode configuration

Click **Finish**.

- 9 Ensure that barcode rules are enabled for the appropriate robotic libraries.

Enabling or disabling barcode rules by using robotic library properties

You can enable or disable barcode rules for a robotic library. By default, barcode rules are disabled for robotic libraries.

To enable or disable barcode rules by using robotic library properties

- 1 On the **Storage** tab, double-click the robotic library for which you want to enable barcode rules.
- 2 In the left pane, click **Properties**.
- 3 In the **Barcode rules** field, in the drop-down menu, do one of the following:

To enable all barcode rules for this robotic library Click **Enabled**.

To disable all barcode rules for this robotic library Click **Disabled**.

- 4 Click **OK**.

Adding, editing, and deleting barcode rules by using the Backup Exec global settings

You can add, edit, or delete the barcode rules for a robotic library in the global settings for Backup Exec. All barcode rules changes apply to all of the robotic libraries for which barcode rules are enabled.

To add, edit, or delete a barcode rule by using the Backup Exec global settings

- 1** Click the Backup Exec button, click **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2** In the left pane, click **Bar Code Rules**.
- 3** Do one of the following:

To add new barcode rule

Click **New**, and then continue with the next step.

To edit an existing barcode rule

Select the barcode rule that you want to edit, click **Edit**, and then continue with the next step.

To delete an existing barcode rule

Do the following in the order listed:

- Select the barcode rule that you want to delete, and then click **Delete**.
- Click **Yes** to confirm the deletion, and then click **OK**.

The barcode rule is deleted.

4 Add or change any of the following information:

Media type

Click the drop-down menu and select or change the media type.

Vendor

Enter or change the name of the robotic library manufacturer.

A best practice is to use a vendor name only if there is a specific need for it and you have multiple brands of tape drives that are available

This field is limited to 16 characters.

Barcode prefix

Enter or change the prefix of a barcode to which you want to apply this barcode rule.

A standard practice is to use a prefix to indicate a media-handling policy, such as **CLN** for cleaning media, or **I** for internal use, or **O** for cartridges that are to be taken offsite.

This field is limited to 16 characters.

Barcode suffix

Enter or change the suffix of a barcode to which you want to apply this barcode rule.

A standard practice is to use a suffix to indicate the generation of the media, such as L1, L2, and so on.

This field is limited to 16 characters.

5 Click **OK**.

Initializing a robotic library when the Backup Exec service starts

You can initialize a robotic library whenever the Backup Exec services start.

During startup, if media is in the drives in the robotic library, Backup Exec attempts to return the media to its original drive. If the media cannot be returned to the drive, an error message appears that prompts you to eject the media from the drive.

You can also create a job to initialize a robotic library.

See [“Initializing a robotic library”](#) on page 456.

See [“Robotic libraries in Backup Exec”](#) on page 366.

To initialize a robotic library when the Backup Exec services start

- 1 On the **Storage** tab, double-click the robotic library that you want to initialize.
- 2 In the left pane, click **Properties**.
- 3 In the **Startup initialization** field, in the drop-down menu, click **Enabled**.
- 4 Click **Apply**.

Defining a cleaning slot

Before submitting a cleaning job, you must define a cleaning slot that contains the cleaning tape.

Make sure that the cleaning tape is located in the slot that you defined as the cleaning slot. After defining the cleaning slot, you can set up a cleaning job for the robotic library drive.

See [“Cleaning a robotic library drive”](#) on page 458.

Note: Defined cleaning slots are not inventoried when an inventory job runs.

To define a cleaning slot

- 1 On the **Storage** tab, expand the robotic library, and then double-click **Slots**.
- 2 Double-click the slot that contains the cleaning tape.
- 3 In the **Cleaning slot** field, click the drop-down menu, and then click **Yes**.
- 4 Click **Apply**.

Editing robotic library properties

You can view robotic library properties.

See [“Robotic libraries in Backup Exec”](#) on page 366.

To view robotic library properties

- 1 On the **Storage** tab, double-click the robotic library for which you want to view properties.
- 2 In the left pane, click **Properties**.

Creating robotic library partitions

You can group one or more robotic library slots into partitions. Partitioning robotic library slots provides you with more control over which tape is used for backup jobs.

When you create robotic library partitions, Backup Exec creates a storage device pool for each partition. Jobs that you send to a partition's storage device pool run on the media that is located in the partition's slots. For example, if you set up a partition that contains slots 1 and 2 and you want to run a weekly backup only on the media in these slots, you would submit the job to the storage device pool for the partition that contains slots 1 and 2. The storage device pools for robotic library partitions appear in the drop-down menu for the **Storage** field when you edit the backup job. All storage device pools for a robotic library partition have the same name and display the slot ranges for the partition in parentheses within the name.

Partitions can include any number of robotic library slots.

Depending on the robotic library configuration, the first slot may be numbered 1 or 0. If the robotic library uses a zero-based slot configuration, the **Configure Storage** wizard uses slot 0 as the first slot for partition 1 and adjusts the starting slot accordingly for all other partitions.

See ["Reassigning a slot base number for robotic libraries"](#) on page 376.

When the robotic library is partitioned, Backup Exec searches for the oldest recyclable media in the specified partition only. If more than one media is found that meets the requirements, Backup Exec then selects the media in the lowest-numbered slot; for example, media in slot 2 is selected before equivalent media in slot 4.

You can create a partitioning scheme that best suits your environment. For example, some administrators may create partitions based on users and groups, while others may create partitions according to operation types.

To create robotic library partitions

1 Do one of the following:

To view all of the robotic libraries that you can configure partitions for

Do the following in the order listed:

- On the **Storage** tab, in the **Configure** group, click **Configure Storage**.
- Select **Tape storage**, and then click **Next**.
- Select **Configure robotic library partitions**, and then click **Next**.
- Select the robotic library to configure partitions for, and then click **Next**.

To configure partitions for a specific robotic library Do the following in the order listed:

- On the **Storage** tab, right-click the robotic library for which you want to create partitions.
 - Click **Configure partitions**.
- 2 Specify the number of partitions to create, enter the number of slots for each partition, and then click **Next**.
 - 3 Review the storage configuration summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To configure the partitions

Click **Finish**.

See [“Changing the name of a robotic library partition”](#) on page 375.

Changing the name of a robotic library partition

You can change the name of a robotic library partition.

See [“Creating robotic library partitions”](#) on page 373.

To change the name of a robotic library partition

- 1 On the **Storage** tab, double-click the robotic library partition.
- 2 In the left pane, click **Properties**.
- 3 Change the name of the robotic library partition.
- 4 Click **Apply**.

Adding or removing tape drives in a robotic library partition

You can add or remove tape drives in a robotic library partition.

See [“Creating robotic library partitions”](#) on page 373.

To add or remove tape drives in a robotic library partition

- 1 On the **Storage** tab, double-click the robotic library partition.
- 2 In the left pane, click **Properties**.
- 3 Do either of the following:

To add a tape drive to the robotic library partition

Do the following in the order listed:

- Click **Add**.
- In the devices list that appears, select the tape drive that you want to add, and then click **OK**.

To remove a tape drive from the robotic library partition

Select the tape drive that you want to remove, and then click **Remove**.

- 4 Click **Apply**.

Reassigning a slot base number for robotic libraries

Backup Exec automatically assigns slot base numbers for robotic libraries. If necessary, you can reassign how robotic library slots are displayed in Backup Exec. Slot base numbers in some robotic libraries start at 0, while slots in other robotic libraries start at 1. If the robotic library uses a zero-based slot configuration, you can reassign how the slots are displayed.

See [“Robotic libraries in Backup Exec”](#) on page 366.

To reassign a slot base number for robotic libraries

- 1 On the **Storage** tab, double-click the robotic library for which you want to reassign a slot base number.
- 2 In the **First slot number** field, click the drop-down menu to change the base number.
- 3 Click **Apply**.

Editing robotic library slot properties

You can edit the properties of a robotic library slot.

See [“Robotic libraries in Backup Exec”](#) on page 366.

To edit robotic library slot properties

- 1 On the **Storage** tab, double-click the robotic library.
- 2 In the left pane, click **Slots**.

- 3 Double-click the slot for which you want to view properties and then click **Properties**.
- 4 To view another slot's properties, click the drop-down menu at the top center of the window, and then click a slot number.
- 5 Edit any of the following options.

Cleaning slot

Indicates if this slot has been defined as a cleaning slot. If **Yes** is displayed, this slot has been defined as a cleaning slot.

Before you can submit a cleaning job, you must define a cleaning slot that contains the cleaning tape.

Make sure that the cleaning tape is located in the slot that you define as the cleaning slot. After defining the cleaning slot, you can set up a cleaning job for the robotic library drive.

Note: Defined cleaning slots are not inventoried when an inventory job runs.

See ["Cleaning a robotic library drive"](#) on page 458.

Preserve description

Keeps the media description when you select **Yes** in the drop-down menu. The media description is kept until an overwrite backup job runs or an erase or label storage operation job runs.

By default, the media description is not retained. This option is set to **No** by default.

- 6 Click **Apply**.

Removing or reconfiguring robotic library partitions

You can remove or reconfigure robotic library partitions.

To remove robotic library partitions

- 1 On the **Storage** tab, right-click the robotic library that contains the partitions that you want to remove or reconfigure.
- 2 Click **Configure partitions**.

- 3 Click **Remove all robotic library partitions**, and then click **Next**.
- 4 Review the storage configuration summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To remove the partitions

Click **Finish**.

To reconfigure robotic library partitions

- 1 On the **Storage** tab, right-click the robotic library that contains the partitions that you want to remove or reconfigure.
- 2 Click **Configure partitions**.
- 3 Click **Reconfigure robotic library partitions**, and then click **Next**.
- 4 Specify the number of partitions to create, enter the number of slots for each partition, and then click **Next**.
- 5 Review the storage configuration summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To configure the partitions

Click **Finish**.

See [“Reassigning a slot base number for robotic libraries”](#) on page 376.

Managing tapes

For tape media, you can perform the following actions:

- Protect data from being overwritten.
- Set up media rotation strategies.

- Track the location of media.
- Label media automatically.
- Read and track the media labels that have barcodes.
- Collect and report media statistics.

The Advanced Device and Media Management (ADAMM) feature in Backup Exec automatically selects the tape media for jobs. Backup Exec keeps track of all tape media that is loaded into the attached storage device. Backup Exec also keeps track of the media that is offline and the media that has been placed in media vaults.

For data that is kept on tapes, Backup Exec uses media sets to apply overwrite protection periods and append periods to manage the expiration of backup sets.

A media set consists of the following rules that apply to tape media:

- How long to protect the data on the media from overwrite. This is called the overwrite protection period.
- How long to append data to a media. This is called the append period.
- When and where to send media for vaulting.

Media that are associated with a media set are allocated media. Allocated media have current append and overwrite protection periods. Media that are associated with a media set but have expired overwrite protection periods are recyclable media.

See [“Default media sets”](#) on page 379.

See [“Tape media rotation strategies”](#) on page 412.

See [“Creating media sets for tapes”](#) on page 386.

Default media sets

When you install Backup Exec, default system media sets and default user media sets are created automatically. When you add tapes to Backup Exec by importing media, Backup Exec associates the tape with one of the system media sets.

Note: You should not associate scratch media with a media set that you create. When a backup job runs, Backup Exec automatically moves the media from the Scratch Media set to the required media set as needed.

See [“Importing media to Backup Exec ”](#) on page 459.

You cannot modify the properties of system media sets. System media sets are described in the following table:

Table 10-2 Default system media sets

Name	Description
Backup Exec and Windows NT Backup Media	<p>Displays all media that is imported from another installation of Backup Exec.</p> <p>See “Cataloging a storage device” on page 452.</p>
Cleaning Media	<p>Displays all cleaning media.</p> <p>See “Cleaning a robotic library drive” on page 458.</p>
Foreign Media	<p>Displays all media that is imported from a product other than Backup Exec.</p> <p>See “About restoring NetWare SMS volume backups to non-SMS volumes with Backup Exec” on page 246.</p> <p>See “Cataloging a storage device” on page 452.</p>
Retired Media	<p>Displays all media that you have taken out of service, usually because of an excessive number of errors. After you associate a media with the retired media set, Backup Exec does not select it for backup jobs. The media is still available for restore operations, if it has not been damaged. Retired Media protects media from being used (overwritten).</p> <p>If Backup Exec cannot recognize data on a tape, then it moves the tape to Retired Media. If you want to reuse the tape, erase or label the tape. These operations write a new header to the tape that Backup Exec can recognize. After the tape is erased or labeled, Backup Exec moves it to the Scratch Media set.</p> <p>You can delete the media that is in Retired Media to remove it from Backup Exec. You may want to delete media if you have a lot of off-site media that you do not want to recycle. You can also delete media if you throw away the media.</p> <p>See “Erasing tape or disk cartridge media” on page 407.</p> <p>See “Retiring damaged tape media” on page 406.</p>

Table 10-2 Default system media sets (*continued*)

Name	Description
Scratch Media	<p>Displays all media that can be overwritten. New, blank, and erased media are automatically associated with the Scratch Media set.</p> <p>See “Overwrite protection periods and append periods in media sets” on page 382.</p>

Backup Exec creates the following default user media sets:

Table 10-3 Default user media sets

Name	Description
Keep Data for 4 Weeks	<p>Displays all tape media that you associate with this media set. If you use the backup job defaults that are set when you install Backup Exec, the media set Keep Data for 4 Weeks is the default media set for all backup jobs that you send to tape storage. This media set protects data from being overwritten for four weeks and allows the media to be appended to for six days.</p> <p>You can edit and rename Keep Data for 4 Weeks after installation. Therefore, it may not continue to appear in the Media view or in the backup job defaults as Keep Data for 4 Weeks.</p>
Keep Data Infinitely - Do Not Allow Overwrite	<p>Displays all tape media that you associate with this media set.</p> <p>When you associate media with this media set, data is not overwritten unless you perform any of the following actions on the media:</p> <ul style="list-style-type: none">■ Erase■ Label■ Format■ Associate the media with the scratch media set <p>You can append data to this media for an infinite period (until the media is full).</p> <p>You can edit and rename Keep Data Infinitely - Do Not Allow Overwrite after installation. Therefore, it may not continue to appear in the Media view or in the backup job defaults as Keep Data Infinitely - Do Not Allow Overwrite.</p>

You can change the default media set for backup jobs by doing one of the following:

- Create new media sets that have the append and overwrite protection periods set to the time intervals that accommodate your data retention strategy. Then, specify the media set that is most appropriate when you create a backup job. For example, you can create a media set that keeps data for 60 days, and a media set that keeps data for 90 days.
- Select the other default media set **Keep Data Infinitely - Do Not Allow Overwrite** when you create a backup job. The risk that is associated with the media set **Keep Data Infinitely - Do Not Allow Overwrite** is that you can use all of your scratch media. You must continually add new tape or disk cartridge media to Backup Exec.

Note: Symantec recommends that if you need to keep data longer than four weeks, you should duplicate it. You can duplicate the backup data from the original storage device to tape, which you can then send for long-term or off-site storage.

To view all media sets

- ◆ On the **Storage** tab, double-click **All Media Sets**.

See [“Overwrite protection periods and append periods in media sets”](#) on page 382.

See [“Creating media sets for tapes”](#) on page 386.

See [“Creating media vault rules to move tape media to and from media vaults”](#) on page 404.

See [“Duplicating backup sets or a job history manually”](#) on page 221.

See [“Associating tape media with a media set”](#) on page 409.

Overwrite protection periods and append periods in media sets

Each tape media is associated with a media set, which is a set of rules that manage media.

These rules include overwrite protection and append periods.

Table 10-4 Overwrite protection and append periods

Rule	Description
Append period	The amount of time that data can be appended to tape media. It is measured from the time the media was first allocated. It can be specified in hours, days, weeks, or years.

Table 10-4 Overwrite protection and append periods (*continued*)

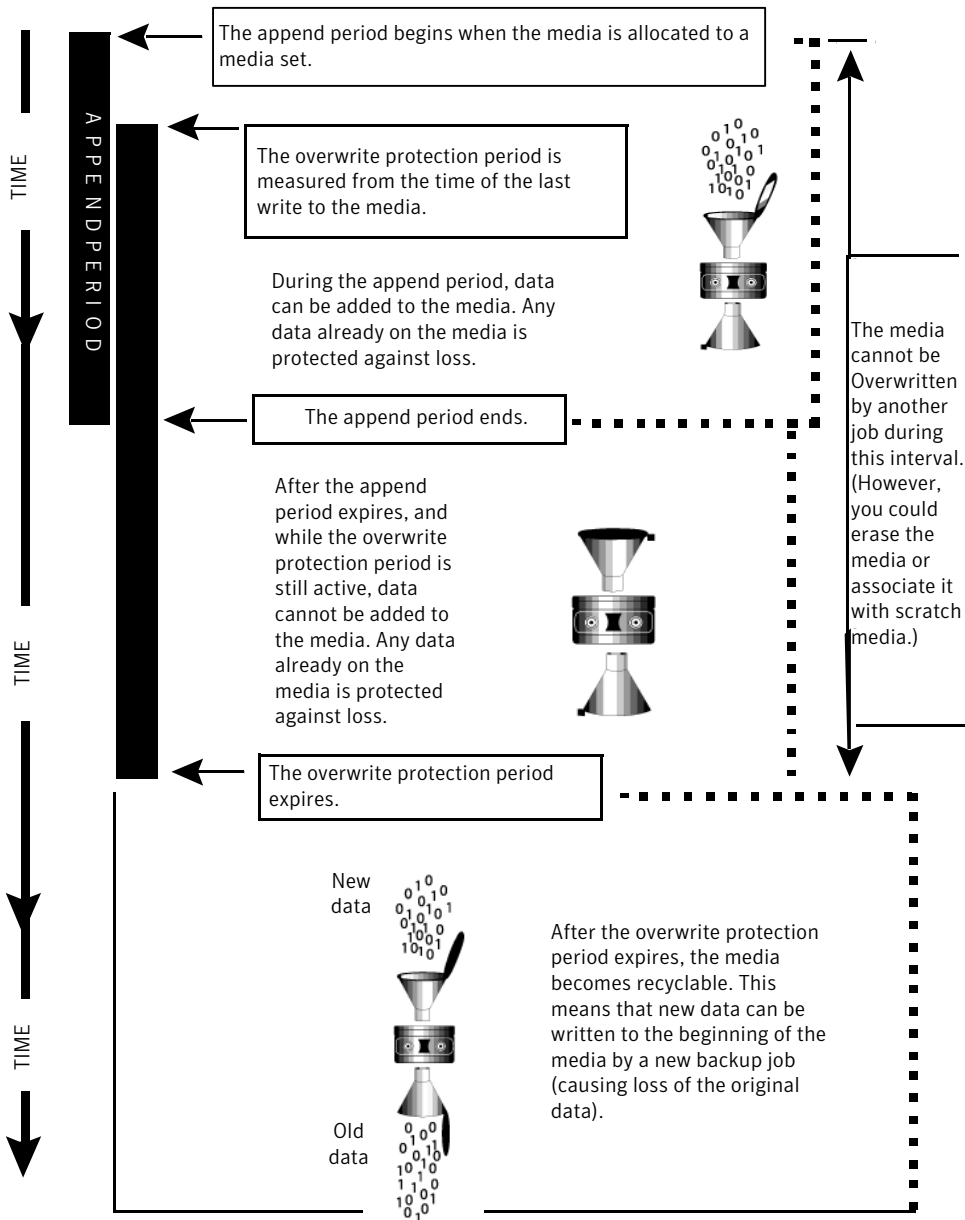
Rule	Description
Overwrite protection period	<p>The amount of time that tape media is protected from being overwritten. The period is measured from the time of the last write to the media, that is, at the end of the last append or overwrite job. It can be specified in hours, days, weeks, or years. When the overwrite protection period is over, the media becomes recyclable and can be overwritten.</p> <p>The overwrite protection period begins when the backup job is completed. If there is an append period, the overwrite protection period begins again each time an append job completes. Because the overwrite protection period does not begin until the job completes, the amount of time that is required to complete the job affects when the media can be overwritten. You may shorten the overwrite protection period to take into account the amount of time a job may run.</p> <p>For example, you set the overwrite protection period for seven days. You also set the append period for four days to ensure that data is not overwritten for at least seven days. The data can be appended to the media for the next four days. The last data that is appended to this media is retained for at least seven days.</p> <p>Note: Any media can be overwritten if the overwrite protection level is set to None.</p> <p>See “Media overwrite protection levels for job media” on page 394.</p>

Your tape media rotation strategy must balance between the need to save data as long as possible, and the fact that tape media are not in infinite supply. The media set rules allow Backup Exec to identify which tape media can be written to and which tape media are overwrite-protected. You should consider the use of disk storage for backup data.

See [“Configuring disk storage”](#) on page 318.

The following graphic shows the relationship between the append period and the overwrite protection period.

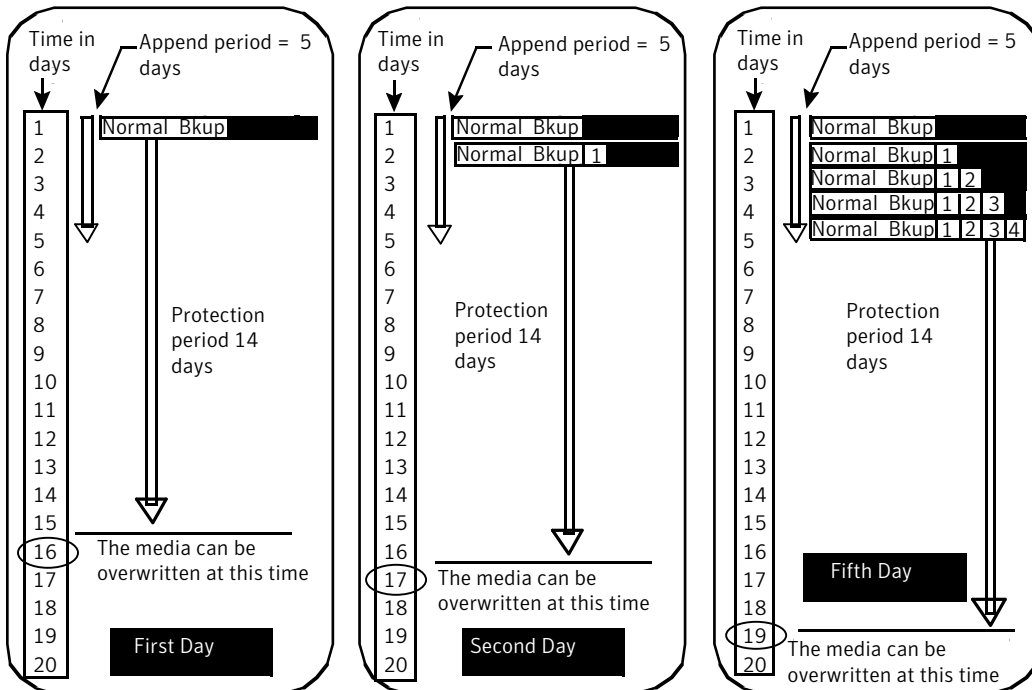
Figure 10-1 Append periods and overwrite protection periods



The append and overwrite protection periods that you specify apply to all the data on the media.

Each time data is written to a media, the time remaining in the overwrite protection period is reset, and the countdown is restarted.

Figure 10-2 How overwrite protection periods are reset



The amount of time that is required to complete the job affects when the media can be overwritten.

For example, suppose that you create a media set named Weekly with an overwrite protection period of seven days. You also specify an append period of zero days, and you schedule a full backup job to run each Friday at 20:00. When it is time for the full backup to run at 20:00 the following Friday, the job cannot run. The first backup job that ran the previous Friday did not complete until 21:10. The overwrite protection period for the Weekly media set still has 70 minutes remaining.

To prevent this situation, you can shorten the overwrite protection period to account for the amount of time a job may run. For this example, the scheduled job recurring at 20:00 can run if the overwrite protection period is set to six days instead of seven days.

Creating media sets for tapes

A media set consists of the rules that specify the following:

- Append periods
- Overwrite protection periods
- Media vaults
- Amount of time to move media to and from the media vault

The media set rules apply to all of the media that you associate with the media set.

Note: You must have already created a media vault before you are prompted to add media vault rules in a media set.

See [“Default media sets”](#) on page 379.

To create a media set for tapes

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media**.
- 2 Double-click **All Media Sets**.
- 3 Under **User Media Sets**, right-click a media set, and then click **Create media set**.
- 4 Enter a name and a description for the media set, and then click **Next**.
- 5 Specify how long you want to keep data in this media set from being overwritten, and then click **Next**.
- 6 Specify how long you want to allow data to be appended to media in this media set, and then click **Next**.
- 7 Review the summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To create the media set

Click **Finish**.

See [“Media set properties”](#) on page 389.

Editing media set properties

You can edit the following properties for media sets:

- Name of a media set.
- Overwrite protection and append periods for a media set.
- Media vault and the vaulting periods associated with a media set.

See [“Default media sets”](#) on page 379.

To edit media set properties

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media** .
- 2 Double-click **All Media Sets**.
- 3 Double-click the media set for which you want to edit properties.
- 4 In the left pane, click **Properties**.
- 5 Change any of the following information:

Name	Displays the name of the media set.
Description	Displays the description of the media set.

Overwrite protection period

Displays the length of time in hours, days, weeks, or years to retain the data on the media before the media can be overwritten.

Regardless of the overwrite protection period that is set, media can be overwritten if you perform the following operations on it:

- Erase
- Label
- Associate it with the **Scratch Media Set**
- Set the **Media Overwrite Protection Level** to **None**
- Format

Because of the method Backup Exec uses to compute time, the unit of time that you enter may be converted. For example, if you enter 14 days, the next time you view this property, it may be displayed as two weeks.

The default period is **Infinite - Don't Allow Overwrite**, which protects the media from being overwritten for 1,000 years.

See [“Overwrite protection periods and append periods in media sets”](#) on page 382.

Append period

Displays the length of time in hours, days, or weeks that data can be added to media. Because of the method Backup Exec uses to compute time, the unit of time that you enter may be converted. For example, if you enter 14 days, the next time you view this property, it may be displayed as two weeks.

The append period starts when the first backup job is written to this media.

The default period is **Infinite - Allow Append**, which allows data to be appended until the media capacity is reached.

Media vault to use with this media set	Displays the media vault that stores the media that is associated with this media set. See "Default media vaults" on page 402.
Move media to this vault after	Displays the time period after which this media is reported as ready to be moved to this vault.
Return media from this vault after	Displays the time period after which this media is reported as ready to be returned from this vault.

6 Click **Apply**.

Media set properties

Properties for media sets provide information about the overwrite protection period, the append period, and media vaults.

See ["Editing media set properties"](#) on page 387.

Table 10-5 Properties for media sets

Item	Description
Name	Displays the name of the media set.
Description	Displays the description of the media set.

Table 10-5 Properties for media sets (*continued*)

Item	Description
Overwrite protection period	<p>Displays the length of time in hours, days, weeks, or years to retain the data on the media before the media can be overwritten.</p> <p>Regardless of the overwrite protection period that is set, media can be overwritten if you perform the following operations on it:</p> <ul style="list-style-type: none">■ Erase■ Label■ Associate it with the Scratch Media Set■ Set the Media Overwrite Protection Level to None■ Format <p>Because of the method Backup Exec uses to compute time, the unit of time that you enter may be converted. For example, if you enter 14 days, the next time you view this property, it may be displayed as two weeks.</p> <p>The default period is Infinite - Don't Allow Overwrite, which protects the media from being overwritten for 1,000 years.</p> <p>See "Overwrite protection periods and append periods in media sets" on page 382.</p>
Append period	<p>Displays the length of time in hours, days, or weeks, that data can be added to media. Because of the method Backup Exec uses to compute time, the unit of time that you enter may be converted. For example, if you enter 14 days, the next time you view this property, it may be displayed as two weeks.</p> <p>The append period starts when the first backup job is written to this media.</p> <p>The default period is Infinite - Allow Append, which allows data to be appended until the media capacity is reached.</p>
Media vault to use with this media set	<p>Displays the media vault that stores the media that is associated with this media set.</p> <p>See "Default media vaults" on page 402.</p>
Move media to this vault after	<p>Displays the time period after which this media is reported as ready to be moved to this vault.</p>

Table 10-5 Properties for media sets (continued)

Item	Description
Return media from this vault after	Displays the time period after which this media is reported as ready to be returned from this vault.

Changing the overwrite protection period or the append period for a media set

You can change the length of time to retain data on media that are associated with a media set. You can also change the length of time to add data to media that are associated with a media set.

See [“Overwrite protection periods and append periods in media sets”](#) on page 382.

To change the overwrite protection period or the append period for a media set

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media**.
- 2 Double-click **All Media Sets**.
- 3 Double-click the media set for which you want to change the overwrite protection period or the append period.
- 4 In the left pane, click **Properties**.
- 5 Change the following information as appropriate:

Overwrite protection period

Displays the length of time in hours, days, weeks, or years to retain the data on the media before the media can be overwritten.

Regardless of the overwrite protection period that is set, media can be overwritten if you perform the following operations on it:

- Erase
- Label
- Associate it with the **Scratch Media Set**
- Set the **Media Overwrite Protection Level** to **None**
- Format

Because of the method Backup Exec uses to compute time, the unit of time that you enter may be converted. For example, if you enter 14 days, the next time you view this property, it may be displayed as two weeks.

The default period is **Infinite - Don't Allow Overwrite**, which protects the media from being overwritten for 1,000 years.

Append period

Displays the length of time in hours, days, or weeks that data can be added to media.

Because of the method Backup Exec uses to compute time, the unit of time that you enter may be converted. For example, if you enter 14 days, the next time you view this property, it may be displayed as two weeks.

The append period starts when the first backup job is written to this media.

The default period is **Infinite - Allow Append**, which allows data to be appended until the media capacity is reached.

Changing the name and description of a media set

You can change the name and description of a media set.

To change the name and description of a media set

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media** .
- 2 Double-click **All Media Sets**.
- 3 Double-click the media set for which you want to change the name or description.
- 4 In the left pane, click **Properties**.
- 5 Change the name and/or the description of the media set, and then click **Apply**.

Changing the media vault or vaulting periods for a media set

You can change the media vault that stores this media set, and you can change the time periods for when you want to move media to a vault or return media from a vault.

See [“Default media vaults”](#) on page 402.

To change the media vault or vaulting periods for a media set

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media** .
- 2 Double-click **All Media Sets**.
- 3 Double-click the media set for which you want to change the media vault or vaulting periods.
- 4 In the left pane, click **Properties**.
- 5 Change the folloiwng information as appropriate:

Media vault to use with this media set	Displays the media vault that stores the media that is associated with this media set.
Move media to this vault after	Displays the time period after which this media is reported as ready to be moved to this vault.
Return media from this vault after	Displays the time period after which this media is reported as ready to be returned from this vault.

- 6 Click **Apply**.

Deleting a media set

If you delete a media set that has scheduled jobs associated with it, you are prompted to associate the jobs to another media set.

Caution: Ensure that the media set that you associate the jobs with has the appropriate overwrite protection and append periods.

See [“Default media sets”](#) on page 379.

To delete a media set

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media**.
- 2 Double-click **All Media Sets**.
- 3 Right-click the media set that you want to delete, and then click **Delete**.
- 4 When you are prompted to delete the media set, click **OK**.

Renaming a media set

When you rename a media set, any jobs that use that media set display the new media set name.

See [“Default media sets”](#) on page 379.

To rename a media set

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media**.
- 2 Double-click **All Media Sets**.
- 3 Right-click the media set that you want to rename, and then click **Details**.
- 4 In the left pane, click **Properties**.
- 5 In the **Name** field, type the new name that you want to assign to this media set, and then click **Apply**.

Media overwrite protection levels for tape media

The media overwrite protection level is a global setting that supersedes the media set's overwrite protection period. Although the terms are similar, the media overwrite protection level and the media overwrite protection period are different. The media overwrite protection period is a time interval that changes from one media set to another. The media overwrite protection level specifies whether to overwrite scratch, imported, or allocated media, regardless of the media's overwrite protection period.

Use the media overwrite protection level to specify the type of media that you want to be available for overwrite backup jobs.

See [“Overwrite protection periods and append periods in media sets”](#) on page 382.

Overwriting allocated or imported tape media

Media that are associated with a media set are called allocated media. Media that are imported from another installation of Backup Exec, or from another product are called imported media. Backup Exec protects allocated and imported media from being overwritten when full or partial overwrite protection is used. However, you can let Backup Exec overwrite allocated and imported media before the data overwrite protection period expires, and without setting the media overwrite protection level to None.

The following methods are available:

- Associate the media with the **Scratch Media Set**. The media is overwritten when it is selected for an overwrite job.
- Erase the media. Erased media is automatically recognized as scratch media and is overwritten immediately.
- Label the media. The **Label Media** operation immediately writes a new media label on the media, which destroys any data that is contained on the media.
- Format the media. Formatting destroys any data that is contained on the media.
- Change the overwrite protection period for the media set so that it is expired.

See [“Managing tapes”](#) on page 378.

See [“Deleting tape media”](#) on page 407.

See [“Editing global settings for storage”](#) on page 439.

How Backup Exec searches for overwritable media in tape drives

Media overwrite options set the order in which Backup Exec searches for overwritable media in tape drives. When Backup Exec searches for overwritable media for a backup job, it searches for either scratch media or media that has an expired overwrite protection period.

You are prompted to select one of the following options that you want Backup Exec to use first:

- Overwrite scratch media before overwriting recyclable media that is contained in the destination media set.

If you choose to overwrite scratch media before recyclable media, more media may be required for the same number of jobs. However, the recyclable media may be preserved longer for possible recovery.

- Overwrite recyclable media that is contained in the destination media set before overwriting scratch media.

If you choose to overwrite recyclable media before scratch media, the same media is re-used more frequently than if you overwrite scratch media before recyclable media.

In a storage device pool for tape drives, Backup Exec selects the oldest recyclable media in the storage device pool to use first.

In a robotic library, Backup Exec selects the oldest recyclable media in the library to use first. If the robotic library is partitioned, Backup Exec searches for the oldest recyclable media in the targeted partition only.

Caution: Symantec recommends that you physically write-protect media containing critical data. Use the write-protect tab on the media cartridge to protect against unintentional move or erase operations, or expired overwrite protection periods.

The following table describes the order in which Backup Exec searches for media to use for an overwrite job.

Table 10-6 How Backup Exec searches for overwritable media in tape drives

Overwrite protection level and overwrite option:	Media is overwritten in tape drives in this order:
Full + Overwrite scratch media first Note: This combination provides the most protection against overwriting media.	<ol style="list-style-type: none">1 Scratch media2 Recyclable media in the destination media set3 Recyclable media in any media set
Full + Overwrite recyclable media first	<ol style="list-style-type: none">1 Recyclable media in the destination media set2 Scratch media3 Recyclable media in any media set
Partial + Overwrite scratch media first	<ol style="list-style-type: none">1 Scratch media2 Recyclable media in the destination media set3 Recyclable media in any media set4 Media that is imported from another installation of Backup Exec, or from another product

Table 10-6 How Backup Exec searches for overwritable media in tape drives
(continued)

Overwrite protection level and overwrite option:	Media is overwritten in tape drives in this order:
Partial + Overwrite recyclable media first	<ol style="list-style-type: none"> 1 Recyclable media in the destination media set 2 Scratch media 3 Recyclable media in any media set 4 Media that is imported from another installation of Backup Exec, or from another product
None - No overwrite protection + overwrite scratch media first Warning: This options is not recommended because it does not protect data from being overwritten.	<ol style="list-style-type: none"> 1 Scratch media 2 Recyclable media in the destination media set 3 Recyclable media in any media set 4 Media that is imported from another installation of Backup Exec, or from another product 5 Allocated media in any media set
None - No overwrite protection + overwrite recyclable media first Warning: This options is not recommended because it does not protect data from being overwritten.	<ol style="list-style-type: none"> 1 Recyclable media in the destination media set 2 Scratch media 3 Recyclable media in any media set 4 Media that is imported from another installation of Backup Exec, or from another product 5 Allocated media in any media set

In addition to setting overwrite protection levels, you must set overwrite options, which set the order in which Backup Exec searches for overwritable media.

The most obvious candidates for backup jobs that require overwritable media are scratch media and recyclable media. Recyclable media have expired overwrite protection periods. Backup Exec searches for these types of media first when a backup requires tape media to overwrite. The search pattern is different according to whether you have selected Full, Partial, or None. The media indicate that a type of media set is examined for availability.

See [“Editing global settings for storage”](#) on page 439.

See [“Default media vaults”](#) on page 402.

Viewing tapes that are used by a job

To see which tapes are used for a job, review the Device and Media Information section of the job log.

See [“Managing tapes”](#) on page 378.

To view the tapes that are used by a job

- 1 Do any of the following:

To view the job log from the **Job Monitor** tab Select the **Job Monitor** tab.

To view the job log from the **Backup and Restore** tab or the **Storage** tab Do the following in the order listed:

- On the **Backup and Restore** tab or the **Storage** tab, double-click the server or the storage device where the job ran.
- In the left pane, click **Job History**.

- 2 Right-click the job, and then select **View Job Log**.
- 3 Review the information in the section **Device and Media Information**.

Labeling tape media

Media labels identify the tapes that you use in Backup Exec. When a new, blank, or unlabeled tape is used during a backup operation, Backup Exec automatically labels the tape media. This label consists of a prefix that identifies the cartridge type, and an incrementing number. For example, if the media is a 4mm tape, then the prefix is 4M, followed by 000001. The next media label generated for an unlabeled 4mm tape would be 4M000002, and so on.

Another type of media label used by Backup Exec is the media ID, which is a unique label assigned by Backup Exec to the individual tape media used in Backup Exec. The media ID is used internally by Backup Exec to keep statistics on each media. Because the media label or barcode label for tape media can be changed, Backup Exec must use the media ID to preserve continuity in record keeping for each individual tape media. You cannot change or erase the media ID. The media ID has no effect on the media label, or on your ability to rename, label, or erase tape media.

At times, you may need to use the media ID to distinguish the tape media that have duplicate media labels. Duplicate labels can be automatically generated in instances when Backup Exec is reinstalled or media from another Backup Exec installation is used. You can view the media ID in a media's property page.

Write the media label on an external label that is fixed to the outside of the physical tape cartridge. Whenever you change the media label, you should also change the external label to match.

The following methods are available in Backup Exec to change a tape media label:

- Write a new media label on the tape media. The Label operation destroys any data on the media.
- Rename the media. Renaming the tape media changes the name of the media in the display, but does not write the new label to the media until an overwrite operation occurs. The data on the media is viable until the media is overwritten.
- Edit the label. Editing the label changes the name of the tape media in the display, but does not write the new label to the media until an overwrite operation occurs. The data on the media is viable until the media is overwritten.

Note: Media that use barcode labels cannot be renamed. When you try to label the media that use barcode labels, the job logs report successfully completed jobs. However, the media label names do not change.

To label media

- 1 On the **Storage** tab, do either of the following:
 - Right-click the drive that contains the tape that you want to label.
 - Double-click **Slots**, and then right-click the slot that contains the tape that you want to label.
- 2 Click **Label**.

The following warning appears:

This operation is performed on the current media in the drive or slot. If the media has changed since the last inventory ran, the media label in the next dialog may not match the media in the selected device.
- 3 Click **OK**.
- 4 Type the name that you want to use as the media label for this media.
- 5 To erase all data on the media and re-label the media, click **OK**.
- 6 Write this same media label on an external label that is fixed to the outside of the physical media.
- 7 (Optional) View the job history for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 451.

See [“Renaming a tape media label ”](#) on page 401.

See [“How barcode labels become media labels for tapes”](#) on page 400.

About labeling imported tape media

Tape media that is imported from another installation of Backup Exec, or from another product is called imported media. Backup Exec does not automatically relabel imported media.

Backup Exec reads the imported tape media's existing label and displays the label in either the **Backup Exec and Windows NT Backup Media** media set or in the **Foreign Media** media set. If the media overwrite protection level is set to Partial or None, the imported media may be selected for a job and be overwritten. The imported media is automatically labeled when it is overwritten during a job. If you want to label a specific imported media while maintaining full media overwrite protection for other imported media, erase the specific media and then label it.

The original media label of the imported tape media is displayed in the media's properties. You can edit the media description in the media's property page to make it a more descriptive label.

See [“Erasing tape or disk cartridge media”](#) on page 407.

How barcode labels become media labels for tapes

If there is a barcode label on the physical tape cartridge, and if the robotic library has a barcode reader, the barcode label automatically becomes the media label.

For example, robotic library 1 has barcode support. During a backup operation, Backup Exec requests a new media or an overwritable media for the operation. A new media with the barcode label 'ABCD' is inserted in the robotic library magazine and the barcode reader scans the barcode ID. Backup Exec selects this media for the operation and detects that a barcode label has been assigned to the media. Backup Exec automatically uses the barcode label and continues the operation.

If a barcode label is the media label, then you cannot change the media label in Backup Exec. To change a barcode label and use a media label, you must remove the physical barcode label from the media cartridge. Or, you can use the media in a device that does not have a barcode reader.

When you change magazines or insert new media in a magazine, you can use the **Scan** operation to quickly update slot information.

See [“Labeling tape media”](#) on page 398.

See [“Scanning a storage device”](#) on page 452.

Renaming a tape media label

You can rename a tape media's label and description. The new label is not written to the tape media until an overwrite operation occurs. All of the data that is on the media is preserved until the next overwrite job. However, the new media label is stored in the database and is displayed for that media. To write a new media label to the media immediately, use the **Label** operation. The media's contents are erased.

If you rename a tape media, and then use it in another installation of Backup Exec, that media is imported to the **Backup Exec and Windows NT Media** media set. The media's original media label is displayed. The renamed label is not transferred to other installations of Backup Exec.

Note: If a barcode label is the media label, then you cannot change the media label in Backup Exec.

See [“Labeling tape media”](#) on page 398.

To rename a tape media label

- 1 On the **Storage** tab, right-click the tape drive or slot that contains the media that you want to relabel, and then click **Details**.
- 2 In the left pane, click **Media Properties**.
- 3 In the **Media label** field, enter a new label name.
- 4 To change the description, enter a new description in the **Media description** field.
- 5 Click **Apply**.

How WORM media is used in Backup Exec

Write once, read many (WORM) data storage is used to keep the data that has a long retention period. Data can be written to WORM media one time only. After it is written to, the media can be appended to, but it cannot be overwritten, erased, or reformatted.

When WORM media is used in a media set, the overwrite protection period is not applied to it, but the append period is applied.

New WORM media is WORM media that has not been written to. When new WORM media is introduced into Backup Exec, it is placed in the **Scratch Media** set. After the WORM media has been written to one time, you cannot move it to the Scratch media set. You can move WORM media to the **Retired Media** set to delete it from Backup Exec, but you cannot erase it or reformat it.

You can format a tape as a WORM tape if the tape drive supports the operation.

See [“Formatting a tape as a WORM tape”](#) on page 456.

When you select the option to use WORM media, Backup Exec verifies that the destination device is or contains a WORM-compatible drive. Backup Exec also verifies that the WORM media is available in the drive. If WORM media or a WORM-compatible drive is not found, an alert is sent and the backup job may fail.

See [“Configuring storage options for backup jobs”](#) on page 527.

Default media vaults

A media vault is a logical representation of the actual physical location of specified tape media. You can create media vaults to keep track of where media is physically stored, such as a scratch bin, or an off-site location. Backup Exec creates default media vaults to let you view all media that are online, offline, or in a media vault.

You must run the **Configure Storage** wizard to update the location of tape media in media vaults. From the **Configure Storage** wizard, you can print reports that detail which tape media are ready to move to and return from the vault. You can also update the location of the media if you choose to move them. However, you must physically collect the tape media, and move the media to and from the vault. The location of the tape media is updated in the Backup Exec Database, but the media is not ejected or exported. If Backup Exec detects that the media is in a robotic library, you are prompted to export media. If you choose to export the media, an export media job runs. If your environment includes remote sites, you should create separate media sets for each remote site. Then, the reports contain details on which media are ready to be moved for a specific site.

Table 10-7 Default media vaults

Default media vault	Description
Online Tape Media	Displays the media that are available in tape drives or robotic libraries. You cannot add or move media to the online media vault. Backup Exec does that automatically. If you move media from the online media vault to another media vault, the media's overwrite protection period and append period remain in effect.

Table 10-7 Default media vaults (*continued*)

Default media vault	Description
Offline Tape Media	Displays the media that are on-site but are not in tape drives or robotic libraries, and are not in media vaults. Media appear in the offline media vault if you use Backup Exec to remove media from a tape drive or robotic library. You can add media to the offline media vault from another media vault. An inventory operation or a catalog operation moves the offline media back to the online media vault. You cannot delete or rename the offline media vault.
Vaulted Tape Media	Displays the media that are not in tape drives or robotic libraries, and have been moved to a media vault. Vaulted Tape Media displays in All Media Vaults details only after you create a media vault.
All Media Vaults	Displays the media that are in media vaults that you create. All Media Vaults displays on the Storage tab only after you create a media vault. You can associate media vaults with media sets that you create. You specify when to move the media from a media set to the media vault. You also specify when the media is to return to the media set from the media vault. See “Creating media vault rules to move tape media to and from media vaults” on page 404.

See [“Changing the name or description of a media vault”](#) on page 403.

Changing the name or description of a media vault

You can edit the name and description of a media vault.

See [“Default media vaults”](#) on page 402.

To change media vaults and media vault rules, edit the properties of the media set that is associated with the media vault.

See [“Editing media set properties”](#) on page 387.

Changing the name or description of a media vault

- 1 On the **Storage** tab, expand **Tape Cartridge Media Sets and Vaults**, and then expand **All Media Vaults**.
- 2 Right-click the media vault for which you want to edit properties, and then click **Details**.
- 3 In the left pane, click **Media Vault Properties**.
- 4 Change the name or description of the media vault.
- 5 Click **Apply**.

Media vault properties

Properties for media vaults include the name and a description of the media vault.

See [“Changing the name or description of a media vault”](#) on page 403.

Table 10-8 Properties for media vaults

Item	Description
Name	Displays the name of the media vault.
Description	Displays a description of the media vault.

Creating media vault rules to move tape media to and from media vaults

Create media vault rules to do the following:

- Associate a media vault with the media set to which you want to send media.
- Specify the amount of time to wait between when the media is allocated and when it is sent to the vault.
- Specify the amount of time to wait between returning the media from the vault and when it was last written to.

See [“Default media vaults”](#) on page 402.

Backup Exec does not update the vault automatically. You must use the **Configure Storage** wizard to update the location of the tape media. You can also print or view the reports that contain details on which media are ready to be moved to and from the vault.

See [“Updating the tape media location in media vaults”](#) on page 405.

To create media vault rules to move tape media to and from media vaults

- 1 On the **Storage** tab, expand **All Media Sets**.
- 2 Right-click **Keep Data for 4 Weeks, Keep Data Infinitely - Do Not Allow Overwrite**, or a media set that you created, and then click **Details**.
- 3 In the left pane, click **Properties**.
- 4 Select the media vault that you want to use with the media set.
- 5 Specify when to move the media to the vault and when to return the media to the media set.

See [“Media set properties”](#) on page 389.

Updating the tape media location in media vaults

You can update the location of tape media that are in vaults. You can also print the reports that detail which media are ready to move to and return from the vault. However, you must physically collect the media, and move the media to and from the vault.

See [“Default media vaults”](#) on page 402.

To update the tape media location in media vaults

- 1 On the **Storage** tab, expand **Tape Cartridge Media Sets and Vaults**, and then double-click **All Media Vaults**.
- 2 Right-click the media vault for which you want to update the media location, and then click **Update vault using wizard**.
- 3 Follow the on-screen prompts.

Deleting a media vault

You can only delete an empty media vault. If any tape media in the vault, you must move it before you can delete the vault. You cannot delete the online media vaults or the offline media vaults.

See [“Default media vaults”](#) on page 402.

To delete a media vault

- 1 On the **Storage** tab, expand **All Media Vaults**.
- 2 Right-click the media vault that you want to delete, and then click **Delete**.
- 3 Click **Yes** when you are prompted to delete the media vault.

Moving tape media to a vault

You can use a barcode scanner to enter the media labels of tape media that you want to move to a vault. You can also type a media label into the dialog box.

See [“Default media vaults”](#) on page 402.

To move tape media to a vault

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media**, and then double-click **All Media Vaults**.
- 2 Right-click the media vault that you want to move media to, and then click **Move media to vault**.
- 3 Follow the on-screen prompts.

Retiring damaged tape media

You can retire damaged tape media so that Backup Exec does not use it for backup jobs. You should associate tape media that meets or exceeds the discard thresholds that are determined by the media manufacturer with the **Retired Media** media set. Backup Exec tracks the soft errors that are generated by the storage device firmware. Media that exceed acceptable levels of these errors are reported as potential candidates to be discarded.

To decide which tape media to retire, run a Media Errors report for the total number of errors for media, or view the properties for a specific media.

Associate any media with an unacceptable level of errors to **Retired Media** so that you are protected against using defective media before critical backup operations begin. After you associate tape media with the **Retired Media** set, it is not used by Backup Exec for future backup jobs. The media is still available to be restored from if it is not damaged.

To retire damaged tape media

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media**.
- 2 Right-click **Online Tape Media**, and then click **Details**.
- 3 Right-click the media that you want to retire, and then click **Retire**.
- 4 Click **Yes** when you are prompted to retire the media.

See [“Deleting tape media”](#) on page 407.

See [“Media Errors report”](#) on page 670.

Deleting tape media

When you delete tape media from Backup Exec, all records of the media are removed from the Backup Exec database. These records include catalog information, media statistics, and other information that is associated with the tape media. You can only delete media when it belongs to the **Retired Media** set.

You may want to delete tape media when the following occurs:

- You have a lot of off-site media that you do not want to recycle.
- You throw away damaged or old media.

If you import deleted media back into Backup Exec, it is added to either the **Backup Exec and Windows NT Media** media set or the **Foreign Media** media set. Before you can restore from the media, you must catalog it.

To delete tape media

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media**.
- 2 Expand **All Media Sets**, right-click **Retired Media**, and then click **Details**.
- 3 Right-click the media that you want to delete, and then click **Delete**.
- 4 Click **Yes** when you are prompted to delete the media.

See [“Retiring damaged tape media”](#) on page 406.

Erasing tape or disk cartridge media

You can erase tape media or disk cartridge media immediately, or you can schedule the erase operation.

Warning: The erase operation is run on whatever media is in the drive or slot at the time when the operation runs. If the media has been changed since the last inventory operation was run, the media label that appears on the Backup Exec Administration Console may not match the media in the selected drive or slot. If media is moved unexpectedly, data loss can occur. Check any scheduled erase jobs carefully.

Table 10-9 Erase operations

Erase operation	Description
Erase media now	<p>Writes an indicator at the beginning of the media that makes the data that is contained on the media inaccessible. For most uses, the Erase media now operation is sufficient.</p> <p>This is the only erase operation available for disk cartridge media.</p>
Long erase media now	<p>Instructs the drive to physically erase the entire media. If you have sensitive information on the media and you want to dispose of it, use the Long erase media now operation. A long erase operation on a media will take several hours to complete depending on the drive and the media capacity.</p> <p>Some devices do not support a long erase operation.</p>
Schedule	<p>Let you schedule either an Erase or Long erase operation, and lets you choose notification options.</p>

The erase operation does not change the media label. To change a tape media label, either run a Label operation or rename the media before you run an Erase operation.

You cannot cancel an Erase operation after it has started. You can cancel an Erase operation that is scheduled or queued.

When you schedule an erase operation, you can configure the time and frequency that you want to run the job.

To erase tape or disk cartridge media now

- 1 On the **Storage** tab, right-click the drive or the robotic library slot that contains the media that you want to erase.
- 2 Click **Erase media now**, and then do one of the following:
 - To run an erase operation immediately Click **Erase media now**.
 - To run a long erase operation immediately Click **Long erase media now**.
- 3 Click **Yes** when you are prompted to erase the media.

To schedule an erase operation for tape or disk cartridge media

- 1 On the **Storage** tab, right-click the drive or the robotic library slot that contains the media that you want to erase.
- 2 Click **Erase media now**, and then do one of the following:

To schedule an erase operation	Click Schedule erase .
To schedule a long erase operation	Click Schedule long erase .
- 3 Click **Yes** when you are prompted to erase the media.
- 4 To send notification when the job completes, in the left pane, click **Notification** and select the appropriate options.
- 5 To schedule the job, in the left pane, click **Schedule** and select the appropriate options.
See [“Scheduling a storage operation job”](#) on page 432.
- 6 Click **OK**.

About cataloging tape or disk cartridge media that contains encrypted backup sets

When you catalog tape media or disk cartridge media that contains encrypted backup sets, Backup Exec attempts to find valid encryption keys for the sets in the Backup Exec database. If Backup Exec does not find a valid key, it issues an alert that instructs you to create one. After you create a valid key, you can respond to the alert to retry cataloging the encrypted set. Alternatively, you can skip the encrypted set and continue to catalog the rest of the media, or cancel the catalog job.

See [“Encryption key management”](#) on page 586.

Associating tape media with a media set

When you create a backup job to tape media, the default media set that Backup Exec selects for you is called **Keep Data for 4 Weeks**. You can select other media sets when you create the backup job, or you can associate the tape media with another media set later.

When you associate tape media with a media set, the tape uses the following properties of that media set:

- Append periods
- Overwrite protection periods
- Media vaults
- Amount of time to move media to and from the media vault.

Note: You should not associate scratch or imported media with a media set. Backup Exec automatically associates scratch or imported media with a media set when a backup job requires it.

See [“Creating media sets for tapes”](#) on page 386.

To associate tape media with a media set

- 1 On the **Storage** tab, expand **Tape and Disk Cartridge Media**.
- 2 Double-click **All Tape Media** to display a list of media.
- 3 Right-click the tape media that you want to associate with a media set, and then click **Associate with media set**.
- 4 Select a media set from the drop-down list, and then click **OK**.

See [“Default media sets”](#) on page 379.

Editing tape media properties

You can view tape media properties and edit some properties.

See [“Managing tapes”](#) on page 378.

To edit tape media properties

- 1 On the **Storage** tab, double-click the drive that contains the media.
- 2 In the left pane, click **Media Properties**.

3 Change any of the following options:

Media label

Displays the media label that Backup Exec assigns automatically, or that the administrator assigned, or that is a pre-assigned barcode label.

You can edit the media label, which is limited to 32 characters. Editing the label changes the name of the media in the display, but does not write the new label to the media until an overwrite operation occurs. When you edit a media label, try to make it a concise identifier that remains constant even when the media is reused. You should write this media label on a label that is fixed to the outside of the physical media.

Duplicate labels can be automatically generated. For example, reinstalling Backup Exec or bringing media from another Backup Exec installation can cause duplication in labels. Duplicate labels are allowed, but not recommended.

If a barcode is available, and a barcode-equipped device is used, then the media label automatically defaults to that barcode.

Media description

Displays the original media label if the media is imported media.

You can edit the media description, which is limited to 128 characters, to make it a more descriptive label.

Preserve description

Keeps the media description when you select **Yes** in the drop-down box. The media description is kept until an overwrite backup job runs or an erase or label storage operation job runs.

By default, the media description is not retained. This option is set to **No** by default.

4 Click **Apply**.

Tape media rotation strategies

Many tape media rotation strategies exist that you can use to back up your data.

The most commonly used tape media rotation strategies include the following:

- Son, which uses the same tape each day to run a full backup.
- Father/Son, which uses multiple tapes, and includes a combination of weekly full and daily differential or incremental backups for a two-week schedule. This strategy provides backups for off-site storage .
- Grandfather, which uses multiple tapes, and includes a combination of weekly and monthly full and daily differential or incremental backups. This strategy also provides backups for off-site storage.

Son media rotation strategy


The Son media rotation strategy requires the following:

Table 10-10 Son media rotation strategy

Item	Description
Number of media required	1 (minimum)
Overwrite protection period	Last backup

The Son strategy involves performing a full backup every day.

Figure 10-3 Son backup strategy



Although the Son strategy is easy to administer, backing up with a single media is not an effective method of backup. Magnetic media eventually wears out after many uses and the data you can restore only spans back to your last backup.

Father/Son media rotation strategy

The Father/Son media rotation strategy requires the following:

Table 10-11 Father/Son media rotation strategy

Item	Description
Number of media required	6 (minimum)
Overwrite protection period	Two weeks


The Father/Son media rotation strategy uses a combination of full and differential or incremental backups for a two-week schedule.

In the Father/Son scenario, four media are used Monday through Thursday for differential or incremental backups. The other two media containing full backups are rotated out and stored off-site every Friday.

The Father/Son strategy is easy to administer and lets you keep data longer than the Son strategy. The Father/Son strategy is not suitable for the stringent data protection needs of most network environments.

When this backup strategy is first implemented, you must start with a full backup.

Figure 10-4 Father/Son backup strategy



Grandfather media rotation strategy

The Grandfather media rotation strategy requires the following:

Table 10-12 Grandfather media rotation strategy


Item	Description
Number of media required	19 (minimum)
Overwrite protection period	One year

The Grandfather method is one of the most common media rotation strategies. The Grandfather method is easy to administer and comprehensive enough to allow easy location of files when they need to be restored.

In the Grandfather scenario, four tapes are used Monday through Thursday for incremental or differential backups; another three tapes are used every Friday for full backups.

The remaining 12 media are used for monthly full backups and are kept off-site.

Figure 10-5 Grandfather backup strategy



The Grandfather strategy is recommended because it offers a good media number to storage life ratio (19 media/1 year). You can easily incorporate more media. For example, you can perform a full backup on the last Saturday of the month to keep permanently.

Storage device pools

This chapter includes the following topics:

- [Creating storage device pools](#)
- [Specifying a default selection method for a device in a storage device pool](#)
- [Viewing jobs, job histories, and active alerts for a storage device pool](#)
- [Adding or removing devices in a storage device pool](#)
- [Storage device pool properties](#)

Creating storage device pools

A storage device pool is a group of similar types of storage devices that enables load-balancing of Backup Exec jobs. The workload is shared across the storage device pool. You can send backup jobs to specific storage devices or to a storage device pool. By default, if the specific storage device is busy, the job must wait until the storage device becomes available.

When you send a job to specific storage devices, Backup Exec cannot automatically route the job to the next available storage device. When you submit a backup job to a storage device pool, by default the job is sent to the first available storage device in that pool. As other jobs are created and started, they can run concurrently on other storage devices in the storage device pool. By dynamically allocating storage devices as jobs are submitted, Backup Exec processes jobs quickly and efficiently. Storage device pools provide fault tolerance if you configure error-handling rules to resubmit a job that fails because of a storage device error.

You can change the default selection method for a device in a storage device pool so that jobs are sent first to one of the following devices in the storage device pool:

- A storage device with the least amount of free space. This setting is beneficial for manual offsite rotation strategies since the job runs to the same device until the device is full.
- A storage device with the most amount of free space. This setting is beneficial for a backup that has Granular Recovery Technology (GRT) enabled, and is sent to a disk storage device pool. GRT jobs that are sent to disk storage devices cannot span.

Note: You must use the Backup Exec Management Command Line Interface to change the default to specify whether jobs are sent to a device that has the most or least free space.

See [“Specifying a default selection method for a device in a storage device pool”](#) on page 419.

Backup Exec creates and maintains system-defined storage device pools for disk storage, tape drives, disk cartridges, and virtual disks. Any storage devices that you configure or that you attach are automatically added to the appropriate system-defined storage device pool. You can select these storage device pools when you select the storage for a backup job. You cannot edit a system-defined storage device pool. Before you can view the default storage device pools on the **Storage** tab, under **All Storage Pools**, you must create a storage device pool. If you have the Central Admin Server Option installed, you can create managed Backup Exec server pools.

Table 11-1 System-defined storage device pools maintained by Backup Exec

System-defined storage device pools that are maintained by Backup Exec	Description
Any disk storage	Contains the fixed-disk storage.
Any tape drive	Contains the tape drives. Backup Exec creates this pool when it detects an attached tape drive or robotic library. In a tape drive storage pool, Backup Exec uses the oldest recyclable media first.
Any disk cartridge	Contains the disk cartridges that you have configured to use as storage. Backup Exec creates this pool the first time that you configure disk cartridge storage.

Table 11-1 System-defined storage device pools maintained by Backup Exec
(continued)

System-defined storage device pools that are maintained by Backup Exec	Description
Any virtual disk	Contains the virtual disks that are on storage arrays. Backup Exec creates this pool when you install the Storage Provisioning Option.

For storage device pools that you create, you must edit the properties of the pool and check the check box next to the device before jobs can use that device in the pool.

Use the **Configure Storage** wizard to create storage device pools.

To create a storage device pool

- On the **Storage** tab, in the **Configure** group, click **Configure Storage**.
- Do one of the following:

If the Central Admin Server Option is not installed

Select **Storage pools**, and then click **Next**.

If the Central Admin Server Option is installed

Do the following in the order listed:
 - Select the Backup Exec server that you want to configure storage for, and then click **Next**.
 - Select **Storage pools**, and then click **Next**.
- Select **Storage device pool**, and then click **Next**.
- Enter a name and description for the pool, and then click **Next**.
- Click the drop-down arrow, select the type of storage device pool that you want to configure, and then click **Next**.
- Select all of the storage devices that you want to include in the pool, and then click **Next**.
- Review the storage configuration summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To configure the storage device pool

Click **Finish**.

See [“Adding or removing devices in a storage device pool”](#) on page 420.

See [“How to use Backup Exec server pools in CASO”](#) on page 1136.

See [“Adding or removing devices in a storage device pool”](#) on page 420.

See [“Error-handling rules for failed or canceled jobs”](#) on page 272.

Specifying a default selection method for a device in a storage device pool

The default selection method for a storage device pool is the first available storage device in a pool. You can use the Backup Exec Management Command Line Interface to specify another default selection method for a disk-based storage device pool. You can change the default setting of an existing storage device pool, or a storage device pool that Backup Exec creates.

You can change the default so that jobs are first sent to one of the following devices:

- A storage device with the least amount of free space. This setting is beneficial for manual off-site rotation strategies since the job runs to the same device until the device is full.
- A storage device with the most amount of free space. This setting is beneficial for a backup job for which Granular Recovery Technology (GRT) is enabled, and that is sent to a disk storage device pool. A GRT job cannot span when it is sent to a disk storage device.

The Backup Exec Management Command Line Interface commands are as follows:

- `New-BEStorageDevicePool` sets the device selection method for a new disk-based storage device pool when you create it.
- `Set-BEStorageDevicePool` sets the device selection method for an existing disk-based storage device pool.

For information on how to use the Backup Exec Management Command Line Interface and the commands, view the help file named BEMCLI, located in the default installation location:

C:\Program Files\Symantec\Backup Exec

See [“Creating storage device pools”](#) on page 416.

Viewing jobs, job histories, and active alerts for a storage device pool

You can view the jobs that are sent to a storage device pool, and the job histories as well as any active alerts. You must create a storage pool to enable **All Storage Pools** to appear on the **Storage** tab.

See [“How to monitor and manage jobs in Backup Exec”](#) on page 253.

See [“Alerts and notifications in Backup Exec”](#) on page 286.

To view jobs, job histories, and active alerts for a storage device pool

- 1 On the **Storage** tab, expand **All Storage Pools**.
- 2 Right-click the storage device pool for which you want to view the jobs, and then click **Details**.
- 3 In the left pane, click **Jobs**, **Job Histories**, or **Active Alerts**.

Adding or removing devices in a storage device pool

You can add or remove devices in a storage device pool. Only similar types of storage devices can belong to the same storage device pool. You cannot edit a system-defined storage device pool, such as **Any disk storage**, or **Any disk cartridge storage**.

You must create a storage device pool before **All Storage Pools** appears on the **Storage** tab.

See [“Creating storage device pools”](#) on page 416.

Note: To view or change the default device selection method for a storage device pool, you must use the Backup Exec Management Command Line Interface.

See [“Specifying a default selection method for a device in a storage device pool”](#) on page 419.

To add or remove devices in a storage device pool

- 1 On the **Storage** tab, expand **All Storage Pools**.
- 2 Double-click the storage pool for which you want to add and remove devices.
- 3 In the left pane, click **Properties**.
- 4 Do one or both of the following:

To add a storage device to the pool

Do the following in the order listed:

- Click **Add**.
- Select the devices that you want to add, and then click **OK**.

To remove a storage device from a pool

Select the device that you want to remove, and then click **Remove**.

- 5 Click **Apply**.

Storage device pool properties

You can edit the properties of any storage device pool that you create. You can change the name and description of a pool, and add and remove specific storage devices.

See [“Adding or removing devices in a storage device pool”](#) on page 420.

Note: To view or change the default device selection method for a storage device pool, you must use the Backup Exec Management Command Line Interface.

See [“Specifying a default selection method for a device in a storage device pool”](#) on page 419.

Table 11-2 Storage device pool properties

Item	Description
Name	Displays the name of the storage device pool.
Description	Displays a description of the storage device pool.
Storage device pool type	Displays the type of storage that is in this pool. You cannot edit this field.

Table 11-2 Storage device pool properties (*continued*)

Item	Description
Storage devices that belong to the pool	<p>Displays the storage devices that are currently in this pool. You can add or remove storage devices from this pool.</p> <p>When you configure a storage device for the first time, Backup Exec automatically adds it to the appropriate storage device pool. If Backup Exec detects an attached storage device, the storage device is automatically added to the appropriate storage device pool.</p>

Storage operations

This chapter includes the following topics:

- [About storage operation jobs](#)
- [Storage tab overview in Backup Exec](#)
- [Using the Configure Storage wizard](#)
- [Viewing details for multiple storage devices](#)
- [Sending a notification when a scheduled storage operation job completes](#)
- [Scheduling a storage operation job](#)
- [Editing global settings for storage](#)
- [Sharing storage devices](#)
- [Deleting a storage device](#)
- [Changing the state of a storage device to online](#)
- [Renaming a storage device](#)
- [Viewing jobs, job histories, backup sets, and active alerts for storage devices](#)
- [Cataloging a storage device](#)
- [Scanning a storage device](#)
- [Inventorying a storage device](#)
- [Inventorying and cataloging a storage device](#)
- [Pausing and unpausing a storage device](#)
- [Disabling and enabling a storage device](#)

- [Initializing a robotic library](#)
- [Formatting a tape as a WORM tape](#)
- [Retensioning a tape](#)
- [Formatting a tape in a tape drive](#)
- [Ejecting media from a disk cartridge or tape drive](#)
- [Cleaning a robotic library drive](#)
- [Importing media to Backup Exec](#)
- [Exporting media and expired media](#)
- [Locking and unlocking the robotic library's front portal](#)
- [Backup Exec server and storage device states](#)

About storage operation jobs

Backup Exec provides the storage operations that help you manage storage devices and media. You can perform most storage operations by right-clicking the storage device, and then selecting the operation. Only the storage operations that are supported for that storage device or media are available on the right-click menu. Not all storage operations are available for all devices.

Backup Exec treats virtual tape libraries and simulated tape libraries as physical robotic libraries. You can identify virtual tape libraries by the label VTL that displays on a library's properties pages. You can identify simulated tape libraries by the label TLS (Tape Library Simulator Utility). The virtual tape libraries and simulated tape libraries do not support all of the storage operations that are available for physical robotic libraries.

You can schedule some storage operations as recurring jobs. You can specify a schedule and a recipient for notification when these jobs run.

You can view all active and scheduled storage operations, and histories of storage operations on the **Job Monitor** tab.

See [“Scheduling a storage operation job”](#) on page 432.

See [“Sending a notification when a scheduled storage operation job completes”](#) on page 432.

Storage tab overview in Backup Exec

On the **Storage** tab, Backup Exec provides overview information for each storage device that you configure. You can view details for each storage device. If you want to view details for all of the storage devices on a server, you can select multiple storage devices.

See [“Viewing details for multiple storage devices”](#) on page 431.

You can customize the columns that appear in the **All Storage** view. Different columns are available from the **Tree** view and from the **List** view. Depending on the view that you select and the storage devices that are available, only some of the columns described in the following table may appear.

See [“How to sort, filter, and copy information on the Backup Exec Administration Console”](#) on page 136.

Table 12-1 All Storage overview

Item	Description
Name	<p>Indicates the name of the storage device.</p> <p>By default, Backup Exec provides a name for the storage device based on the type of storage and an incrementing number, such as Disk storage 0001. You can change the name of the storage device in the storage properties.</p> <p>See “Renaming a storage device” on page 451.</p>
State	<p>Indicates the state of the storage device, such as if it is online, offline, disabled, or if services need to be restarted.</p> <p>See “Backup Exec server and storage device states” on page 475.</p>
Parent Device	<p>Identifies the parent device if there is a relationship between devices. Examples include a relationship between a tape library and a tape drive, or between a storage array and a virtual disk.</p> <p>This column only appears if you select the List view from the Views group at the top of the window.</p>
Backup Exec Server	<p>Identifies the Backup Exec server on which you configured the storage device. If you share the device between multiple Backup Exec servers, the device displays for each Backup Exec server.</p> <p>This column only appears if you select the List view from the Views group at the top of the window.</p>

Table 12-1 All Storage overview (continued)

Item	Description
Storage Type	<p>Displays the type of storage that is associated with the device. The storage type can include tape drives, robotic libraries, disk storage, media sets, media vaults, and so on.</p> <p>This column only appears if you select the List view from the Views group at the top of the window.</p>
Active Alerts	<p>Indicates that an event or condition in Backup Exec has occurred for which a message is displayed or a response is required.</p> <p>See “Alerts and notifications in Backup Exec” on page 286.</p>
Storage Trending	<p>Indicates the estimate for the number of days of storage that is left for disk storage and virtual disk storage.</p> <p>See “Storage trending statuses for disk storage and virtual disks” on page 316.</p>

Table 12-1 All Storage overview (*continued*)

Item	Description
Capacity	<p>Displays storage capacity. Backup Exec provides overview information of used and free storage capacity, as well as capacity details for each storage that you configure. Storage capacity information is rolled up for any items that are collapsed under a storage type, such as a robotic library. The information that displays in the Capacity column includes all of the storage capacity of all of the collapsed items. When you expand the items, individual storage capacity information displays.</p> <p>Before capacity information can display for storage, you must inventory and catalog the storage.</p> <p>You can view storage capacity in the following places:</p> <ul style="list-style-type: none">■ On the Storage tab, in the Capacity column. When you hover the mouse over the capacity bar, additional details display in the tool tip.■ On the Backup and Restore tab, when you specify the storage for a backup job.■ On the Home tab, in Storage Status.■ On the Storage tab, when you view properties for disk storage devices. <p>For disk storage and disk cartridge storage, you can set the thresholds for low disk space on the device property pages.</p> <p>See “Editing disk storage properties” on page 321.</p> <p>See “Editing disk cartridge properties” on page 329.</p> <p>See “How Backup Exec catalogs work” on page 248.</p> <p>See “Inventorying a storage device” on page 453.</p>
Total Capacity	<p>Displays the total amount of storage space that is available on the device.</p> <p>For disk storage, this column indicates the size of the volume on which the disk storage is located.</p> <p>For disk cartridge storage, this column indicates the size of the cartridge in the disk cartridge.</p> <p>For tapes, this column indicates the used native capacity and total native capacity of the media.</p>
Used Space	<p>Displays the amount of space that is used as storage, after any compression or deduplication operations have occurred.</p>

Table 12-1 All Storage overview (*continued*)

Item	Description
Backup Data Written	Displays how much raw backup data is backed up, before compression or deduplication have occurred. For example, if you back up 100 MB of data ten times to a deduplication disk folder, the used space is 100 MB but the amount of the backup data written is 1 GB.
Available Space	Displays the difference between Total Capacity and Used Space .
% of Available Space	Displays the difference as a percentage between Total Capacity and Used Space .
Active and Scheduled Jobs	Displays the number of storage operation jobs, backup jobs, and restore jobs that are running on this device. Click the text in this column to see more details of all jobs that are running or that are scheduled to run.
Compression Ratio	Displays the ratio of the uncompressed size of a file over its compressed size.
Average job rate	Displays the average speed of the jobs that run on this device.

Using the Configure Storage wizard

Use the Configure Storage wizard to set up different types of storage to which you can back up data. The Configure Storage wizard creates the storage with the best possible defaults for your environment. However, you can customize all of the device's setting in the device properties.

You can find a list of compatible types of storage devices in Backup Exec's Hardware Compatibility List:

<http://entsupport.symantec.com/umi/V-269-2>

Click **Configure Storage** on the **Storage** tab to start the Configure Storage wizard.

After Backup Exec is installed and the Backup Exec services are started, any storage that is attached to the Backup Exec server is automatically detected. However, you must use the Configure Storage wizard to configure the storage for backups.

Table 12-2 Storage that you can configure in the Configure Storage wizard

Type of storage	Description
Disk-based storage	<p>Storage that remains attached to the server.</p> <p>Types of disk-based storage include the following:</p> <ul style="list-style-type: none">■ Disk storage A location on a locally attached internal hard drive, a USB device, a FireWire device, or a NAS (network-attached storage) device. See “Configuring disk storage” on page 318.■ Disk cartridge storage Storage that usually remains attached to the server while you remove the media. Disk cartridges use disk cartridge media such as an RDX device, or devices that appear in Windows as removable storage. See “Configuring disk cartridge storage” on page 328.■ Deduplication disk storage A location on a hard drive that reduces the size of backups by storing only unique data. See “Creating or importing deduplication disk storage” on page 809.■ Storage arrays See “Features of the Storage Provisioning Option” on page 1268. Disk arrays that contain multiple disk drives that support data redundancy and failover.■ Virtual disks Virtual disk storage that consists of multiple physical disks in a storage array. See “About virtual disks in the Storage Provisioning Option” on page 1278.■ Vault store Disk storage for the archived data that the Backup Exec Archiving Option archives from one server. See “About vault stores in the Archiving Option” on page 1351.

Table 12-2 Storage that you can configure in the Configure Storage wizard
(continued)

Type of storage	Description
Network storage	<p>Network storage includes the following:</p> <ul style="list-style-type: none">■ NDMP servers Network attached storage (NAS) that supports the Network Data Management Protocol to allow the use of devices that are attached to the servers. See “Features of the NDMP Option” on page 1170.■ OpenStorage devices Network-attached storage that supports Symantec’s OpenStorage technology. See “Configuring an OpenStorage device” on page 804.■ Remote Media Agent for Linux Storage that lets you back up data from remote computers to the storage devices that are directly attached to a Linux server. You can also back up to a simulated tape library on a Linux server. See “About the Remote Media Agent for Linux ” on page 1241.
Tape storage	<p>Tape storage includes the following:</p> <ul style="list-style-type: none">■ Stand-alone tape drives Storage that uses a tape cartridge for reading and writing data. See “Support for tape drives and robotic libraries” on page 357.■ Robotic libraries Storage that contains tape drives, slots, and an automated method for loading tapes. See “Robotic libraries in Backup Exec” on page 366.■ Barcode rules See “Configuring barcode rules for a robotic library ” on page 368.■ Virtual tape libraries (VTLs) See “About the Virtual Tape Library Unlimited Drive Option ” on page 358.

Table 12-2 Storage that you can configure in the Configure Storage wizard
(continued)

Type of storage	Description
Storage pools	<p>Storage pools include the following:</p> <ul style="list-style-type: none">■ Storage device pools■ Managed Backup Exec server pools <p>See “Creating storage device pools” on page 416.</p>
Media sets and vaults	<p>Media sets and vaults are for tape media only, and include the following:</p> <ul style="list-style-type: none">■ Append period■ Overwrite protection period■ Vaulting rules <p>You can also run wizards to update media vaults.</p> <p>See “Managing tapes” on page 378.</p>
Archiving	<p>Archiving lets you configure the following items that are required when you use the Archiving Option:</p> <ul style="list-style-type: none">■ Vault stores■ Vault store partitions <p>See “About vault stores in the Archiving Option” on page 1351.</p> <p>See “About vault store partitions in the Archiving Option” on page 1354.</p>

See [“About storage operation jobs”](#) on page 424.

Viewing details for multiple storage devices

You can select multiple storage devices for which you want to view all jobs, job histories, and active alerts. Viewing details for multiple storage devices lets you see all of the activity for a specific Backup Exec server.

See [“Storage tab overview in Backup Exec”](#) on page 425.

To view details for multiple storage devices

- 1 On the **Storage** tab, Shift + click or Ctrl + click the storage devices, and then right-click one of the selected storage devices.
- 2 Click **Details**.
- 3 In the left pane, click **Jobs**, **Job History**, or **Active Alerts**.

Sending a notification when a scheduled storage operation job completes

You can assign recipients to be notified when a scheduled storage operation job completes. Recipients must be set up before you can set up notification.

See [“About storage operation jobs”](#) on page 424.

To send a notification when a scheduled storage operation job completes

- 1 Create a new scheduled storage operation job or edit an existing one.
See [“About storage operation jobs”](#) on page 424.
- 2 On the storage operation job dialog box, in the left pane, click **Notification**.
- 3 Select the check box for each recipient that you want to notify when each type of storage operation job completes.
- 4 You can continue selecting other options, or click **OK**.

See [“Scheduling a storage operation job”](#) on page 432.

Scheduling a storage operation job

When you schedule a storage operation job, you can configure the time and frequency that you want to run the job.

See [“About storage operation jobs”](#) on page 424.

To schedule a storage operation job

- 1 On the **Storage** tab, right-click the device for which you want to schedule a storage operation job.

If the storage operation can be scheduled, a small arrow appears next to the operation name.
- 2 Click the storage operation, and then click **Schedule**.
- 3 On the storage operation job dialog box, in the left pane, click **Schedule**.

4 Select any of the following options:

Recurrence

Specify a recurrence schedule for the job.

Hours

Create a recurrence pattern that is measured in hours or minutes.

When you select **Hours**, you can configure the following options:

- **Every X hour/minute**

Indicates the number of hours or minutes between the start time of a job and the start time of the next job instance.

- **From**

Designates the starting time for a job to run.

- **Between**

Restricts the job to certain hours and days. For example, if you only want the job to run during business hours, you can select 9:00 AM to 5:00 PM on Monday, Tuesday, Wednesday, Thursday, and Friday.

The start time and end time can span a maximum of 24 hours, however they can cross over midnight into the following day.

Days

Create a recurrence pattern that is measured in days.

When you select **Days**, you must choose between the following options:

- **Every X day**

Indicates the number of days between the start time of a job and the start time of the next job instance.

- **Every weekday**

Specifies that the job should run on Mondays, Tuesdays, Wednesdays, Thursdays, and Fridays.

Weeks

Create a recurrence pattern that is measured in weeks.

When you create a recurrence pattern that is measured in weeks, you must configure the **Every X week on** field. The **Every X week on** specifies the number of weeks between the start time of a job and the start time of the next job instance. It also specifies the days of the week on which the job should run.

Months

Create a recurrence pattern that is measured in months.

When you select **Months**, you must choose between the following options:

- **Day X of every X month**

Specifies the day on which the job should run. It also indicates the number of months between the start time of a job and the start time of the next job instance.

- **Every X X of every X month**

Specifies the day on which the job should run. It also indicates the number of months between the start time of a job and the start time of the next job instance.

- **Selected days of the month**

Specifies the weeks and days of the month on which Backup Exec runs the job. You select the days and weeks on a grid. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current week and day of the month. For example, if you create the job on the third Monday of the month, the default setting is for the job to run once a month on the third Monday.

You can change the default or select additional days on which the job should run. Any additional days that you select are added to the monthly recurrence pattern.

- **Selected dates of the month**

Specifies the dates of the month on which Backup Exec runs the job. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current date of the month. For example, if you create the job on the 15th, the default setting is for the job to run once a month on the 15th.

You can change the default or select additional days on which the job should run. Any additional days that you select are added to the monthly recurrence pattern.

If you select the 31st, the job runs on the last day of the month in months that do not have 31 days. For example, if you configure the job to run on the 31st, in September the job runs on the 30th instead.

Years

Create a recurrence pattern that is measured in years.

When you select **Years**, you can configure the following options:

- **Every X year**

Specifies the number of years between the start time of a job and the start time of the next job instance.

- **On X**

Specifies the date on which Backup Exec runs the job. The date that you select in this field corresponds to the number of years that you selected in the **Every X year** field. So if you selected to run the job every 2 years and you selected June 28th in this field, the job runs every 2 years on June 28th.

- **On the X of X**

Specifies the day and month on which Backup Exec runs the job. The date that you select in this field corresponds to the number of years that you selected in the **Every X year** field. So if you selected to run the job every 2 years and you selected the fourth Thursday of June in this field, the job runs every 2 years on the fourth Thursday of June.

at

Designate the starting time for the first job in the recurrence pattern.

Calendar	View all scheduled jobs on a calendar to check for scheduling conflicts.
Reschedule the job if it does not start x hours after its scheduled start	Specify the amount of time past the job's scheduled start time at which Backup Exec changes the job completion status to Missed. The job is rescheduled to run based on the time window that you configured.
Cancel the job if it is still running x hours after it scheduled start time	Specify the amount of time after the job's scheduled start time at which you want to cancel the job if it is still running. Backup Exec changes the job completion status to Canceled, timed out.
Include dates with the schedule of this job	Specify dates to include with the job schedule. The job runs on all of the dates that you select using this option, in addition to the dates that are part of its normal schedule recurrence. The job resumes its normal schedule on the next day that it is scheduled after an include date.
Exclude dates from the schedule for this job	Specify dates to exclude from the job schedule. The job does not run on any of the dates that you select using this option. It resumes its normal schedule on the next day that a job is scheduled after an excluded date.
Run now with no recurring schedule	Run the job immediately without scheduling any more instances of it for the future.
Run on	Run the job at the time and date that you specify.

Create without a schedule

Create a job without scheduling it. When you use this option, the job does not run at the time of creation and it does not have a recurring schedule. The job remains unscheduled until you choose to run it. You can use a third-party job automation or task scheduling tool to run the job later.

If you use this option to create a job, you cannot place the job on hold. You cannot place jobs on hold unless they are scheduled.

- 5 Click **OK**.

Editing global settings for storage

You can edit the global settings that apply to the robotic libraries, tape media, and disk-based storage that are in your environment.

To edit global settings for storage

- 1 Click the Backup Exec button, click **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, click **Storage**.
- 3 Select the appropriate options.

Inventory robotic libraries when Backup Exec services start

Enable Backup Exec to inventory all of the slots in a robotic library when Backup Exec services start. Depending on the number of slots and robotic libraries, this process may take a few minutes.

This option is not enabled by default.

Full - protect allocated and imported media

Select this option to prevent Backup Exec from overwriting tape media that are in media sets and the media that are imported from another installation of Backup Exec or from another product.

See [“Overwrite protection periods and append periods in media sets”](#) on page 382.

This is the safest option to choose because the tape media that is protected cannot be overwritten until one of the following actions occur:

- The overwrite protection period for the media expires.
- You move the media that belongs to an active media set to scratch media.
- You erase, format, or label the media.
- You move imported media to **Scratch Media**.

Partial - protect only allocated media

Select this option to let Backup Exec overwrite tape media that are imported from another installation of Backup Exec or from another product, or scratch media. Media in a media set that has an overwrite protection period that has not expired (allocated media) cannot be overwritten.

Symantec recommends this option if you want to use media from another installation of Backup Exec or from another product

This option is enabled by default.

Prompt before overwriting imported media

Select this option to be prompted before Backup Exec overwrites tape media that is imported from another installation of Backup Exec or from another product. You must select the option **Partial - protect only allocated media**.

The job cannot run until you respond to this prompt.

None

Select this option to disable the media overwrite protection feature for media in tape drives. With this option, you are responsible for making sure that the media in tape drives are not accidentally overwritten.

When an overwrite job is submitted to a tape drive and the media overwrite protection level is **None**, the media are overwritten.

Note: This option is not recommended because it does not protect data from being overwritten.

Prompt before overwriting allocated or imported media

Select this option to be prompted before Backup Exec overwrites allocated or imported media in tape drives. If you selected **None**, Symantec recommended that you select this option to be prompted before overwriting allocated or imported media.

The job cannot run until you respond to this prompt.

Overwrite scratch media before overwriting recyclable media contained in the targeted media set

Select this option to let Backup Exec overwrite scratch media first in a tape drive when an overwrite job occurs.

See [“How Backup Exec searches for overwritable media in tape drives ”](#) on page 395.

If no scratch media are found in any of the tape drives, Backup Exec overwrites recyclable media in the selected media set.

If no recyclable media are found in the selected media set, Backup Exec searches for recyclable media in any media set.

If no recyclable media are found, Backup Exec automatically searches for other media to overwrite. The media that is overwritten depends on the level of the overwrite protection that you set. If you select this option, more media may be required for the same number of jobs than if you choose to overwrite recyclable media first.

This option affects the order in which Backup Exec overwrites media. If you choose to overwrite scratch media first, the recyclable media may be preserved longer for possible recovery.

This option is enabled by default.

**Overwrite recyclable media contained
in the targeted media set before
overwriting scratch media**

Select this option to let Backup Exec overwrite recyclable media in a tape drive in the selected media set first when an overwrite job occurs.

If no recyclable media are found in any of the tape drives, Backup Exec overwrites scratch media.

If no recyclable media or scratch media are found, Backup Exec searches for media to overwrite. The media that is overwritten depends on the level of the overwrite protection that you set.

See [“How Backup Exec searches for overwritable media in tape drives ”](#) on page 395.

If you choose to overwrite recyclable media in the selected media set first, the same media is re-used more frequently than if you choose to overwrite scratch media first.

Limit Backup Exec to read-only operations on a disk-based storage device if it has been detached for

Select this option to prevent Backup Exec from reclaiming disk space from expired backup sets on any disk-based storage device that is attached after being absent for a number of days. Backup jobs that you send to this device fail. You can change this setting per disk-based storage device in the device's properties.

Before you disable this setting, you may want to view the expiration dates of the backup sets and decide if you want to keep some backup sets longer. You can change the expiration date of a backup set, or you can retain a backup set indefinitely.

See ["Editing disk storage properties"](#) on page 321.

See ["Backup sets "](#) on page 342.

See ["How to restore data from a reattached or reinserted disk-based storage device"](#) on page 327.

See ["How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage"](#) on page 335.

Number of days

Specify how long before Backup Exec is limited to read-only operations on a disk-based storage device when you reattach it.

The default setting is 14 days.

Limit Backup Exec to read-only operations on a disk cartridge if it has not been inserted for

Prevents Backup Exec from reclaiming disk space from expired backup sets on any disk cartridge that is inserted after being absent for a number of days. Backup jobs that you send to this device fail. You can change this setting per disk cartridge in the device's properties.

Before you disable this setting, you may want to view the expiration dates of the backup sets and decide if you want to keep some backup sets longer. You can change the expiration date of a backup set, or you can retain a backup set indefinitely.

See ["Editing disk cartridge properties"](#) on page 329.

See ["Backup sets "](#) on page 342.

See ["How to restore data from a reattached or reinserted disk-based storage device"](#) on page 327.

See ["How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage"](#) on page 335.

Number of days

Specify the number of days that a device can be absent from the Backup Exec server, after which Backup Exec is limited to read-only operations on the device when you re-insert it.

The default setting is 30 days.

**Allow Backup Exec to delete all expired
backup sets**

Select this option to let Backup Exec delete the last full, incremental, and differential backup sets that are necessary to restore a server if the backup sets have expired. By default, Backup Exec keeps the most recent backup sets that are necessary to restore a server, even if the backup sets expire. If you allow Backup Exec to delete all of the expired backup sets for a server, you may not be able to restore that server.

This option is useful if you do not want to keep data after a period of time, usually several years.

However, when you enable this option, you can lose backup sets in the following conditions:

- When the length of time that the backup data or backup sets is kept is less than the frequency of the backup. That is, the backup sets from the last full backup job expire before the next full backup runs. Ensure that when you create jobs, the backup data is kept longer than the amount of time between full backups.
- When the backup job fails or is missed, and is not rerun before the backup sets expire. Monitor any failed or missed jobs, and ensure that you rerun them before the backup sets from the previous full backup expire.

This option is not selected by default.

Note: In a Central Admin Server Option (CASO) environment, this option is only available on the central administration server. If you enable this option on the central administration server, DLM deletes all expired backup sets on the central administration server as well as on all of the managed Backup Exec servers. This option deletes all expired backup sets on both centrally managed and locally managed Backup Exec servers in a CASO environment.

See [“Backup sets ”](#) on page 342.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 335.

4 Click **OK**.

Sharing storage devices

In environments in which there is more than one Backup Exec server, those servers can share storage devices. For example, multiple Backup Exec servers in a CASO environment can share storage devices. In these environments, Backup Exec maintains a database of the shared storage device. Otherwise, the backup data that one server submits to the storage device can overwrite the data that another server submits.

Note: The Enterprise Server Option must be installed before you can share storage devices between Backup Exec servers.

Backup Exec servers can share the following types of storage:

- Storage that is attached to an NDMP server
- Deduplication disk storage
- OpenStorage devices
- Virtual disks
- Disk storage
- Remote Media Agents
- The Backup Exec agents that are configured to send data directly to storage

For disk storage devices and virtual disk, you must specify a UNC path by which the Backup Exec servers can access the storage device. Disk cartridges cannot be shared.

When you share a storage device, you can select which Backup Exec servers can access the storage device. The Backup Exec server from which you added the storage device is automatically enabled to share the storage device. However, you can remove the sharing capability from that Backup Exec server at any time. For example, if you add a storage device to a central administration server, then that server can use the storage device. However, if your environment does not allow

the central administration server to operate as a managed Backup Exec server, then you can remove the sharing capability from the central administration server.

If you have multiple Backup Exec servers and multiple types of storage in your environment, you can select a Backup Exec server and manage the storage for it.

To share a storage device

- 1 On the **Storage** tab, right-click the storage device that you want to share.
- 2 Click **Share**.
- 3 To share a disk storage device or a virtual disk, enter a UNC path by which the servers can access the storage device that you want to share.
- 4 Check the Backup Exec servers or managed Backup Exec servers that you want to share this storage device.

See [“Shared Storage options”](#) on page 449.

- 5 Click **OK**.

Shared Storage options

You can let multiple Backup Exec servers access a storage device.

See [“Sharing storage devices”](#) on page 448.

Table 12-3 Shared Storage options

Item	Description
Enter a UNC path to share X	Indicates the UNC path that the servers can use to access the storage device. Note: This field only displays when you share disk storage or a virtual disk.
Server	Indicates the names of the Backup Exec servers or managed Backup Exec servers that share the storage device.
Status	Indicates the status of the shared storage device, or the status of the Backup Exec server that you selected to share the device. The server must be able to access the UNC path and the storage device. See “Backup Exec server and storage device states” on page 475.

Deleting a storage device

You can delete a storage device from the Backup Exec database. If the storage device is a legacy backup-to-disk folder, a disk storage device, or a deduplication disk storage device, then Backup Exec prompts you to delete the backup sets from the Administration Console view. You can no longer view or select those backup sets from the Administration Console. However, the backup sets remain on the storage device. You must run inventory and catalog operations on the storage device before you can restore from it.

You may want to delete the backup sets from the Administration Console if you move a storage device to another Backup Exec installation. However, if the move is only temporary, then you probably do not want to delete the backup sets. Keeping the backup sets lets you avoid running the inventory and catalog operations on the device when you move it back. You should also keep the backup sets if you plan to recreate the storage device.

You can also use Windows Explorer to navigate to a legacy backup-to-disk folder or to a disk storage, and then delete it. If you use this method, then you cannot recreate the storage in Backup Exec.

To delete a storage device

- 1 On the **Storage** tab, right-click the device that you want to delete, and then click **Disable**.
- 2 Right-click the device again, and then click **Delete**.
- 3 When you are prompted to delete the storage device, click **Yes**.

See [“Disabling and enabling a storage device”](#) on page 455.

See [“Recreating a legacy backup-to-disk folder and its contents”](#) on page 353.

See [“Backup sets ”](#) on page 342.

Changing the state of a storage device to online

Usually, when a device goes offline, an alert appears. The alert message provides a specific reason why the device is offline, and may include a link to the Symantec Knowledge Base for more information.

Correct the problem that caused the device to go offline. Then, for tape drives, robotic libraries, and some other types of devices, you must manually change the state of the device to online.

For disk storage, disk cartridge, storage arrays, and virtual disk devices, Backup Exec detects that the device is online within five minutes and automatically changes the state to online.

See [“Troubleshooting hardware-related issues in Backup Exec”](#) on page 687.

To change the state of a storage device to online

- 1 On the **Storage** tab, right-click the storage device that you want to change to online.
- 2 Click **Offline** to clear the check mark.

Renaming a storage device

You can rename a storage device that is in your environment.

You cannot rename system-defined storage device pools, but you can rename any storage device pools that you create.

See [“About storage operation jobs”](#) on page 424.

To rename a storage device

- 1 On the **Storage** tab, double-click the storage device that you want to rename.
- 2 In the storage device properties, in the **Name** field, type the new name.
- 3 Click **Apply**.

Viewing jobs, job histories, backup sets, and active alerts for storage devices

You can view information that is related to a storage device.

See [“About storage operation jobs”](#) on page 424.

To view jobs, job histories, backup sets, and active alerts for storage devices

- 1 On the **Storage** tab, double-click the storage device for which you want to view the job history, backup sets, or active alerts.
- 2 In the left pane, click **Jobs**, **Job Histories**, **Backup Sets**, or **Active Alerts**.

See [“About the Job History”](#) on page 264.

See [“How to monitor and manage jobs in Backup Exec”](#) on page 253.

See [“About the Job Monitor”](#) on page 255.

See [“Alerts and notifications in Backup Exec”](#) on page 286.

See [“Backup sets ”](#) on page 342.

Cataloging a storage device

You can run a catalog operation to do the following:

- Log the contents of a media that was created by another installation of Backup Exec.
- Create a new catalog on the local hard drive if the catalog for the storage device no longer exists.

Before you can restore or verify data on a storage device, a catalog for that device must exist. If Backup Exec has not used this storage device before, you must run an **Inventory and Catalog** storage operation on the device first.

Note: If a media password was used from a previous release of Backup Exec, Backup Exec catalogs the media as if it is not password-protected. Symantec recommends that you encrypt data instead.

See [“Using encryption with Backup Exec”](#) on page 583.

See [“Inventorying and cataloging a storage device”](#) on page 454.

See [“How Backup Exec catalogs work”](#) on page 248.

To catalog storage

- 1 On the **Storage** tab, right-click the storage device for which you want to create a catalog.
- 2 Click **Catalog**.
- 3 On the catalog dialog box, click **General**, and then enter a name for the job.
- 4 Click **OK**.
- 5 (Optional) View the job log or click the **Job Monitor** tab for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 451.

Scanning a storage device

The scan operation gets information about the media that are in the slots, including barcode information if it is available. Then, the scan operation updates the Backup Exec database with the latest information about where the media are located. When you change magazines or insert new media in a magazine in a robotic library, use the scan operation to update the slot information.

The scan job log reports the barcoded media that are in the drives and portals. If the robotic library is busy, the scan job log may not be able to read the drive and portal status. If the drives and portals can't be read, then none are displayed in the job log. For best results, run the scan when the robotic library is idle.

See [“About storage operation jobs”](#) on page 424.

To scan a storage device now

- 1 On the **Storage** tab, right-click the robotic library or slot that you want to scan.
- 2 Click **Scan**, and then click **Scan now**.
- 3 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

To schedule a scan operation for a storage device

- 1 On the **Storage** tab, right-click the robotic library or slot that you want to scan.
- 2 Click **Scan**, and then click **Schedule**.
- 3 To send notification when the job completes, in the left pane, click **Notification** and select the appropriate options.
- 4 To schedule the job, in the left pane, click **Schedule** and select the appropriate options.

See [“Scheduling a storage operation job”](#) on page 432.

- 5 (Optional) View the scan job log to see which barcoded media are in the slots, drives, and portals of a robotic library, or click the **Job Monitor** tab for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 451.

Inventorying a storage device

You can run an inventory operation to have Backup Exec read a storage device and update the Backup Exec database with information about the media that is on that device.

For robotic libraries, you can inventory all of the slots in the robotic library when you change tapes. You can also select specific slots to inventory. You are not required to re-inventory slots when you add the tapes that Backup Exec requests. For example, if the data that you want to restore is on a tape that is not in the robotic library, you are prompted to insert the correct tape for the restore operation. In this case, you are not required to re-inventory the slot where the tape is inserted. When you add or remove a tape that Backup Exec does not request, you should run an

inventory operation on the changed slots. You can select specific slots to inventory. If you swap tapes often, you may want to run an inventory operation on the robotic library magazine each time that you restart the Backup Exec services.

For tape drives, you can run an inventory operation to mount the media in tape drives and to read the media label. If you change the media that is in a drive, run an inventory operation so that the current media's label appears in the properties. Otherwise, the previous media continues to appear in the properties. There may be a delay as the media is mounted and inventoried in a robotic library.

To inventory a storage device now

- 1 On the **Storage** tab, right-click the storage device that you want to inventory.
- 2 Click **Inventory** and then click **Inventory now** again.

The inventory operation runs. You can view the job log or click the **Job Monitor** tab for details about the job.

To schedule an inventory job for a storage device

- 1 On the **Storage** tab, right-click the storage device that you want to inventory.
- 2 Click **Inventory**, and then click **Schedule**.
- 3 To send notification when the job completes, in the left pane, click **Notification** and select the appropriate options.
- 4 To schedule the job, in the left pane, click **Schedule** and select the appropriate options.

See [“Scheduling a storage operation job”](#) on page 432.

- 5 Click **OK**.
- 6 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 451.

See [“Inventorying robotic libraries when Backup Exec services start”](#) on page 368.

Inventorying and cataloging a storage device

You can run the inventory and the catalog operations together on a storage device, if the device supports both operations.

See [“About storage operation jobs”](#) on page 424.

To inventory and catalog a storage device

- 1 On the **Storage** tab, right-click the storage device that you want to inventory and catalog.
- 2 Click **Inventory and Catalog**.
- 3 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 451.

Pausing and unpausing a storage device

You can pause a storage device to prevent scheduled jobs and new jobs from running on the storage while you perform maintenance activities. Active jobs are not affected if they start before the storage device is paused.

See [“About storage operation jobs”](#) on page 424.

To pause and unpause a storage device

- 1 On the **Storage** tab, right-click the storage device that you want to pause or unpause.
- 2 Do one of the following:
 - To pause the storage device, click **Pause**.
 - To unpause the storage device, right-click it, and then click **Pause** to clear the check mark.

Disabling and enabling a storage device

You can disable a storage device to prevent new jobs from running on it. Backup Exec does not discover disabled NDMP storage devices when the Backup Exec services start.

To disable and enable a storage device

- 1 On the **Storage** tab, right-click the storage device that you want to disable or enable.
- 2 Do one of the following:
 - To disable the storage device, click **Disable**.
 - To enable the storage device, right-click it, and then click **Disable** to clear the check mark.

Initializing a robotic library

You can initialize the robotic library, which sends a startup command to the library.

To initialize a robotic library

- 1 On the **Storage** tab, right-click the robotic library that you want to initialize.
- 2 Click **Initialize**.
- 3 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 451.

Formatting a tape as a WORM tape

You can convert a tape to a write-once, read-many (WORM) tape if the tape drive supports the operation. DLT tape drives support the **Format WORM** operation.

To format a tape as a WORM tape

- 1 On the **Storage** tab, right-click the tape drive that contains the tape that you want to convert to a WORM tape.
- 2 Click **Format WORM**.
- 3 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

See [“How WORM media is used in Backup Exec”](#) on page 401.

Retensioning a tape

Before you run a backup job to a tape, you can run the tape in the tape drive from beginning to end at a fast speed. Retensioning helps the tape wind evenly and run more smoothly past the tape drive heads. Refer to the documentation that came with your tape drive to see how often to run this operation.

This operation is only available if the tape drive supports retensioning.

See [“About storage operation jobs”](#) on page 424.

To retension a tape

- 1 On the **Storage** tab, do either of the following:
 - Right-click the drive that contains the tape that you want to retension.

- Double-click **Slots**, and then right-click the slot that contains the tape that you want to retension.
- 2 Click **Retension**.
 - 3 (Optional) View the job log or click the **Job Monitor** tab for details about the job.
- See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 451.

Formatting a tape in a tape drive

Backup Exec can format the tape in a drive if the drive supports formatting. Formatting a tape may take hours. Most tape drives do not support formatting.

Caution: Formatting erases the tape. All data on the tape is destroyed.

The media label that is displayed was read during the last inventory operation. The media label does not change until another inventory operation occurs. If you change the tape that is in the device but do not inventory the device, the media label that displays may not match the actual media that is in the device.

See [“About storage operation jobs”](#) on page 424.

To format a tape in a tape drive

- 1 On the **Storage** tab, do either of the following:
 - Right-click the tape drive that contains the tape that you want to format.
 - Double-click **Slots**, and then right-click the slot that contains the tape that you want to format.
 - 2 Click **Format**.
 - 3 To format the tape that is displayed, click **Yes**.
 - 4 (Optional) View the job log or click the **Job Monitor** tab for details about the job.
- See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 451.

Ejecting media from a disk cartridge or tape drive

Backup Exec can eject the media that is in a disk cartridge or tape drive. Some devices do not support a software-driven media eject. If the media is a tape, the tape is rewound and you may be instructed to manually remove it.

See [“About storage operation jobs”](#) on page 424.

To eject media now from a disk cartridge or tape drive

- 1 On the **Storage** tab, right-click the disk cartridge or tape drive that you want to eject the media from.
- 2 Click **Eject** and then click **Eject now**.
- 3 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

To schedule an eject operation for a disk cartridge or tape drive

- 1 On the **Storage** tab, right-click the disk cartridge or tape drive that you want to eject the media from.
- 2 Click **Eject**, and then click **Schedule**.
- 3 To send notification when the job completes, in the left pane, click **Notification** and select the appropriate options.
- 4 To schedule the job, in the left pane, click **Schedule** and select the appropriate options.

See [“Scheduling a storage operation job”](#) on page 432.

- 5 Click **OK**.
- 6 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 451.

Cleaning a robotic library drive

You can create and schedule a cleaning job for a robotic library drive.

See [“About storage operation jobs”](#) on page 424.

To create a cleaning job now

- 1 Ensure that you specify the slot that contains the cleaning tape.
See [“Defining a cleaning slot ”](#) on page 373.
- 2 Ensure that the cleaning tape is in the defined cleaning slot and is in the same library that contains the drive that you want to clean.
- 3 On the **Storage** tab, right-click the drive that you want to clean, click **Clean**, and then click **Clean now**.
- 4 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

To schedule a cleaning job

- 1 Ensure that you specify the slot that contains the cleaning tape.
- 2 Ensure that the cleaning tape is in the defined cleaning slot and is in the same library that contains the drive that you want to clean.
- 3 On the **Storage** tab, right-click the drive that you want to clean, click **Clean now**, and then click **Schedule**.
- 4 To send a notification when the job completes, in the left pane, click **Notification** and select the options you want.
- 5 To schedule the job, in the left pane, click **Schedule** and select the options that you want.
See [“Scheduling a storage operation job”](#) on page 432.
- 6 Click **OK**.
- 7 (Optional) View the job history or click the **Job Monitor** tab for details about the job.
See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 451.

Importing media to Backup Exec

You can import media to a robotic library to add tapes to Backup Exec, or to import media that is required for a restore job. When you insert media into a robotic library, you must create an import storage operation job. The import storage operation updates the Backup Exec database with the information about the media. Backup Exec associates the media that you import with a system media set.

See [“Default media sets”](#) on page 379.

Note: You should not associate scratch media with a media set that you create. Backup Exec automatically moves the media to the required media set as needed.

Before you import media, note the following:

- If the media does not have a barcode, you must run the **Inventory after import** operation so that the current media's label appears in the properties. You can only select this operation after you select **Import media now**.
- If the robotic library uses a media magazine, ensure that no jobs are currently running. Before you swap the magazine, ensure that all media are ejected from the drive and are back in the magazine slots.

You can select any number of slots to import media to.

The import storage operation supports robotic libraries with portals. When this storage operation job runs, Backup Exec checks the selected slots for media. If media is found, it is exported to the portals. After all of the media is exported, you are prompted to insert new media into the portal so it can be imported. This process continues until all of the requested media have been imported into the robotic library.

You can also run a scan operation to update the slot information when you insert new media in a robotic library. The scan job log reports the barcoded media that are in the drives and portals.

See [“Scanning a storage device”](#) on page 452.

To import media now

- 1 On the **Storage** tab, do one of the following
 - Expand the robotic library, right-click **Slots**, and then click **Import media now**.
 - Right-click the robotic library, and then click **Import media now**.
- 2 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

To schedule an import media job

- 1 On the **Storage** tab, do one of the following:
 - Expand the robotic library, right-click **Slots**, and then click **Import media now**.
 - Right-click the robotic library, and then click **Import media now**, and then click **Schedule**.
- 2 In the left pane, click **Storage operations**.

- 3 Click the drop-down menu, and select storage operation that you want to schedule:

Import	Updates the Backup Exec database with information about the media.
Inventory after import	Mounts the media in the drive, reads the media label, and updates the Backup Exec database. This operation is necessary for media that do not have barcodes.

- 4 To send a notification when the job completes, in the left pane, click **Notification** and select the appropriate options:

Recipient name	Show the names of the individual and group recipients.
Recipient type	Indicate Recipient for an individual recipient or Group for a group recipient.
Manage Recipients	Add, edit, or delete recipients.
Properties	View or change the properties of a selected recipient.

- 5 To schedule the job, in the left pane, click **Schedule** and select the appropriate options.

Recurrence	Specify a recurrence schedule for the job.
-------------------	--

Hours

Create a recurrence pattern that is measured in hours or minutes.

When you select **Hours**, you can configure the following options:

- **Every X hour/minute**
Indicates the number of hours or minutes between the start time of a job and the start time of the next job instance.
- **From**
Designates the starting time for a job to run.
- **Between**
Restricts the job to certain hours and days. For example, if you only want the job to run during business hours, you can select 9:00 AM to 5:00 PM on Monday, Tuesday, Wednesday, Thursday, and Friday.
The start time and end time can span a maximum of 24 hours, however they can cross over midnight into the following day.

Days

Create a recurrence pattern that is measured in days.

When you select **Days**, you must choose between the following options:

- **Every X day**
Indicates the number of days between the start time of a job and the start time of the next job instance.
- **Every weekday**
Specifies that the job should run on Mondays, Tuesdays, Wednesdays, Thursdays, and Fridays.

Weeks

Create a recurrence pattern that is measured in weeks.

When you create a recurrence pattern that is measured in weeks, you must configure the **Every X week on** field. The **Every X week on** specifies the number of weeks between the start time of a job and the start time of the next job instance. It also specifies the days of the week on which the job should run.

Months

Create a recurrence pattern that is measured in months.

When you select **Months**, you must choose between the following options:

- **Day X of every X month**

Specifies the day on which the job should run. It also indicates the number of months between the start time of a job and the start time of the next job instance.

- **Every X X of every X month**

Specifies the day on which the job should run. It also indicates the number of months between the start time of a job and the start time of the next job instance.

- **Selected days of the month**

Specifies the weeks and days of the month on which Backup Exec runs the job. You select the days and weeks on a grid. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current week and day of the month. For example, if you create the job on the third Monday of the month, the default setting is for the job to run once a month on the third Monday.

You can change the default or select additional days on which the job should run. Any additional days that you select are added to the monthly recurrence pattern.

- **Selected dates of the month**

Specifies the dates of the month on which Backup Exec runs the job. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current date of the month. For example, if you create the job on the 15th, the default setting is for the job to run once a month on the 15th.

You can change the default or select additional days on which the job should run. Any additional days that you select are added to the monthly recurrence pattern.

If you select the 31st, the job runs on the last day of the month in months that do not have 31 days. For example, if you configure the job to run on the 31st, in September the job runs on the 30th instead.

Years

Create a recurrence pattern that is measured in years.

When you select **Years**, you can configure the following options:

- **Every X year**

Specifies the number of years between the start time of a job and the start time of the next job instance.

- **On X**

Specifies the date on which Backup Exec runs the job. The date that you select in this field corresponds to the number of years that you selected in the **Every X year** field. So if you selected to run the job every 2 years and you selected June 28th in this field, the job runs every 2 years on June 28th.

- **On the X of X**

Specifies the day and month on which Backup Exec runs the job. The date that you select in this field corresponds to the number of years that you selected in the **Every X year** field. So if you selected to run the job every 2 years and you selected the fourth Thursday of June in this field, the job runs every 2 years on the fourth Thursday of June.

at

Designate the starting time for the first job in the recurrence pattern.

Calendar

View all scheduled jobs on a calendar to check for scheduling conflicts.

Reschedule the job if it does not start x hours after its scheduled start

Specify the amount of time past the job's scheduled start time at which Backup Exec changes the job completion status to Missed. The job is rescheduled to run based on the time window that you configured.

Cancel the job if it is still running x hours after it scheduled start time	Specify the amount of time after the job's scheduled start time at which you want to cancel the job if it is still running. Backup Exec changes the job completion status to Canceled, timed out.
Include dates with the schedule of this job	Specify dates to include with the job schedule. The job runs on all of the dates that you select using this option, in addition to the dates that are part of its normal schedule recurrence. The job resumes its normal schedule on the next day that it is scheduled after an include date.
Exclude dates from the schedule for this job	Specify dates to exclude from the job schedule. The job does not run on any of the dates that you select using this option. It resumes its normal schedule on the next day that a job is scheduled after an excluded date.
Run now with no recurring schedule	Run the job immediately without scheduling any more instances of it for the future.
Run on	Run the job at the time and date that you specify.
Create without a schedule	<p>Create a job without scheduling it. When you use this option, the job does not run at the time of creation and it does not have a recurring schedule. The job remains unscheduled until you choose to run it. You can use a third-party job automation or task scheduling tool to run the job later.</p> <p>If you use this option to create a job, you cannot place the job on hold. You cannot place jobs on hold unless they are scheduled.</p>

- 6 Click **OK**
- 7 (Optional) View the job history or click the **Job Monitor** tab for details about the job.
- See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 451.

Exporting media and expired media

The export media operation supports robotic libraries that have portals. When this operation is run on one or more robotic library slots, the exported media is placed

in the portals. If you select more media than there are portals, the robotic library fills as many slots as possible. Then, you are prompted to remove the media from the portal. This process continues until all of the selected media have been removed from the robotic library. You can also export expired media from a robotic library.

The export expired media operation lets you automate media handling in robotic libraries. This operation removes the media that Backup Exec cannot write to. You can then use the **Import after export** operation to add scratch media to the robotic library to prepare for the next backup.

After you export the expired media from the robotic library, the expired media appears in **Offline Tape**. If the media is in a media set that has an applicable vault media rule, then the media appears in the vault location.

You can export cleaning media with the export expired media storage operation. You can include all cleaning media, or all cleaning media that has been used more than a specified number of times.

See [“Importing media to Backup Exec ”](#) on page 459.

Note: Not all storage operations are available for all devices.

To export media or export expired media

- 1 On the **Storage** tab, do one of the following:
 - Expand the robotic library, right-click **Slots**, and then click **Export media**.

- Right-click the robotic library, and then click **Export**.
- 2 Do one of the following:
- | | |
|--|--|
| To immediately export only the media that Backup Exec cannot write to and place it in the portal | Click Export expired media now .
The operation runs. You can view the job history for details about the job. |
| To immediately export the media to the portal | Click Export media now .
The operation runs. You can view the job history for details about the job. |
| To immediately import media after the export | Click Import after export .
The operation runs. You can view the job history for details about the job. |
- 3 (Optional) View the job history or click the **Job Monitor** tab for details about the job.
- See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 451.

To schedule an export media or export expired media operation

- 1 On the **Storage** tab, do one of the following:
 - Expand the robotic library, right-click **Slots**, and then click **Export media**.
 - Right-click the robotic library, and then click **Export**.
- 2 Click **Schedule**.
- 3 In the left pane, click **Storage operations**.

- 4 Click the drop-down menu and select one of the following storage operations that you want to schedule:

Export	<p>Places the media into the robotic library's portals.</p> <p>If you select more media than there are portals, the robotic library fills as many slots as possible. Then you are prompted to remove the media from the portal. This process continues until all of the selected media have been removed from the robotic library.</p>
Import media after export	<p>Adds the scratch media to the robotic library to prepare for the next backup.</p>
Export expired media	<p>Places the expired media into the robotic library's portals.</p> <p>This operation lets you automate media handling in robotic libraries by removing the media that Backup Exec cannot write to. After you export the expired media from the robotic library, the expired media appears in Offline Tape. If the media is in a media set that has an applicable vault media rule, then the media appears in the vault location.</p>

- 5 To send notification when the job completes, in the left pane, click **Notification** and select the appropriate options.
- 6 To schedule the job, in the left pane, click **Schedule** and select the appropriate options.
- 7 Click **OK**.
- 8 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

Locking and unlocking the robotic library's front portal

By default, the robotic library portal is not locked, even when you run the lock storage operation. For instructions on how to lock the library portal, go to:

<http://www.symantec.com/docs/TECH67698>

You must create a job to unlock the robotic library's front portal.

To unlock the robotic library's front portal

- 1 On the **Storage** tab, right-click the robotic library that has the front portal that you want to unlock.
- 2 Click **Unlock**, and then click **Unlock now**.
- 3 (Optional) View the job history or click the **Job Monitor** tab for details about the job.

To schedule a job to unlock the robotic library's front portal

- 1 On the **Storage** tab, right-click the robotic library that has the front portal that you want to unlock.
- 2 Click **Unlock**.
- 3 Click **Schedule**.
- 4 To send notification when the job completes, in the left pane, click **Notification** and select any of the following options:

Recipient name	Show the names of the individual and group recipients.
Recipient type	Indicate Recipient for an individual recipient or Group for a group recipient.
Manage Recipients	Add, edit, or delete recipients.
Properties	View or change the properties of a selected recipient.

- 5 To schedule the job, in the left pane, click **Schedule** and select the appropriate options.

Recurrence	Specify a recurrence schedule for the job.
------------	--

Hours

Create a recurrence pattern that is measured in hours or minutes.

When you select **Hours**, you can configure the following options:

- **Every X hour/minute**
Indicates the number of hours or minutes between the start time of a job and the start time of the next job instance.
- **From**
Designates the starting time for a job to run.
- **Between**
Restricts the job to certain hours and days. For example, if you only want the job to run during business hours, you can select 9:00 AM to 5:00 PM on Monday, Tuesday, Wednesday, Thursday, and Friday.
The start time and end time can span a maximum of 24 hours, however they can cross over midnight into the following day.

Days

Create a recurrence pattern that is measured in days.

When you select **Days**, you must choose between the following options:

- **Every X day**
Indicates the number of days between the start time of a job and the start time of the next job instance.
- **Every weekday**
Specifies that the job should run on Mondays, Tuesdays, Wednesdays, Thursdays, and Fridays.

Weeks

Create a recurrence pattern that is measured in weeks.

When you create a recurrence pattern that is measured in weeks, you must configure the **Every X week on** field. The **Every X week on** specifies the number of weeks between the start time of a job and the start time of the next job instance. It also specifies the days of the week on which the job should run.

Months

Create a recurrence pattern that is measured in months.

When you select **Months**, you must choose between the following options:

- **Day X of every X month**

Specifies the day on which the job should run. It also indicates the number of months between the start time of a job and the start time of the next job instance.

- **Every X X of every X month**

Specifies the day on which the job should run. It also indicates the number of months between the start time of a job and the start time of the next job instance.

- **Selected days of the month**

Specifies the weeks and days of the month on which Backup Exec runs the job. You select the days and weeks on a grid. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current week and day of the month. For example, if you create the job on the third Monday of the month, the default setting is for the job to run once a month on the third Monday.

You can change the default or select additional days on which the job should run. Any additional days that you select are added to the monthly recurrence pattern.

- **Selected dates of the month**

Specifies the dates of the month on which Backup Exec runs the job. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current date of the month. For example, if you create the job on the 15th, the default setting is for the job to run once a month on the 15th.

You can change the default or select additional days on which the job should run. Any additional days that you select are added to the monthly recurrence pattern.

If you select the 31st, the job runs on the last day of the month in months that do not have 31 days. For example, if you configure the job to run on the 31st, in September the job runs on the 30th instead.

Years

Create a recurrence pattern that is measured in years.

When you select **Years**, you can configure the following options:

- **Every X year**
Specifies the number of years between the start time of a job and the start time of the next job instance.
- **On X**
Specifies the date on which Backup Exec runs the job. The date that you select in this field corresponds to the number of years that you selected in the **Every X year** field. So if you selected to run the job every 2 years and you selected June 28th in this field, the job runs every 2 years on June 28th.
- **On the X of X**
Specifies the day and month on which Backup Exec runs the job. The date that you select in this field corresponds to the number of years that you selected in the **Every X year** field. So if you selected to run the job every 2 years and you selected the fourth Thursday of June in this field, the job runs every 2 years on the fourth Thursday of June.

at

Designate the starting time for the first job in the recurrence pattern.

Calendar

View all scheduled jobs on a calendar to check for scheduling conflicts.

Reschedule the job if it does not start x hours after its scheduled start

Specify the amount of time past the job's scheduled start time at which Backup Exec changes the job completion status to Missed. The job is rescheduled to run based on the time window that you configured.

Cancel the job if it is still running x hours after it scheduled start time	Specify the amount of time after the job's scheduled start time at which you want to cancel the job if it is still running. Backup Exec changes the job completion status to Canceled, timed out.
Include dates with the schedule of this job	Specify dates to include with the job schedule. The job runs on all of the dates that you select using this option, in addition to the dates that are part of its normal schedule recurrence. The job resumes its normal schedule on the next day that it is scheduled after an include date.
Exclude dates from the schedule for this job	Specify dates to exclude from the job schedule. The job does not run on any of the dates that you select using this option. It resumes its normal schedule on the next day that a job is scheduled after an excluded date.
Run now with no recurring schedule	Run the job immediately without scheduling any more instances of it for the future.
Run on	Run the job at the time and date that you specify.
Create without a schedule	<p>Create a job without scheduling it. When you use this option, the job does not run at the time of creation and it does not have a recurring schedule. The job remains unscheduled until you choose to run it. You can use a third-party job automation or task scheduling tool to run the job later.</p> <p>If you use this option to create a job, you cannot place the job on hold. You cannot place jobs on hold unless they are scheduled.</p>

6 Click **OK**.

See [“Viewing jobs, job histories, backup sets, and active alerts for storage devices”](#) on page 451.

Backup Exec server and storage device states

Backup Exec servers and storage devices display a state that indicates their current condition.

Table 12-4 Possible states for Backup Exec servers and storage devices

State	Description
All the Backup Exec services need to be restarted on <Backup Exec server>	The Backup Exec services and the Backup Exec deduplication services must be restarted. See “Starting and stopping Backup Exec services” on page 621.
An error occurred while discovering this device. Cycle the services on <Backup Exec server> to retry device discovery.	The Backup Exec services must be restarted. See “Starting and stopping Backup Exec services” on page 621.
Active	The storage device is in use by a job.
Configuration failed	Configuration has failed for a local disk storage device or virtual disk.
Configuring	A local disk storage device or virtual disk is in the process of configuration.
Disabled	The storage device is disabled and Backup Exec cannot use it. The device is available for other applications.
Disabled; Active	The storage device's status was changed to Disabled while a job was running to the device.
Low disk space; Active	The storage device is in a low disk space condition, but is currently in use by a job.
Low disk space	The storage device has low disk space. See “Properties for virtual disks on storage arrays” on page 1297.
No communication	Communications have stopped between a managed Backup Exec server and a central administration server in a Central Admin Server Option environment. See “What happens when CASO communication thresholds are reached” on page 1131.
Not configurable	The disk cannot be configured because it is in a bad state, or it has failed.

Table 12-4 Possible states for Backup Exec servers and storage devices
(continued)

State	Description
Not configured	The disk is available for configuration but has not yet been configured.
Offline	<p>The storage device is offline.</p> <p>A storage device can appear offline if any of the following actions occur:</p> <ul style="list-style-type: none">■ The device was turned off after Backup Exec started.■ The device was being used by another application when Backup Exec started.■ The device is removed from the server.■ The device reports a critical error.■ The firmware of the device was updated. <p>Usually, when a device goes offline, an alert appears. The alert message provides a specific reason why the device is offline, and may include a link to the Symantec Knowledge Base for more information.</p> <p>Correct the problem that caused the device to go offline. Then, for tape drives, robotic libraries, and some other types of devices, you must manually change the state of the device to online. For disk storage, disk cartridge, storage arrays, and virtual disk devices, Backup Exec detects that the device is online within five minutes and automatically changes the state to online</p> <p>See "Changing the state of a storage device to online" on page 450.</p>
Online	The storage device is online.
Paused	<p>The storage device is paused.</p> <p>See "Pausing and unpausing a storage device" on page 455.</p>
Paused; Active	The storage device is paused, but is currently in use by a job.

Table 12-4 Possible states for Backup Exec servers and storage devices
(continued)

State	Description
Stalled	Communications have stalled during communications between a managed Backup Exec server and a central administration server in a Central Admin Server Option environment. See “What happens when CASO communication thresholds are reached” on page 1131.
The Backup Exec deduplication services need to be restarted on <Backup Exec server>	The Backup Exec deduplication services should be restarted. The deduplication services are separate from the Backup Exec services so the Backup Exec services are not affected. See “Starting and stopping Backup Exec services” on page 621.
The Backup Exec services on <Backup Exec server> need to be restarted	The Backup Exec services must be restarted. See “Starting and stopping Backup Exec services” on page 621.
This device has not been discovered correctly. Cycle the services on <Backup Exec server> to retry device discovery.	A state that can occur after you add a new storage device to Backup Exec. You must restart the Backup Exec services so that the device discovery process can run again.
Uninitialized	The device has not been initialized.

Conversion to virtual machines

This chapter includes the following topics:

- [How conversion of physical computers to virtual machines works in Backup Exec](#)
- [Requirements for conversion to virtual machine jobs](#)
- [Converting to a virtual machine simultaneously with a backup job](#)
- [Converting to a virtual machine after a backup job](#)
- [Adding a conversion to virtual machine stage to a backup job](#)
- [Converting to a virtual machine from a point in time](#)
- [Creating a one-time conversion to a virtual machine](#)
- [Setting default options for conversion to virtual machine jobs](#)

How conversion of physical computers to virtual machines works in Backup Exec

Backup Exec provides the ability to convert a physical computer to a virtual machine in the following ways:

- Back up a physical computer and simultaneously convert it to a virtual machine.
- Back up a physical computer and schedule a conversion to a virtual machine to run after the backup job runs.
- Convert existing backup sets to a virtual machine.

- Convert a running physical computer to a virtual machine without running a backup job.

Note: Conversion of a physical Hyper-V host into a virtual machine is not supported.

The newly created virtual machine is bootable and is identical to the physical computer from which the virtual machine was converted, with the exception of the network cards and settings. Conversion to a virtual machine enables business continuity for both Hyper-V and VMware environments.

This topic includes the following information:

[Conversion to virtual machine options](#)

[How backup selections are processed during conversion to virtual jobs](#)

[How full, incremental, and differential backups work in conversion to virtual jobs](#)

[Conversion of disks or volumes larger than 2 terabytes](#)

[Notes about conversion to a virtual machine running on a Windows Server 2012 Hyper-V host](#)

[Notes about conversion of Exchange servers](#)

Conversion to virtual machine options

You use one of the following options on the **Backup and Restore** tab to set up a conversion to a virtual machine:

Table 13-1 Conversion to virtual machine options

Name of option	Description
Back up to Disk and Simultaneously Convert to Virtual Machine Back up to Deduplication Disk Storage and Simultaneously Convert to Virtual Machine	<p>These options run the conversion simultaneously with the backup job. Because two operations are performed at the same time, this job may take longer to run than a regular backup job. A large backup window is recommended for this option.</p> <p>A conversion from a full backup creates the new virtual machine. Incremental and differential backups update the virtual machine that was created from the full backup.</p> <p>Note: Incremental backups are preferred over differential backups because the differential backups are inefficient for conversion as compared to incremental backups.</p> <p>Although the backup runs simultaneously with the conversion, the backup is the primary job. Therefore, if the backup fails, then the conversion fails also. However, if the conversion fails, the backup continues to run. For a conversion failure, the job is marked as a success with exceptions. In the case of a failed conversion, the conversion process runs again during the next full backup.</p> <p>See “Converting to a virtual machine simultaneously with a backup job” on page 488.</p>

Table 13-1 Conversion to virtual machine options (*continued*)

Name of option	Description
Back Up to Disk and then Convert to Virtual Machine Back Up to Deduplication Disk Storage and then Convert to Virtual Machine	<p>These options let you schedule the conversion to run after the backup job. These options require a smaller backup window than the simultaneous conversion options.</p> <p>A conversion from a full backup creates the new virtual machine. Incremental and differential backups update the virtual machine that was created from the full backup.</p> <p>Note: Incremental backups are preferred over differential backups because the differential backups are inefficient for conversion as compared to incremental backups.</p> <p>See “Converting to a virtual machine after a backup job” on page 492.</p>
Convert to Virtual Machine from Point-in-Time	<p>A conversion to a virtual machine from a point in time converts existing backup sets from a backup job in which all components that are necessary for a virtual machine conversion were selected. When all necessary components are selected for a backup job, Backup Exec identifies that job as Fully selected and the Simplified Disaster Recovery option has a status of ON. The option to convert to a virtual machine from a point in time is useful in a disaster recovery situation in which you want to quickly recover a failed server. The backup sets contain all of the critical components of the server. Additionally, you can select application data or user data to include in the conversion.</p> <p>Note: The option Convert from Point-in-Time becomes available for selection only after you run at least one full backup that includes all critical system components.</p> <p>See “Converting to a virtual machine from a point in time” on page 503.</p>

Table 13-1 Conversion to virtual machine options (*continued*)

Name of option	Description
One-Time Convert to Virtual Machine	<p>This option converts a running physical computer to a virtual machine without a separate backup job. A one-time conversion job can be scheduled to run at a later time, but it cannot be scheduled to run more than one time.</p> <p>Only Full (Copy) backups are supported for this type of conversion. Incremental and differential backups are not supported for one-time conversions.</p> <p>Note: Copy backups are not supported in Hyper-V 2008 or 2008 R2. A crash-consistent conversion is created instead.</p> <p>See “Creating a one-time conversion to a virtual machine” on page 505.</p>
Add Stage	<p>You can add a stage to a backup job to convert to a virtual machine. Two types of stages are available: Convert to Virtual Machine After Backup and Convert to Virtual Machine Simultaneously with Backup.</p> <p>See “Adding a conversion to virtual machine stage to a backup job” on page 498.</p>

Note: Regardless of the option that is used to initiate the conversion, Backup Exec does not power on the virtual machine after creating it.

Backup Exec creates a snapshot of the virtual machine at the end of the conversion process. The snapshot is removed before the next job runs as long as the virtual machine is not powered on and the only snapshot on the virtual machine is the one that Backup Exec created. If you want to start using the virtual machine, you must manually remove the snapshot.

If the converted virtual machine's host fails and you bring the virtual machine online, the existing conversion job continues to run and then fails. In this situation, you must create a new conversion job.

How backup selections are processed during conversion to virtual jobs

When you set up a conversion to virtual job, you select the items to back up before the conversion or simultaneously with the conversion. Backup Exec may automatically exclude or include data in certain situations. The job log lists the items that were excluded from or included in the conversion.

Selections are processed as follows:

- If you exclude a volume from the backup job, then that volume is automatically excluded from the corresponding conversion job.
- If you include an application in the backup job, the disk and volume on which that application resides are automatically included in the conversion job.
- If you exclude a volume from the backup, the disk that is part of the volume is automatically excluded if the volume is the only thing on the disk.

How full, incremental, and differential backups work in conversion to virtual jobs

Conversion-to-virtual-machine jobs create a virtual machine from a full backup. Subsequent incremental and differential backup jobs update the virtual machine that was created by the full backup. Although differential backups are supported for conversion jobs, incremental backups are the preferred method for updating virtual machines. Differential backups are inefficient for conversion when compared to incremental backups.

When using the incremental backup method for conversion to virtual machine jobs, keep in mind the following points:

- Data from volumes is backed up at a file\folder level. Even if only a portion of the file has changed, the entire file is backed up.
- The entire System State is backed up. Incremental backups are supported only for the system file components of System State; the other System State components are backed up as a full backup.
- Block-level backup methods must be selected for SQL. After the initial incremental backup runs for SQL, block-level backups are performed for any subsequent full or incremental backups for which the **Simplified Disaster Recovery** option is **ON**. Conversion will not be performed if a block-level backup method is not selected.
- Full backups are always performed for Active Directory. Neither incremental backups nor differential backups are supported for Active Directory.

Conversion of disks or volumes larger than 2 terabytes

Backup Exec supports conversion of disks or volumes that are larger than 2 terabytes (TB) for VMware hosts and Hyper-V 2012 hosts.

Note: If the boot volume or the system volume on the source physical computer is larger than 2 TB and you are converting to a VMware or Hyper-V 2008 R2 virtual machine, then Backup Exec cannot convert it. A boot volume cannot be split or created as a dynamic spanned volume. This limitation also applies if the boot volume on the source is larger than the destination datastore's maximum supported disk size.

Backup Exec converts all disks and volumes on the destination to simple or spanned dynamic volumes. Even if the source disk is a basic disk, the converted disks on the destination will be dynamic. Since Backup Exec supports only simple and spanned dynamic volumes, if the source volume type is striped, mirror, or raid5, then Backup Exec converts that volume to either simple or spanned on the destination.

The following additional information applies only to VMware:

- Conversion of a volume that is larger than 2 TB is supported in the following situations:
 - If the volume is not a boot or system volume
 - If the volume was created on a dynamic disk
 - If the volume is an MBR disk
- Conversion of GPT disks is not supported, regardless of the size of the volume on the disk. Conversions of GPT disks fail.
- In situations where the VMware source disk size is larger than the destination datastore's maximum supported size, the source disk is split into multiple disks. Backup Exec splits the disk based on the destination datastore's maximum allowed size.

Notes about conversion to a virtual machine running on a Windows Server 2012 Hyper-V host

Before you create a job to convert a physical computer to a Windows Server 2012 Hyper-V host, review the following information:

- Disk data is stored in vhdx files for conversion of a physical computer to a virtual machine running on a Windows Server 2012 Hyper-V host. The vhdx files can have a maximum capacity of 64 TB. The physical computer's disk sector size is maintained during the conversion.

- The conversion of physical computers that have simple GPT disks is supported.
- The conversion of physical computers that have dynamic disks is not supported.
- Storage Spaces and Storage Pools are not supported.
- If the physical computer runs Windows Server 2012 with an ReFS volume, conversion to a Windows Server 2012 Hyper-V host is supported. Conversion to any previous versions of a Windows Hyper-V host is not supported, so those jobs fail
- If the physical computer runs Windows Server 2012 with one or more Windows deduplication volumes, conversion to a Hyper-V host is possible, but it may fail. The converted disk data is not deduplicated. In other words, an unoptimized data transfer is performed. For this reason, the conversion may fail if the amount of unoptimized data is greater than the capacity of the destination volume.

Notes about conversion of Exchange servers

Backup Exec disables Exchange services on a newly-created virtual machine after a conversion-to-virtual job. If Exchange databases are present on a converted server, after a conversion-to-virtual job you must open the Microsoft Services Control Manager and manually restart the following Exchange services:

- MExchangeDagMgmt
- MExchangeADTopology
- MExchangeAntispamUpdate
- MExchangeDiagnostics
- MExchangeEdgeSync
- MExchangeFrontEndTransport
- MExchangeHM
- MExchangeMailboxAssistants
- MExchangeDelivery
- MExchangeSubmission
- MExchangeMigrationWorkflow
- MExchangeMailboxReplication
- MExchangeRPC
- MExchangeFastSearch
- HostControllerService
- MExchangeServiceHost

- MExchangeThrottling
- MExchangeTransport
- MExchangeTransportLogSearch
- MExchangeUM
- MExchangeUMCR

See [“Requirements for conversion to virtual machine jobs”](#) on page 487.

Requirements for conversion to virtual machine jobs

Before you use the conversion to virtual machine feature, review the following requirements:

- The option **Simplified Disaster Recovery** must have a status of **ON** on the **Browse** tab of the **Backup Selections** dialog box.

Note: The Agent for VMware and Hyper-V is not required for conversion to virtual machines.

- Only Windows servers are supported.
- In a VMware environment, if you convert a physical server to a virtual server and then want to back up the converted server, you must push-install the Agent for Windows to the converted server. Push-installing the Agent for Windows installs the Backup Exec VSS provider. Note that you need to install the Agent for Windows on the converted server even if you installed the Agent for Windows on the physical server before you converted it to a virtual server.
- For conversion in a Hyper-V environment, the following additional requirements apply:
 - The Agent for Windows must be installed on the Hyper-V host to which the conversion is sent.
 - The maximum disk size is 2 TB for Hyper-V 2008/2008 R2 and earlier. Disks larger than 2TB are supported for Hyper-V 2012.
 - Only basic disks are supported. Dynamic disks are not supported for Hyper-V conversions.

Note: Conversion of a physical Hyper-V host into a virtual machine is not supported.

- For specific operating system requirements for conversion to virtual machines, see the Backup Exec Software Compatibility List at the following URL:
<http://entsupport.symantec.com/umi/V-269-1>

See “How conversion of physical computers to virtual machines works in Backup Exec” on page 479.

Converting to a virtual machine simultaneously with a backup job

With this type of a conversion, the backup and the conversion run at the same time.

Note: If the backup fails, then the conversion fails also. However, if the conversion fails, the backup continues to run. For a conversion failure, the job is marked as a success with exceptions and the conversion process runs again during the next full backup.

To convert to a virtual machine simultaneously with a backup job

- 1 On the **Backup and Restore** tab, right-click the server that contains the data you want to back up and convert.
- 2 Select **Backup**, and then select **Back up to Disk and Simultaneously Convert to Virtual Machine** or **Back up to Deduplication Disk Storage and Simultaneously Convert to Virtual Machine**, depending on the type of storage device that you want to use.
- 3 Do any of the following:

To change the backup selections

Do the following:

- In the <Name of Server> box, click **Edit**, and then select the items to back up.
Note: The option **Simplified Disaster Recovery** must have a status of **ON**.
- Click **OK**.

To change the backup options

Do the following:

- In the **Backup** box, click **Edit**, and then change the backup options as needed.
- Click **OK**.

- 4 In the **Conversion to Virtual** box, click **Edit** to set the options for conversion.
- 5 In the **Convert for** field, select either **Hyper-V** or **VMware ESX/vCenter server**.
- 6 Configure the conversion options:

If you selected Hyper-V in step 5

Do the following to configure the conversion options:

- Click the arrow in the **Hyper-V server name** field, select the name of the server where you want to create the virtual machine, and then click **Add**.
- In the **Destination drive or path** field, enter the location on the physical computer where the virtual disks should be created. Enter a drive letter and path.
- If you want to change the default virtual machine name, type the new name in the **Virtual machine name** field.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of Hyper-V Integration Components ISO image** field, enter the location of your Hyper-V Integration Components ISO image. The ISO image is needed to make the virtual machine bootable.
- If you want to change the CPU count or amount of physical RAM for the destination virtual machine, on the **Server configuration** tab, enter the new amounts in the **Desintation virtual machine** fields.
- If you want to change the disk type, the controller, or the virtual disk location for one of the disks, on the **Disk configuration** tab, click **Edit disk configuration**, and then enter the new information.

If you selected VMware ESX/vCenter
server in step 5

Do the following to configure the conversion options:

- Click the arrow in the **ESX/vCenter server name** field, and click the name of the server where you want to create the virtual machine.
- Click the arrow in the **Logon account** field, and then select the appropriate logon account for the server that you selected.
- Click **Select** next to the ESX/vCenter server name field. Backup Exec fills in the remaining information about the server.
- Browse to select the virtual machine folder and resource pool that are associated with the server that you selected.

Note: If you target an ESX server and select a resource pool, the newly created virtual machine is not added to the resource pool automatically. You can manually move the virtual machine into the appropriate resource pool after the conversion.

- If you want to change the default virtual machine name, type the new name in the **Virtual machine name** field.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of VMware Tools ISO image** field, enter the location of your VMware Tools ISO image. The path should be accessible with the default credentials. The path should also be local to the Backup Exec server. The ISO image is needed to make the virtual machine bootable.

Note: In a CASO environment, this path should be local to the managed Backup Exec server to which the job is targeted.

- If you want to change the CPU count or amount of physical RAM for the destination virtual machine, on the **Server configuration** tab, enter the new amounts in the **Desintation virtual machine** fields.
- If you want to change the disk type, the controller, or the virtual disk location for one of the disks, on the **Disk configuration** tab, click **Edit disk configuration**, and then enter the new information.

7 Click **OK**.

8 On the **Backup Definition Properties** dialog box, click **OK** to create the job.

See [“How conversion of physical computers to virtual machines works in Backup Exec”](#) on page 479.

Converting to a virtual machine after a backup job

Backup Exec sets up this type of conversion as a stage that runs after the backup job runs. The backup sets that are created from the backup job are used to create the virtual machine.

To convert to a virtual machine after a backup job

- 1 On the **Backup and Restore** tab, right-click the server that contains the data you want to back up and convert.
- 2 Select **Backup**, and then select **Back Up to Disk and then Convert to Virtual Machine** or **Back Up to Deduplication Disk Storage and then Convert to Virtual Machine**, depending on the type of storage device that you want to use.
- 3 Do any of the following:

To change the backup selections

Do the following:

- In the <Name of Server> box, click **Edit**, and then select the items to back up.

Note: The option **Simplified Disaster Recovery** must have a status of **ON**.

- Click **OK**.

To change the backup options

Do the following:

- In the **Backup** box, click **Edit**, and then change the backup options as needed.
- Click **OK**.

4 In the **Conversion to Virtual** box, click **Edit**.

5 In the left pane, select **Schedule**, and then select one of the following options:

To schedule the conversion to run at a specific time

Do the following:

- Select **According to schedule**.
- Click the arrow in the **Source** field to select either all backups or the most recent full backup as the source to initiate the conversion.
- Select **Recurrence**, and then click the arrow to set the recurrence pattern.

To create the conversion job without scheduling it

Do the following:

- Select **According to schedule**.
- Click the arrow in the **Source** field to select either all backups or the most recent full backup as the source to initiate the conversion.
- Select **Create without a schedule**.
 When you use this option, the job does not run at the time of creation and it does not have a recurring schedule. The job remains in a pending state until you choose to run it. You can use a third-party job automation or task scheduling tool to run the job later.

To run the conversion immediately after the backup job completes

Select **Convert to virtual immediately after source task completes**.

- 6 Optional: In the left pane, select **Notification** to notify selected recipients when the job completes.
- 7 In the left pane, select **Conversion Settings** to set the options for the conversion.
- 8 In the **Convert for** field, select either **Hyper-V** or **VMware ESX/vCenter server**.
- 9 Configure the conversion options:

If you selected Hyper-V in step 8

Do the following to configure the conversion options:

- Click the arrow in the **Hyper-V server name** field, select the name of the server where you want to create the virtual machine, and then click **Add**.
- In the **Destination drive or path** field, enter the location on the physical computer where the virtual disks should be created. Enter a drive letter and path.
- If you want to change the default virtual machine name, type the new name in the **Virtual machine name** field.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of Hyper-V Integration Components ISO image** field, enter the location of your Hyper-V Integration Components ISO image. The ISO image is needed to make the virtual machine bootable.
- If you want to change the CPU count or amount of physical RAM for the destination virtual machine, on the **Server configuration** tab, enter the new amounts in the **Desintation virtual machine** fields.
- If you want to change the disk type, the controller, or the virtual disk location for one of the disks, on the **Disk configuration** tab, click **Edit disk configuration**, and then enter the new information.

If you selected VMware ESX/vCenter
server in step 8

Do the following to configure the conversion options:

- Click the arrow in the **ESX/vCenter server name** field, and click the name of the server where you want to create the virtual machine.
- Click the arrow in the **Logon account** field, and then select the appropriate logon account for the server that you selected.
- Click **Select** next to the ESX/vCenter server name field. Backup Exec fills in the remaining information about the server.
- Browse to select the virtual machine folder and resource pool that are associated with the server that you selected.

Note: If you target an ESX server and select a resource pool, the newly created virtual machine is not added to the resource pool automatically. You can manually move the virtual machine into the appropriate resource pool after the conversion.

- If you want to change the default virtual machine name, type the new name in the **Virtual machine name** field.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of VMware Tools ISO image** field, enter the location of your VMware Tools ISO image. The path should be accessible with the default credentials. The path should also be local to the Backup Exec server. The ISO image is needed to make the virtual machine bootable.

Note: In a CASO environment, this path should be local to the managed Backup Exec server to which the job is targeted.

- If you want to change the CPU count or amount of physical RAM for the destination virtual machine, on the **Server configuration** tab, enter the new amounts in the **Desintation virtual machine** fields.
- If you want to change the disk type, the controller, or the virtual disk location for one of the disks, on the **Disk configuration** tab, click **Edit disk configuration**, and then enter the new information.

10 Click **OK** to save your selections.

11 On the **Backup Definition Properties** dialog box, click **OK** to create the job.

See [“How conversion of physical computers to virtual machines works in Backup Exec”](#) on page 479.

Adding a conversion to virtual machine stage to a backup job

You can add a stage to a backup definition to convert a backup to a virtual machine. A conversion to virtual machine job requires that the **Simplified Disaster Recovery** option on the backup selections has a status of **ON**. This status means that all components that are necessary for virtualization are selected. Backup Exec automatically selects the necessary components when you add a stage to convert to a virtual machine.

To add a conversion to virtual machine stage to a backup job

- 1 Create a backup job, or edit an existing job.
- 2 In the **Backup** box, click **Add Stage**.
- 3 Select **Convert to Virtual Machine** to set up a conversion to run after the backup job completes, or select **Convert to Virtual Simultaneously With Backups** to run the conversion at the same time as the backup job.
- 4 In the **Conversion to Virtual** box, click **Edit**.

- 5 If you selected the **Convert to Virtual Machine** option in step 3, do the following. If you selected **Convert to Virtual Simultaneously With Backups** in step 3, skip to step 6.
 - In the left pane, select **Schedule** to schedule the conversion, and then indicate if you want to schedule the job or run it immediately after the backup job completes.
 - (Optional) In the left pane, select **Notification** if you want to notify a recipient when the job completes.
- 6 In the left pane, select **Conversion Settings** to set the options for the conversion.
- 7 In the **Convert for** field, select either **Hyper-V** or **VMware ESX/vCenter server**.
- 8 Configure the conversion options:

If you selected Hyper-V in step 7

Do the following to configure the conversion options:

- Click the arrow in the **Hyper-V server name** field, select the name of the server where you want to create the virtual machine, and then click **Add**.
- In the **Destination drive or path** field, enter the location on the physical computer where the virtual disks should be created. Enter a drive letter and path.
- If you want to change the default virtual machine name, type the new name in the **Virtual machine name** field.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of Hyper-V Integration Components ISO image** field, enter the location of your Hyper-V Integration Components ISO image. The ISO image is needed to make the virtual machine bootable.
- If you want to change the CPU count or amount of physical RAM for the destination virtual machine, on the **Server configuration** tab, enter the new amounts in the **Desintation virtual machine** fields.
- If you want to change the disk type, the controller, or the virtual disk location for one of the disks, on the **Disk configuration** tab, click **Edit disk configuration**, and then enter the new information.

If you selected VMware ESX/vCenter
server in step 7

Do the following to configure the conversion options:

- Click the arrow in the **ESX/vCenter server name** field, and click the name of the server where you want to create the virtual machine.
- Click the arrow in the **Logon account** field, and then select the appropriate logon account for the server that you selected.
- Click **Select** next to the ESX/vCenter server name field. Backup Exec fills in the remaining information about the server.
- Browse to select the virtual machine folder and resource pool that are associated with the server that you selected.

Note: If you target an ESX server and select a resource pool, the newly created virtual machine is not added to the resource pool automatically. You can manually move the virtual machine into the appropriate resource pool after the conversion.

- If you want to change the default virtual machine name, type the new name in the **Virtual machine name** field.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of VMware Tools ISO image** field, enter the location of your VMware Tools ISO image. The path should be accessible with the default credentials. The path should also be local to the Backup Exec server. The ISO image is needed to make the virtual machine bootable.

Note: In a CASO environment, this path should be local to the managed Backup Exec server to which the job is targeted.

- If you want to change the CPU count or amount of physical RAM for the destination virtual machine, on the **Server configuration** tab, enter the new amounts in the **Desintation virtual machine** fields.
- If you want to change the disk type, the controller, or the virtual disk location for one of the disks, on the **Disk configuration** tab, click **Edit disk configuration**, and then enter the new information.

9 Click **OK** to save your selections.

10 On the **Backup Definition Properties** dialog box, edit the backup job properties, and then click **OK** to create the job.

See [“How conversion of physical computers to virtual machines works in Backup Exec”](#) on page 479.

Converting to a virtual machine from a point in time

A conversion to a virtual machine from a point in time converts existing backup sets from a backup job in which the Simplified Disaster Recovery option was enabled. The Simplified Disaster Recovery option enables all of the critical system components for a virtual machine conversion to be included in the backup job.

Note: The option **Convert to virtual machine from point-in-time** becomes available for selection only after you run at least one full backup that includes all critical system components.

The option to convert to a virtual machine from a point in time is useful in a disaster recovery situation in which you want to quickly recover a failed server. The backup sets contain all of the necessary components of the system. Additionally, you can select application data or user data to include in the conversion.

To convert to a virtual machine from a point in time

- 1 On the **Backup and Restore** tab, select the server that contains the backup sets you want to convert.
- 2 In the **Conversions** group, click **Convert to Virtual**, and then click **Convert to Virtual Machine from Point In Time**.
- 3 On the **Options** dialog box, in the **Selected Point-in-Time** box, click **Edit**.
- 4 Select the items that you want to include in the conversion, and then click **OK**.
- 5 On the **Options** dialog box, in the **Convert to Virtual** box, click **Edit**.
- 6 In the left pane, select **Schedule** to select when to run the conversion job:

To run the job immediately

Click **Run now**.

To schedule the conversion to run at a specific time

Click **Run on**, and then enter the date and time to run the job.

To create the conversion job without scheduling it

Select **Create without a schedule**. When you use this option, the job does not run at the time of creation and it does not have a recurring schedule. The job remains in a pending state until you choose to run it. You can use a third-party job automation or task scheduling tool to run the job later.

- 7 (Optional) In the left pane, select **Notification** if you want to notify a recipient when the job completes.
- 8 In the left pane, select **Conversion Settings** to set the options for the conversion.
- 9 In the **Point in time** field, select the point in time that you want to use for the conversion.
- 10 In the **Name** field, select all of the components for inclusion in the conversion.
- 11 In the **Application data or non-system user data** field, select additional data to include in the conversion.
- 12 Click **OK** to save your selections.
- 13 On the **Options** dialog box, click **OK**.

Creating a one-time conversion to a virtual machine

You can create a one-time conversion job to convert a running physical computer to a virtual machine without a separate backup job. A one-time conversion job can be scheduled to run at a later time, but it cannot be scheduled to run more than one time.

Only Full (Copy) backups are supported for this type of conversion. Incremental and differential backups are not supported for one-time conversions.

Note: Copy backups are not supported in Hyper-V 2008 or 2008 R2. A crash-consistent conversion is created instead.

To create a one-time conversion to a virtual machine

- 1 On the **Backup and Restore** tab, select the server that you want to convert to a virtual machine.
- 2 In the **Conversions** group, click **Convert to Virtual**, and then click **One-Time Convert to Virtual Machine**.
- 3 On the **One-Time Convert to Virtual Machine Properties** dialog box, in the **Convert to Virtual** box, click **Edit**.
- 4 In the left pane, select **Schedule** to select when to run the conversion job:

To run the job immediately

Click **Run now**.

To schedule the conversion to run at a specific time

Click **Run on**, and then enter the date and time to run the job.

To create the conversion job without scheduling it

Select **Create without a schedule**. When you use this option, the job does not run at the time of creation and it does not have a recurring schedule. The job remains in a pending state until you choose to run it. You can use a third-party job automation or task scheduling tool to run the job later.

- 5 (Optional) In the left pane, select **Notification** if you want to notify a recipient when the job completes.
- 6 In the left pane, select **Conversion Settings** to set the options for the conversion.
- 7 In the **Convert for** field, select either **Hyper-V** or **VMware ESX/vCenter server**.
- 8 Configure the conversion options:

If you selected Hyper-V in step 7

Do the following to configure the conversion options:

- Click the arrow in the **Hyper-V server name** field, select the name of the server where you want to create the virtual machine, and then click **Add**.
- In the **Destination drive or path** field, enter the location on the physical computer where the virtual disks should be created. Enter a drive letter and path.
- If you want to change the default virtual machine name, type the new name in the **Virtual machine name** field.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of Hyper-V Integration Components ISO image** field, enter the location of your Hyper-V Integration Components ISO image. The ISO image is needed to make the virtual machine bootable.
- If you want to change the CPU count or amount of physical RAM for the destination virtual machine, on the **Server configuration** tab, enter the new amounts in the **Desintation virtual machine** fields.
- If you want to change the disk type, the controller, or the virtual disk location for one of the disks, on the **Disk configuration** tab, click **Edit disk configuration**, and then enter the new information.

If you selected VMware ESX/vCenter
server in step 7

Do the following to configure the conversion options:

- Click the arrow in the **ESX/vCenter server name** field, and click the name of the server where you want to create the virtual machine.
- Click the arrow in the **Logon account** field, and then select the appropriate logon account for the server that you selected.
- Click **Select** next to the **ESX/vCenter server name** field. Backup Exec fills in the remaining information about the server.
- Browse to select the virtual machine folder and resource pool that are associated with the server that you selected.

Note: If you target an ESX server and select a resource pool, the newly created virtual machine is not added to the resource pool automatically. You can manually move the virtual machine into the appropriate resource pool after the conversion.

- If you want to change the default virtual machine name, type the new name in the **Virtual machine name** field.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of VMware Tools ISO image** field, enter the location of your VMware Tools ISO image. The path should be accessible with the default credentials. The path should also be local to the Backup Exec server. The ISO image is needed to make the virtual machine bootable.

Note: In a CASO environment, this path should be local to the managed Backup Exec server to which the job is targeted.

- If you want to change the CPU count or amount of physical RAM for the destination virtual machine, on the **Server configuration** tab, enter the new amounts in the **Desintation virtual machine** fields.
- If you want to change the disk type, the controller, or the virtual disk location for one of the disks, on the **Disk configuration** tab, click **Edit disk configuration**, and then enter the new information.

9 Click **OK** to save your selections.

10 On the **One-Time Convert to Virtual Machine Properties** dialog box, click **OK**.

See [“How conversion of physical computers to virtual machines works in Backup Exec”](#) on page 479.

Setting default options for conversion to virtual machine jobs

You can set default options for all conversion to virtual machine jobs. However, you can override the default options for individual jobs.

To set default options for conversion to virtual machine jobs

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Job Defaults**, and then select **Convert to Virtual**.
- 3 Select the default schedule options that you want conversion to virtual jobs to use:

To schedule conversion to virtual jobs to run at a specific time

Do the following:

- Select **According to schedule**.
- Click the arrow in the **Source** field to select either all backups or the most recent full backup as the source to initiate the conversion.
- Select **Recurrence**, and then select the arrow to set the recurrence pattern.

To create conversion to virtual jobs without scheduling them

Do the following:

- Select **According to schedule**.
- Click the arrow in the **Source** field to select either all backups or the most recent full backup as the source to initiate the conversion.
- Select **Create without a schedule**.
 When you use this option, the job does not run at the time of creation and it does not have a recurring schedule. The job remains in a pending state until you choose to run it. You can use a third-party job automation or task scheduling tool to run the job later.

To run the conversion immediately after the backup job completes

Select **Convert to virtual immediately after source task completes**.

- 4 (Optional) In the left pane, select **Notification**, and then select the recipients who should receive notifications about conversion to virtual jobs.
- 5 In the left pane, select **Conversion Settings**.
- 6 In the **Convert for** field, select either **Hyper-V** or **VMware ESX/vCenter server**.
- 7 Configure the conversion options:

If you selected Hyper-V in step 6

Do the following to configure the conversion options:

- Click the arrow in the **Hyper-V server name** field, select the name of the server where you want to create the virtual machines, and then click **Add**.
- In the **Destination drive or path** field, enter the location on the physical computer where the virtual disks should be created. Enter a drive letter and path.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of Hyper-V Integration Components ISO image** field, enter the location of your Hyper-V Integration Components ISO image. The ISO image is needed to make the virtual machine bootable.

If you selected VMware ESX/vCenter server in step 6

Do the following to configure the conversion options:

- Click the arrow in the **ESX/vCenter server name** field, and click the name of the server where you want to create the virtual machines.
- Click the arrow in the **Logon account** field, and then select the appropriate logon account for the server that you selected.
- If you want to enable Backup Exec to overwrite a virtual machine if a virtual machine with the same name already exists, verify that the **Overwrite the virtual machine if it already exists** option is selected. If this option is not selected and the virtual machine name already exists, then the job fails.
- In the **Full path of VMware Tools ISO image** field, enter the location of your VMware Tools ISO image. The path should be accessible with the default credentials. The path should also be local to the Backup Exec server. The ISO image is needed to make the virtual machine bootable.

Note: In a CASO environment, this path should be local to the managed Backup Exec server to which the job is targeted.

8 Click **OK**

Configuration and settings

This chapter includes the following topics:

- [Changing default backup job settings](#)
- [Configuring schedules for backup jobs](#)
- [Configuring storage options for backup jobs](#)
- [Configuring automatic test run jobs for backup jobs](#)
- [Configuring automatic verify operations for backup jobs](#)
- [Configuring delayed catalog operations to improve backup performance](#)
- [Configuring Advanced Open File options for backup jobs](#)
- [Configuring checkpoint restart](#)
- [Configuring Backup Exec to work with Symantec Endpoint Protection](#)
- [Configuring pre/post commands for backup jobs](#)
- [Configuring file and folder options for backup jobs](#)
- [Setting default schedule options for rule-based jobs and run now jobs](#)
- [Excluding dates from the backup schedule for all backups](#)
- [Removing dates from the list of excluded dates](#)
- [Exporting a list of dates that are excluded from all backups to another server](#)
- [Changing the default preferences](#)
- [Configuring the default setting for backing up multiple servers or applications](#)
- [Configuring database maintenance](#)

- Scheduling Backup Exec to check logon accounts
- Configuring Backup Exec to discover data to back up
- Adding discovered servers to the list of servers in Backup Exec
- Backup networks
- Changing network and security options for Backup Exec
- Using Backup Exec with firewalls
- Using encryption with Backup Exec
- Encryption key management
- Creating encryption keys
- Replacing an encryption key
- Deleting encryption keys
- Granular Recovery Technology
- Setting default Granular Recovery Technology (GRT) options
- DBA-initiated job templates
- Creating DBA-initiated job templates
- Editing DBA-initiated job templates
- Deleting DBA-initiated job templates
- Backup Exec logon accounts
- Starting and stopping Backup Exec services
- Changing the credentials for a service account
- Changing startup options for Backup Exec services
- Configuring audit logs
- Viewing the audit log
- Removing entries from the audit log
- Saving an audit log to a text file
- Copying configuration settings to another Backup Exec server
- Viewing server properties

Changing default backup job settings

Backup Exec is preconfigured with default settings for backup jobs. You can change the default settings for your backup jobs. When you create a backup job, the job inherits the default settings that you configure. You can override the default settings for backup jobs when you create them. Backup job settings include storage, security, and file system options for backup jobs, among other things.

You can set unique backup job defaults for the following types of backup jobs:

- Back Up to Deduplication Disk Storage Device
- Back Up to Disk
- Back Up to Tape
- Duplicate to Deduplication Disk Storage Device
- Duplicate to Tape
- Convert to Virtual

Note: Backup Exec displays only the types of backup jobs for which your system is configured. For example, if you do not have a tape drive, you do not see the Back Up to Tape option in the list of backup job types.

To change default backup job settings

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Job Defaults**.
- 2 Select the type of backup for which you want to set default options.

For example, if you want to set up the default options for backups to disk, select **Back Up to Disk**. The options that appear vary depending on what types of storage devices you have configured. Different default options can be configured for backup jobs to different types of storage.

- 3 In the left pane, select the setting for which you want to configure default options.

Schedule

Select this option to configure default settings for the time and frequency with which you want to run backup jobs.

See [“Configuring schedules for backup jobs”](#) on page 519.

Storage

Select this option to configure default settings for the storage device that you want to use for backup jobs.

See [“Configuring storage options for backup jobs”](#) on page 527.

Notification

Select this option to configure Backup Exec to notify specified recipients when backup jobs are completed.

Each type of backup job can be configured with different notification recipients. Backup Exec can notify people by email or text message.

See [“Notification options for jobs”](#) on page 305.

Test Run

Select this option to configure a test job that automatically tests storage capacity, credentials, and media integrity.

The test job can help you determine if there are any problems that might keep backup jobs from completing successfully.

See [“Configuring automatic test run jobs for backup jobs”](#) on page 533.

Verify

Select this option to create a job that automatically verifies whether all of the data was successfully backed up when jobs are completed.

A verify job can also help you determine whether the media you use is defective.

See [“Configuring automatic verify operations for backup jobs”](#) on page 534.

Delayed Catalog

Select this option to configure delayed catalog operations for any GRT-enabled jobs. You can choose to run a delayed catalog operation, schedule the catalog operation for another time, or choose to not run a delayed catalog operation.

See [“Configuring delayed catalog operations to improve backup performance”](#) on page 536.

Advanced Open File

Select this option to configure the snapshot settings that Backup Exec uses to process backup jobs. Snapshot technology lets Backup Exec capture any files that are open when a backup job runs.

You can also enable checkpoint restart, which lets you resume interrupted backup jobs.

See [“Configuring Advanced Open File options for backup jobs”](#) on page 540.

See [“Configuring checkpoint restart”](#) on page 543.

Advanced Disk-based Backup

Select this option to configure off-host backup processing for backup jobs.

See [“Setting default backup options for the Advanced Disk-based Backup Option”](#) on page 1159.

Security

Select this option to configure Backup Exec to work with Symantec Endpoint Protection. You can configure backup jobs to run automatically when the Symantec ThreatCon reaches a specified level. Integrating your backup jobs with Symantec Endpoint Protection helps make sure that your vital data is safely backed up as soon as global threats are detected.

See [“Configuring Backup Exec to work with Symantec Endpoint Protection”](#) on page 545.

Pre/Post Commands

Select this option to configure any commands that you want to run either before backup jobs begin or after backup jobs are completed.

See [“Configuring pre/post commands for backup jobs”](#) on page 547.

Files and Folders

Select this option to configure how Backup Exec processes file system attributes such as junction points and symbolic links.

See [“Configuring file and folder options for backup jobs”](#) on page 552.

Enterprise Vault	<p>Select this option to configure Enterprise Vault options for backup jobs.</p> <p>See “Enterprise Vault backup options” on page 1027.</p>
Linux and Macintosh	<p>Select this option to configure options for any Linux or Macintosh computers that are included in backup jobs.</p> <p>See “Linux backup options” on page 1205.</p> <p>See “Default backup job options for Macintosh systems” on page 1235.</p>
Lotus Domino	<p>Select this option to configure options for any Lotus Domino data that is included in backup jobs.</p> <p>See “Lotus Domino Agent backup options” on page 1070.</p>
Microsoft Active Directory	<p>Select this option to configure options for any Microsoft Active Directory data that is included in backup jobs.</p> <p>See “Editing options for Active Directory and ADAM/AD LDS backup jobs” on page 1089.</p>
Microsoft Exchange	<p>Select this option to configure options for any Microsoft Exchange data that is included in backup jobs.</p> <p>See “Setting default backup options for Exchange Server” on page 957.</p>
Virtual Machines	<p>Select this option to configure options for any virtual machines that are included in backup jobs.</p> <p>See “Setting default backup options for virtual machines” on page 843.</p> <p>See “Setting default backup options for Hyper-V” on page 878.</p>
Microsoft SharePoint	<p>Select this option to configure options for any Microsoft SharePoint data that is included in backup jobs.</p> <p>See “Setting default backup options for SharePoint” on page 978.</p>
Microsoft SQL	<p>Select this option to configure options for any Microsoft SQL data that is included in backup jobs.</p> <p>See “Setting default backup options for SQL Server” on page 911.</p>

NDMP	<p>Select this option to configure options for any NDMP data that is included in backup jobs.</p> <p>See “NDMP backup options for NDMP servers” on page 1174.</p>
Oracle	<p>Select this option to configure options for any Oracle data that is included in backup jobs.</p> <p>See “Oracle backup options” on page 1011.</p>
Archive Option	<p>Select this option to configure how Backup Exec works with any archiving stages in backup definitions.</p> <p>See “Archive Option backup options” on page 1380.</p>
Exclusions	<p>Select this option to exclude specific files or specific types of files from backup jobs.</p> <p>See “Excluding files from backups” on page 178.</p>

- 4 Select the appropriate options.
- 5 When you are finished configuring default options, click **OK**.

Configuring schedules for backup jobs

Backup Exec lets you configure the time and the frequency for which you want to run jobs. You can run jobs immediately, once on a specific day and time, or more than once according to a schedule. Backup Exec lets you use minutes, hours, days, weeks, months, or years as measurements of time to create a recurring pattern for the schedule. Alternatively, you can select specific days of the month to create a recurring schedule on which jobs should run.

See [“How job scheduling works in Backup Exec”](#) on page 215.

You can configure default options for schedules, which all your jobs inherit when you create them. Or you can override the default schedule settings when you create jobs.

To configure schedules for jobs

- 1 Do one of the following:

- | | |
|--|---|
| To configure default schedule settings for all backup jobs | <ul style="list-style-type: none"> ■ Click the Backup Exec button, and then select Configuration and Settings. ■ Select Job Defaults, and then select the type of backup for which you want to configure schedule settings. |
|--|---|

- | | |
|--|---|
| To configure a schedule for specific backup jobs | <ul style="list-style-type: none"> ■ Create a new backup definition or edit an existing backup definition. ■ In the Backup box, click Edit. |
|--|---|

2 In the left pane, click **Schedule**.

3 In the **Job template name** field, type the name of the job template for which you want to configure a schedule.

Job templates are the collection of settings that Backup Exec uses to create jobs. Backup job settings can include scheduling options, storage device options, or backup methods for selected types of data, for example. When you want to run a backup, Backup Exec combines the job template with the backup selections to create a backup job that runs according to the options that you specified.

The job template name that you enter in this field is used to create the job name.

4 In the **Job name** field, type the name of the job for which you want to configure a schedule.

The unique job name helps you to identify backup jobs in Backup Exec. This is only available to set per job, not as a default backup job setting.

5 For each job for which you want to configure a schedule, do one of the following:

To configure recurring jobs	Complete the following steps:
-----------------------------	-------------------------------

- Select **Recurrence**.
- Proceed to step 6 to configure the recurrence pattern.

To configure jobs to run immediately without any recurrences

Complete the following steps:

- Select **Run now with no recurring schedule**.
- Proceed to step 9.

Note: This option is only available for full backup jobs.

To configure a job without a schedule

Complete the following steps:

- Select **Create without a schedule**.
- Proceed to step 10.

When you use this option, the jobs do not run at the time of creation and they do not have a recurring schedule. The jobs remain unscheduled until you choose to run them. You can use a third-party job automation or task scheduling tool to run the jobs later.

You can run unscheduled backup jobs later by using the **Run Next Backup Now** option or you can manually run the job by using the **Run Now** option.

If you use this option to create a job, you cannot place the job on hold. You cannot place jobs on hold unless they are scheduled.

6 To configure the recurrence pattern, complete the following options:

Hours

Select this option to create a recurrence pattern that is measured in hours or minutes.

When you select **Hours**, you can configure the following options:

- **Every X hour/minute**

Indicate the number of hours or minutes between the start time of a job and the start time of the next job instance.

- You must choose between the following options:

- **From**

Designate the starting time for a job to run.

- **Between**

Restrict the job to certain hours and days. For example, if you only want the job to run during business hours, you can select 9:00 AM to 5:00 PM on Monday, Tuesday, Wednesday, Thursday, and Friday.

The start time and end time can span a maximum of 24 hours. However, they can cross over midnight into the following day.

Days

Select this option to create a recurrence pattern that is measured in days.

When you select **Days**, you must choose between the following options:

- **Every X day**

Indicate the number of days between the start time of a job and the start time of the next job instance.

- **Every weekday**

Select this option to enable the job to run on Mondays, Tuesdays, Wednesdays, Thursdays, and Fridays.

Weeks

Select this option to create a recurrence pattern that is measured in weeks.

When you create a recurrence pattern that is measured in weeks, you must configure the **Every X week on** field. The **Every X week on** specifies the number of weeks between the start time of a job and the start time of the next job instance. It also specifies the days of the week on which the job should run.

Months

Select this option to create a recurrence pattern that is measured in months.

When you select **Months**, you must choose between the following options:

- **Day X of every X month**

Specify the day on which the job should run. Then indicate the number of months between the start time of a job and the start time of the next job instance.

- **Every X X of every X month**

Specify the day on which the job should run. Then indicate the number of months between the start time of a job and the start time of the next job instance.

- **Selected days of the month**

Specify the weeks and days of the month on which Backup Exec runs the job. You select the days and weeks on a grid. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current week and day of the month. If you create the job on the third Monday of the month, the default setting is for the job to run once a month on the third Monday.

You can change the default or select additional days on which the job should run. For example, if you want the job to run on the last Friday of every month, select the check box for Friday in the last row of the grid. Any additional days that you select are added to the monthly recurrence pattern.

- **Selected dates of the month**

Specify the dates of the month on which Backup Exec runs the job. The recurrence pattern that you select repeats itself every month.

The default setting is for the job to run every month on the current date of the month. If you create the job on the 15th, the default setting is for the job to run once a month on the 15th.

You can change the default or select additional days on which the job should run. For example, if you want the job to run on the 1st and 15th of every month, select only those dates in the calendar. Any additional days that you select are added to the monthly recurrence pattern.

If you select the 31st, the job runs on the last day of the month in months that do not have 31 days. For example, if you configure the job to run on the 31st, in September the job runs on the 30th instead.

Years

Select this option to create a recurrence pattern that is measured in years.

When you select **Years**, you can configure the following options:

- **Every X year**

Specify the number of years between the start time of a job and the start time of the next job instance.

- You must choose between the following options:

- **On X**

Specify the date on which Backup Exec runs the job. The date that you select in this field corresponds to the number of years that you selected in the **Every X year** field. So if you selected to run the job every 2 years and you selected June 28th in this field, the job runs every 2 years on June 28th.

- **On the X of X**

Specify the day and month on which Backup Exec runs the job. The date that you select in this field corresponds to the number of years that you selected in the **Every X year** field. So if you selected to run the job every 2 years and you selected the fourth Thursday of June in this field, the job runs every 2 years on the fourth Thursday of June.

at

Enter the starting time for the first job in the recurrence pattern.

Starting on

Enter the date on which you want the recurrence pattern to begin.

The date that you enter in this field is the date on which the schedule goes into effect. You can select any date in the past or the future. If you select a date that occurred in the past, Backup Exec calculates the date of the next upcoming job and begins running recurring jobs on that date.

Calendar

Click this option to view all scheduled backup jobs on a calendar to check for scheduling conflicts.

- | | |
|---|--|
| Run initial full backup now in addition to the selected schedule | Select this option to run the initial full backup as soon as the job is created without affecting the schedule of future jobs. |
| Reschedule the job if it does not start X hours after its scheduled start time | <p>Specify the amount of time past the job's scheduled start time at which you want Backup Exec to change the job completion status to Missed. The job is rescheduled to run based on the time window that you configured.</p> <p>See "List of job statuses in Backup Exec" on page 277.</p> |
| Cancel the job if it is running X hours after its scheduled start time | <p>Specify the amount of time after the job's scheduled start time at which you want to cancel the job if it is still running. Backup Exec changes the job completion status to Canceled, timed out.</p> <p>See "List of job statuses in Backup Exec" on page 277.</p> |

- 7 To configure specific dates to include with the recurring job schedule, select the **Include/exclude dates** tab, and then complete the following steps:

- Click **Include dates**.
- Select the dates that you want to include with the recurring job schedule.
- Click **OK**.

The job runs on all of the dates that you select using this option, in addition to the dates that are part of its normal schedule recurrence. The job resumes its normal schedule on the next day that it is scheduled after an include date.

See ["Including a specific date in the schedule for a backup job"](#) on page 216.

- 8 To configure specific dates to exclude from the job schedule, select the **Include/exclude dates** tab, and then complete the following steps:

- Click **Exclude dates**.
- Select the dates that you want to exclude from the recurring job schedule.
- Click **OK**.

The job does not run on any of the dates that you select using this option. It resumes its normal schedule on the next day that a job is scheduled after an excluded date.

See ["Preventing backup jobs from running on a specific date"](#) on page 218.

- 9 To submit jobs with an on-hold status, select **Submit job on hold**.

You should select this option if you want to submit the job, but you do not want the job to run until a later date. The job runs later when you change the job's hold status.

- 10 Click **OK**.

See [“Changing default backup job settings”](#) on page 515.

See [“Backing up data”](#) on page 158.

Configuring storage options for backup jobs

Use the storage options to select the storage and media set on which you want the backup job to run. You can configure different storage devices for each backup job. For example, you may select disk storage for a full backup and a storage pool for an incremental backup in the same backup definition.

You can configure storage options as default settings for all backup jobs. If you do not want to use the default storage options for a particular backup job, you can override the default setting when you create the backup job. You do not have to configure default storage options for all backup jobs, however. If you want to configure different storage options for specific backup jobs, you can configure the storage options when you create those backup jobs.

To configure storage options for backup jobs

- 1 Do one of the following:

To configure default storage options for all backup jobs

Complete the following steps:

- Click the Backup Exec button, and then select **Configuration and Settings**.
- Select **Job Defaults**, and then select the type of backup for which you want to configure storage options.

To configure storage options for specific backup jobs

Complete the following steps:

- Create a new backup definition or edit an existing backup definition.
- In the **Backup** box, click **Edit**.

- 2 In the left pane, click **Storage**.

3 Complete the following options as necessary:

Note: Some of these options display only in Central Admin Server Option (CASO) environments.

Priority	<p>Select the priority of access to the storage devices for backup jobs.</p> <p>This is only available to set per job, not as a default backup job setting.</p> <p>See “Changing the priority for a scheduled job” on page 262.</p>
Backup Exec server or Backup Exec server pool	<p>Select whether you want a job to run on devices on a specific managed Backup Exec server or on devices that are on a group of managed Backup Exec servers.</p> <p>This option displays only if you have the Central Admin Server Option installed. This option is an additional filter that lets you control where certain jobs are delegated. For example, to always run backups of Exchange databases only on the devices that are attached to managed Backup Exec servers in a pool named Exchange Backups, select this option. Then select the Exchange Backups Backup Exec server pool.</p>
Storage	<p>Select the storage device to which you want to send backup data.</p> <p>See “Creating storage device pools” on page 416.</p> <p>See “About the Any Virtual Disk Storage device pool in the Storage Provisioning Option” on page 1278.</p> <p>See “About the Remote Media Agent for Linux ” on page 1241.</p> <p>See “Features and types of disk-based storage and network-based storage” on page 313.</p>

- 4 If you selected to configure an OpenStorage device or a deduplication disk storage device in the **Storage** field, select from the following options:

Enable the remote computer to directly access the storage device and to perform client-side deduplication, if it is supported

Select this option to enable a remote computer to send data directly to an OpenStorage device or a deduplication disk storage device, and to perform client-side deduplication if the device supports it. The Backup Exec server is bypassed, which leaves the Backup Exec server free to perform other operations. If client-side deduplication cannot be performed, then either Backup Exec server deduplication or Appliance deduplication is performed.

This option appears if the Deduplication Option is installed and an OpenStorage device or a deduplication disk storage device is selected in the **Storage** field.

See [“How to use client-side deduplication”](#) on page 821.

Enable the remote computer to access the storage device through the Backup Exec server and to perform Backup Exec server-side deduplication, if it is supported

Select this option to enable a remote computer to send data through the Backup Exec server to an OpenStorage device or a deduplication disk storage device, and to perform Backup Exec server-side deduplication if it is supported. If the Backup Exec server does not support deduplication, the data is deduplicated on an intelligent disk device, such as Symantec PureDisk or a device from a third-party vendor.

This option appears if the Deduplication Option is installed and an OpenStorage device or a deduplication disk storage device is selected in the **Storage** field.

See [“About the Deduplication Option”](#) on page 795.

- 5 In the **Keep for** field, enter the amount of time for which you want to keep the backup sets or job history.
- 6 If you selected to configure a tape device in the **Storage** field, complete the following options as necessary:

Media set

Select the media set to use for the backup job. The media set specifies the overwrite protection period and the append period for the backup data on the media.

If you want to create a new media set for this backup job, click the icon to the right of the media set drop-down menu.

This option is available only if you selected a tape device in the **Storage** field.

See [“Default media sets”](#) on page 379.

Overwrite media

Select this option to place the backed up data on overwritable media. Ensure that appropriate media is in the storage device that you select.

Appropriate media for an overwrite job include the following:

- Scratch media
- Media for which the overwrite protection period has expired

Allocated or imported media may also be overwritten depending on the media overwrite protection level that is set.

Depending on your configuration, overwritable media is selected from scratch media or recyclable media.

If the media in the storage device is not overwritable, an alert appears to prompt you to insert overwritable media.

This option is available only if you selected a tape device in the **Storage** field.

See [“Managing tapes”](#) on page 378.

See [“Media overwrite protection levels for tape media”](#) on page 394.

See [“How Backup Exec searches for overwritable media in tape drives ”](#) on page 395.

Append to media, overwrite if no appendable media is available	<p>Select this option to append the backed up data to the specified media set if an appendable media is available. Otherwise, Backup Exec searches for an overwritable media and adds it to the media set.</p> <p>If the append operation fills a media, the backup job continues on an overwritable media. If the media in the storage device is not overwritable, an alert prompts you to insert overwritable media.</p> <p>This option is available only if you selected a tape device in the Storage field.</p>
Append to media, terminate job if no appendable media is available	<p>Select this option to append the backed up data to the specified media set if an appendable media is available. Otherwise, Backup Exec terminates the job.</p> <p>This option is available only if you selected a tape device in the Storage field.</p>
Eject the media after job completes	<p>Select this option to eject the media from the drive or slot when the operation completes. You can also schedule a job to eject media.</p> <p>This option is available only if you selected a tape device in the Storage field.</p> <p>See “Ejecting media from a disk cartridge or tape drive” on page 458.</p>
Retension media before backup	<p>Select this option to run the tape in the drive from beginning to end at a fast speed. Retensioning helps the tape wind evenly and run more smoothly past the tape drive heads. This option is available only if you select a tape drive that supports retensioning.</p>
Use Write once, read many (WORM) media	<p>Select this option to use WORM (write once, read many) media for this backup job. Backup Exec confirms that the destination device is or contains a WORM-compatible drive, and that the WORM media is available in the drive. If WORM media or a WORM-compatible drive is not found, an alert is sent.</p> <p>See “How WORM media is used in Backup Exec” on page 401.</p>

Export media to vault after job completes

Select this option to logically move the media from the robotic library to the specified media vault.

This operation moves media from robotic library slots into a portal. An alert reminds you to remove the media from the portal or from a slot. If a job requires multiple media, the export media operation starts after the backup job completes, not after each media is filled.

This option is available only if you selected a tape device in the **Storage** field.

See [“Default media vaults”](#) on page 402.

7 In the **Compression** field, select from the following options:

- None**

Select this option to copy the data to the media in its original form (uncompressed). Using some form of data compression can help expedite backups and preserve storage space.

Hardware data compression should not be used in environments where storage devices that support hardware compression are used interchangeably with devices that do not have that functionality. In this situation, hardware compression is automatically disabled. You can manually turn on hardware compression on the drives that support it, but this results in media inconsistency. If the drive that supports hardware compression fails, the compressed media cannot be restored with the non-compression drive.
- Software**

Select this option to use STAC software data compression, which compresses the data before it is sent to the storage device.
- Hardware (if available, otherwise none)**

Select this option to use hardware data compression if the storage device supports it. If the drive does not feature data compression, the data is backed up uncompressed.
- Hardware (if available, otherwise software)**

Select this option to use hardware data compression if the storage device supports it. If the drive does not feature hardware data compression, STAC software compression is used.

8 To configure encryption, complete the following options:

Encryption type	Select the type of encryption that you want to use, if any. See “Using encryption with Backup Exec” on page 583.
Encryption key	Select the encryption key that you want to use, if you selected to use encryption.
Manage Keys	Click this option to manage your encryption keys. You can delete or replace existing encryption keys. You can also create a new encryption key. This option is available only if you select an encryption type. See “Encryption key management” on page 586.

9 Click **OK**.

See [“Changing default backup job settings”](#) on page 515.

See [“Backing up data”](#) on page 158.

Configuring automatic test run jobs for backup jobs

Test run jobs attempt to determine if a scheduled backup job could possibly fail when you run it. When you run a test job, no data is backed up. Instead, Backup Exec checks your storage capacity, credentials, and media to find potential errors. If there is an error, the job continues to run until it is completed. The error appears in the job log. You can also configure Backup Exec to send a notification to a designated recipient.

During a test run job, the following things may cause a job to fail:

- Logon credentials are incorrect.
- Storage capacity is not sufficient.
- Tape cartridge media or disk cartridge media is not available.
- Overwritable media is not available for an overwrite job.
- Appendable media is not available for an append job.

A test run job checks the media capacity that is available for the selected job. However, you can check if there is enough available media for multiple test run jobs in the Test Run Results Report.

See [“Test Run Results report”](#) on page 683.

You can manually run a test run job at any time.

See [“Running a test run job manually”](#) on page 226.

You can also configure test run jobs to run automatically before your scheduled backup jobs as a default setting. If you do not want to run a test run job for a particular backup job, you can override the default setting when you create the backup job. You do not have to enable test run jobs as a default for all backup jobs, however. If you want to run test run jobs only for specific backup jobs, you can configure the test run job when you create those backup jobs.

To configure automatic test run jobs for backup jobs

1 Do one of the following:

To enable test run jobs as a default for all backup jobs

Complete the following steps:

- Click the Backup Exec button, and then select **Configuration and Settings**.
- Select **Job Defaults**, and then select the type of backup for which you want to configure test run jobs.

To enable a test run job for specific backup jobs

Complete the following steps:

- Create a new backup definition or edit an existing backup definition.
- In the **Backup** box, click **Edit**.

2 In the left pane, click **Test Run.**

3 Select **Enable test run.**

4 Click **OK.**

See [“Changing default backup job settings”](#) on page 515.

See [“Backing up data”](#) on page 158.

Configuring automatic verify operations for backup jobs

Backup Exec can perform a verify operation to make sure that the media can be read after a backup job has been completed. Symantec recommends that you verify all backed up data to ensure the integrity of the collection of data and the media on which it resides.

You can manually run a verify operation on a backup set or a job history at any time. You can verify backup sets if you want to verify only the data that was backed up in a specific backup job instance. If you want to verify a backup definition and all of its dependent backup sets, you can verify a job history. For example, if you want to verify a backup definition that used incremental backups, Backup Exec verifies all incrementals dating back to, and including, the last full backup.

See [“Verifying backed up data manually”](#) on page 227.

By default, Backup Exec automatically verifies backed up data at the end of a backup job. However, you can also schedule the automatic verify operation to take place at a later time or disable the verify operation altogether. You can change Backup Exec's verify options as part of the default backup settings or for individual backup jobs.

To configure automatic verify operations for backup jobs

1 Do one of the following:

To configure automatic verify operations for all backup jobs Complete the following steps:

- Click the Backup Exec button, and then select **Configuration and Settings**.
- Select **Job Defaults**, and then select the type of backup for which you want to configure verify operations.

To configure automatic verify operations for specific backup jobs Complete the following steps:

- Create a new backup definition or edit an existing backup definition.
- In the **Backup** box, click **Edit**.

2 In the left pane, click **Verify**.

3 Complete the following options:

At the end of the job	Select this option to run a verify operation automatically when the backup job is completed.
After job finishes, as a separate job	<p>Select this option to create a verify operation and schedule it to run as a separate job when the backup job is completed.</p> <p>You can use the Edit option to configure options for the separate verify job.</p>
As a separate scheduled job	<p>Select this option to create a verify operation and schedule it to run as a separate job at a later time.</p> <p>You can use the Edit option to configure options for the separate verify job.</p>
Do not verify data for this job	Select this option to disable the verify operation for the backup job.

4 Click **OK**.

See [“Changing default backup job settings”](#) on page 515.

See [“Backing up data”](#) on page 158.

Configuring delayed catalog operations to improve backup performance

When you back up data, Backup Exec creates a catalog that contains information about the backup sets and about the storage device on which the backup sets are stored.

See [“How Backup Exec catalogs work”](#) on page 248.

Backup jobs that are enabled for Granular Recovery Technology (GRT) require more time to catalog because of the amount of granular information that they contain. For GRT-enabled backup jobs, the catalog operation is delayed and run as a separate operation to have less of an effect on your backup window. Because the catalog operation runs separately from the backup job, it does not prevent another scheduled backup job from starting on time.

Note: The delayed catalog feature is not supported for backup to tape jobs. If you create a GRT-enabled backup to tape job for Microsoft Exchange, Microsoft SharePoint, Microsoft Hyper-V, or VMware data, the catalog operation runs as part of the backup job.

See [“Granular Recovery Technology”](#) on page 590.

When you enable GRT for Microsoft Exchange, Microsoft SharePoint, Microsoft Hyper-V, or VMware backups, the delayed catalog operation runs after the backup job by default. For Exchange and SharePoint backups, the delayed catalog operation runs immediately after all full backups. It runs once every 24 hours for all incremental backups and differential backups, even if you schedule more than one GRT-enabled job to run in the 24-hour period.

For Hyper-V and VMware backups, the delayed catalog operation runs immediately after all full, incremental, and differential backups by default.

You can also configure the delayed catalog operation to run on a schedule if you do not want it to run immediately after the backup job. The catalog operation can be time consuming. It requires access to the storage device that is used for the backup. You may want to schedule the delayed catalog operation to run outside of your backup window so that it does not interfere with your system resources. If you schedule the delayed catalog operation, it runs only for the most recent backup set since the last catalog operation. In this situation, only the most recent backup set since the last catalog operation can be used for granular recovery.

For example, if you schedule incremental backups to run every 11 hours and schedule the delayed catalog operation to run at midnight, you would have the following backup sets:

- Full (11:00 A.M.)
- Incremental 1 (10:00 P.M.)
- Catalog 1 (Midnight) This job catalogs Incremental 1.
- Incremental 2 (9:00 A.M.)
- Incremental 3 (8:00 P.M.)
- Catalog 2 (Midnight). This job catalogs Incremental 3. Incremental 2 is not cataloged.
- Incremental 4 (7:00 A.M.)
- Incremental 5 (6:00 P.M.)
- Catalog 3 (Midnight) This job catalogs Incremental 5. Incremental 4 is not cataloged.

- Incremental 6 (5:00 A.M.) This backup is not cataloged.

In the example, you can perform a granular recovery from Incremental 5, Incremental 3, or Incremental 1. You cannot perform a granular recovery from Incremental 2, Incremental 4, or Incremental 6. However, you can perform a full restore from any of the backup sets.

You can also choose to not run a delayed catalog operation. If you do not run a delayed catalog operation for GRT-enabled Exchange and SharePoint backups, the granular information that is backed up is not cataloged for granular restore. You will not be able to search the backup sets for granular data. However, you can browse the backup sets. If you decide to restore granular data from the backup sets, Backup Exec reads and displays the granular information in the backup sets as you browse for items to restore. It takes longer to restore granular data if you do not run the delayed catalog operation.

If you choose to not run the delayed catalog operation for GRT-enabled Hyper-V and VMware backups, the catalog operation runs as part of the backup job instead of delaying it to a later time.

You can configure delayed catalog options as default settings for all GRT-enabled backup jobs. If the default settings are not appropriate for a particular job, you can override them when you create the job. You do not have to create default settings for delayed catalogs, however. If you want to configure delayed catalogs only for specific jobs, you can configure the settings when you create those jobs.

To configure a delayed catalog operation

- 1 Do one of the following:

To configure delayed catalog options for all backup jobs Complete the following steps:

- Click the Backup Exec button, and then select **Configuration and Settings**.
- Select **Job Defaults**, and then select the type of backup for which you want to configure delayed catalog settings.

To configure delayed catalog options for specific backup jobs Complete the following steps:

- Create a new backup definition or edit an existing backup definition.
- In the **Backup** box, click **Edit**.

- 2 In the left pane, click **Delayed Catalog**.

3 Select one of the following options:

Do not run a delayed catalog operation

Select this option if you do not want to run a delayed catalog operation for GRT-enabled backup jobs.

If you select this option when you back up Microsoft Exchange or Microsoft SharePoint data, the granular information that is backed up is not cataloged for granular restore. You will not be able to search the backup sets for granular data. However, you can browse the backup sets. If you decide to restore granular data from the backup sets, Backup Exec browses the backup sets for granular data as you browse for items to restore. If you select this option, it increases the amount of time that it takes to restore granular Exchange or SharePoint data.

If you select this option when you back up a virtual machine, the catalog operation runs as part of the backup job. Running the catalog operation as part of the backup job increases the amount of time that it takes to complete the backup job and may delay the start of any pending backup jobs.

Immediately after a backup job finishes, as a separate job

Select this option to run the delayed catalog operation immediately after a backup job finishes. The catalog operation runs as a separate job. This option is the default setting for GRT-enabled backup jobs.

For Exchange and SharePoint backups, the delayed catalog operation runs immediately after all full backups. It runs once every 24 hours for all incremental backups and differential backups.

For Hyper-V and VMware backups, the delayed catalog operation runs immediately after all full, incremental, and differential backups.

As a separate job, according to this schedule

Select this option to run the delayed catalog operation as a separate, scheduled job. Then select the start time and the days of the week on which you want the delayed catalog operation to run.

If you schedule the delayed catalog operation, it runs only for the most recent backup set since the last catalog operation. In this situation, only the most recent backup set since the last catalog operation can be used for granular recovery.

Note: You cannot browse backup sets for granular data or restore granular items from the backup sets until the data has been cataloged. However, you can perform a full restore from the backup sets.

4 Click **OK**.

See [“Changing default backup job settings”](#) on page 515.

See [“Backing up data”](#) on page 158.

Configuring Advanced Open File options for backup jobs

Backup Exec's Advanced Open File feature lets you use snapshot technology to capture any files that are open when a backup runs. You can configure Advanced Open File options as default settings for all backup jobs. If the default settings are not appropriate for a particular job, you can override them when you create the job. You do not have to create default settings for Advanced Open File options, however. If you want to use Advanced Open File options only for specific jobs, you can configure the settings when you create those jobs.

To configure Advanced Open File options for backup jobs

1 Do one of the following:

To configure default Advanced Open File options for all backup jobs

Complete the following steps:

- Click the Backup Exec button, and then select **Configuration and Settings**.
- Select **Job Defaults**, and then select the type of backup for which you want to configure Advanced Open File options.

To configure Advanced Open File options for specific backup jobs Complete the following steps:

- Create a new backup definition or edit an existing backup definition.
- In the **Backup** box, click **Edit**.

2 In the left pane, click **Advanced Open File**.

3 Complete the following options:

Use snapshot technology Select this option to enable the use of snapshot technology for backup jobs.

Snapshot provider Select one of the following snapshot providers for jobs:

- **Automatic** - Allow VSS to select the snapshot provider.
 Select this option to enable VSS to select the best provider for the selected volume. The order in which a snapshot provider is selected is hardware provider and then the system provider.
- **System** - Use Microsoft Software Shadow Copy Provider.
- **Hardware** - Use technology provided by the hardware manufacturer.

If you select multiple volumes, you must use the same type of provider to snap all of the volumes. You can snap multiple volumes with the same provider or you can use multiple providers, but you cannot use system and hardware providers as part of the same snapshot.

Process logical volumes for backup one at a time

Select this option to enable the backup of multiple volumes in one job, with only one logical volume being snapped at a time. To ensure database integrity, or if a volume contains mount points, multiple volumes may need to be snapped one at a time. A volume with mount points to other volumes is considered a logical volume for snapshot purposes. Therefore, that volume and the mount point volumes are snapped together simultaneously.

After the logical volume is snapped and backed up, the snapshot is detected before the next logical volume is snapped. This option increases the ability to meet the minimum quiet time that is needed to complete a snapshot.

A logical volume can comprise multiple physical volumes. A single logical volume can encompass all of the volumes on which databases reside.

If this option is not selected, then a snapshot for all volumes in the backup job is created simultaneously. All volumes must meet the minimum quiet time.

This option is only available for local volumes.

The Shadow Copy Components snapshots are created using VSS, which is reported in the job log.

Enable checkpoint restart

Select this option to enable the checkpoint restart option. Checkpoint restart enables Backup Exec to automatically restart a job that is interrupted. The job restarts from the point where it was interrupted instead of starting over at the beginning. Backup Exec waits two minutes after the job stops and then attempts to restart the interrupted job one time. If the job cannot be restarted automatically or if checkpoint restart is disabled, you must restart it manually. A manual restart starts the job at the beginning instead of at the point where the job was interrupted.

See [“Configuring checkpoint restart”](#) on page 543.

4 Click **OK**.

See [“Changing default backup job settings”](#) on page 515.

See [“Backing up data”](#) on page 158.

Configuring checkpoint restart

Checkpoint restart is a backup job setting that enables Backup Exec to automatically restart a job that is interrupted. The job restarts from the point where it was interrupted instead of starting over at the beginning. Backup Exec waits two minutes after the job stops and then attempts to restart the interrupted job. If the job cannot be restarted automatically or if checkpoint restart is disabled, you must restart it manually. A manual restart starts the job at the beginning instead of at the point where the job was interrupted.

Note: Checkpoint restart cannot restart a backup job until it has backed up at least 32 MB of data. If a backup job fails before it has backed up at least that much data, you must run it again manually.

You can enable checkpoint restart for any jobs that are configured with the storage option **Overwrite media**. However, if there is no available media to overwrite and the job is interrupted, it may be placed on hold until overwritable media is available. When the media is available, the job restarts from where it was interrupted.

Backup Exec automatically cancels any jobs that run for too long according to the schedule settings that you selected when you created the job. If Backup Exec automatically cancels a job, it is not eligible to be restarted. If you manually cancel a job, Backup Exec does not automatically try to restart it.

Note: If you use the Central Admin Server Option (CASO), any jobs that are restarted run on the same managed Backup Exec server on which the job failed. If the original Backup Exec server is not available, Backup Exec selects a different Backup Exec server on which to run the restarted job.

You can enable or disable checkpoint restart in the **Advanced Open File** options when you create backup jobs or in the backup job defaults.

See [“Configuring Advanced Open File options for backup jobs”](#) on page 540.

This topic includes the following information:

[Technologies supported for checkpoint restart](#)

[Things to consider before you use checkpoint restart](#)

[Changing the default checkpoint restart settings](#)

Technologies supported for checkpoint restart

Checkpoint restart is only supported for NTFS volumes. The only type of snapshot technology that is supported for checkpoint restart is VSS.

Checkpoint restart is not supported for the following:

- FAT volumes
- FAT32 volumes
- UNIX computers
- Cluster Shared Volumes (CSV)
- Application agents
- Incremental or differential backups
- Jobs that use catalogs to determine if a file has been backed up
See [“How Backup Exec determines if a file has been backed up”](#) on page 199.

Things to consider before you use checkpoint restart

You should consider the following things before you use checkpoint restart:

- If the failure occurs in the middle of an append job, the media is no longer appendable. The media is not appendable until it is erased or overwritten, or the retention period expires. When the restart occurs, Backup Exec uses new media. You should select an appropriate media overwrite protection level to ensure that the restart does not overwrite the media that was used before the job failure.
- If the failure occurs during a verify job or a database consistency check job, the job restarts at the beginning.
- Full backups that were interrupted and resumed from the point of failure do not display in the Simplified Disaster Recovery **Recover This Computer** Wizard. However, you can restore these backup sets manually after you make the initial recovery by using the **Recover This Computer** Wizard.
- You can enable the checkpoint restart option for a full backup job that uses the **Delete selected files and folders after successful backup** option. If the job fails and is resumed, the files are not deleted from the source volume after the backup completes.

Changing the default checkpoint restart settings

You can change the default checkpoint restart settings in the error-handling rules settings. You can specify the number of times that you want checkpoint restart to retry a failed job, the interval between restart attempts, and the final job disposition for any jobs that cannot be successfully restarted.

To change default checkpoint restart settings

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Error-Handling Rules**.
- 2 Select **Checkpoint Restart**, and then click **Edit**.
- 3 Select **Enable error-handling rule**.
- 4 Select **Retry job**.
- 5 Complete the following fields:

Maximum retries	Enter the maximum number of times that you want Backup Exec to retry a job that fails.
Retry interval	Enter the interval of time in minutes that you want Backup Exec to wait before it attempts to restart a job.

- 6 In the **Final job disposition** group box, select one of the following options:

Place job on hold until error condition has been manually cleared	Select this option to have Backup Exec place the job on hold if the job cannot be completed successfully after the maximum number of retries. The job remains on hold until the error condition has been manually cleared.
Reschedule job for its next scheduled service	Select this option to have Backup Exec reschedule the job for its next scheduled occurrence if the job cannot be completed successfully after the maximum number of retries.

- 7 (Optional) In the **Notes** field, type any additional notes about the error-handling rule.
- 8 Click **OK**.

Configuring Backup Exec to work with Symantec Endpoint Protection

You can use Symantec Endpoint Protection version 11.0 or later with Backup Exec to provide extra security when the threat of viruses or malware is high. Symantec Endpoint Protection uses a ThreatCon level to provide an overall view of global Internet security. Symantec's ThreatCon levels are based on a 1-4 rating system, with level 4 being the highest threat level.

You can find more information about Symantec ThreatCon levels at the following URL:

<http://www.symantec.com>

You can configure Backup Exec to automatically run a backup job when the ThreatCon reaches a level that you specify. You may want to configure special jobs for your most crucial data, for example. This strategy helps make sure that your vital data is safely backed up as soon as global threats are detected.

You should consider the types of jobs that you want to trigger automatically and the potential effect they can have on your system resources. The ThreatCon level is updated frequently and it can be raised at any time without warning. If you configure large or resource-intensive jobs to launch automatically, they may interfere with your normal business operations.

The Backup Exec server must be connected to the Internet to monitor the ThreatCon level. If the Backup Exec server is not connected to the Internet, backup jobs are not triggered when the ThreatCon level elevates.

See the *Administrator's Guide for Symantec Endpoint Protection* for more information about Symantec Endpoint Protection.

You can configure Backup Exec to work with Symantec Endpoint Protection as a default setting for all backup jobs. If you do not want to configure integration with Symantec Endpoint Protection for a particular job, you can override the default setting when you create the job. You do not have to create default settings for Symantec Endpoint Protection integration, however. If you want to use Symantec Endpoint Protection integration only for specific jobs, you can configure the settings when you create those jobs.

To configure Backup Exec to work with Symantec Endpoint Protection

1 Do one of the following:

To configure Symantec Endpoint Protection to work with all backup jobs

Complete the following steps:

- Click the Backup Exec button, and then select **Configuration and Settings**.
- Select **Job Defaults**, and then select the type of backup for which you want to configure Symantec Endpoint Protection.

To configure Symantec Endpoint Protection to work with specific backup jobs

Complete the following steps:

- Create a new backup definition or edit an existing backup definition.
- In the **Backup** box, click **Edit**.

2 In the left pane, click **Security**.

3 Complete the following options:

Run a backup job immediately when an elevated Symantec ThreatCon level is reached

Select this option to run automatic backups when the Symantec ThreatCon reaches the level that you specify in the **Symantec ThreatCon level** field. You must have Symantec Endpoint Protection 11.0 or later installed on the same computer as Backup Exec to use this feature.

Backup job

Select the type of backup job that you want to run when the Symantec ThreatCon level is raised.

Symantec ThreatCon level

Specify the ThreatCon level at which you want automatic backups to run.

You can find more information about Symantec ThreatCon levels at the following URL:

<http://www.symantec.com>

4 Click **OK**.

See [“Changing default backup job settings”](#) on page 515.

Configuring pre/post commands for backup jobs

You can configure commands that you want to run before or after all backup jobs. For example, you may want to create a pre-command to shut down a database before a backup job runs. You could also create a post-command that restarts the database after the job is completed.

Conditions that you can set for these commands include the following:

- Run the backup job only if the pre-command is successful
- Run the post-command only if the pre-command is successful
- Run the post-command even if the backup job fails
- Allow Backup Exec to check the return codes (or exit codes) of the pre- and post-commands to determine if the commands completed successfully

If the pre- or post-command returns an exit code of zero, Backup Exec considers the job to have completed successfully. Backup Exec considers any non-zero exit codes to mean that the job encountered an error.

If it is critical that the job does not run if the pre-command fails, configure Backup Exec to check the return codes. Backup Exec uses the return codes to determine if the pre-command failed or completed successfully.

For example, if a pre-command that shuts down a database before a backup is run fails, the database could be corrupted when the backup runs. In this situation, it is critical that the backup job does not run if the pre-command fails.

If Backup Exec is configured to check the return codes and the post-command returns a non-zero code, the job log reports that the post-command failed. You may have also selected to run the job only if the pre-command is successful. Even if both the pre-command and the job run successfully, Backup Exec marks the job as failed if the post-command fails.

For example, the pre-command can run successfully and shut down the database. The backup job can also run successfully. But if the post-command cannot restart the database, Backup Exec marks the job and the post-command as failed in the job log.

If you select the option **On each server backed up**, the pre- and post-command selections apply to each server independently. The pre- and post-commands are run and completed for one server at a time before they are run on the next selected server.

You can configure pre-commands and post-commands as default settings for all backup jobs. If the default settings are not appropriate for a particular job, you can override them when you create the job. You do not have to create default settings for pre-commands and post-commands, however. If you want to use pre-commands and post-commands only for specific jobs, you can configure the settings when you create those jobs.

To configure pre/post commands for backup jobs

1 Do one of the following:

To configure default pre/post commands for all backup jobs

Complete the following steps:

- Click the Backup Exec button, and then select **Configuration and Settings**.
- Select **Job Defaults**, and then select the type of backup for which you want to configure default pre/post commands.

To configure pre/post commands for specific backup jobs

Complete the following steps:

- Create a new backup definition or edit an existing backup definition.
- In the **Backup** box, click **Edit**.

- 2 In the left pane, click **Pre/Post Commands**.

3 Complete the following options:

Type a command to run before the backup runs	<p>Runs a command on the specified server before the backup job runs. Use local paths, and make sure that the paths exist on each server and are correct.</p> <p>Commands that require user interaction, such as prompts, are not supported.</p>
Run job only if pre-command is successful	<p>Runs the backup job only if the pre-command is successful. If the pre-command fails, the job does not run, and is marked as failed.</p> <p>If it is critical that the job does not run if the pre-command fails, then select Let Backup Exec check the exit codes of the commands to determine if the commands completed successfully. If a non-zero code is returned, Backup Exec interprets it to mean that the pre-command did not run successfully. The job does not run and the job status is marked as Failed.</p>
Type a command to run after the backup runs	<p>Runs a command on the specified server after the backup job runs. Use local paths, and make sure that the paths exist on each server and are correct.</p> <p>Commands that require user interaction, such as prompts, are not supported.</p>
Run post-command after job verification completes	<p>Runs the post-command after the verification completes, if you configured a verify operation for the job.</p> <p>See “Configuring automatic verify operations for backup jobs” on page 534.</p>
Run post-command only if pre-command is successful	<p>Runs the post-command only if the pre-command is successful.</p> <p>If it is critical that the post-command does not run if the pre-command fails, then select Let Backup Exec check the exit codes of the commands to determine if the commands completed successfully. If a non-zero code is returned for the pre-command, Backup Exec interprets it to mean that the pre-command did not run successfully. The post-command does not run.</p> <p>If you also select Run job only if pre-command is successful, and both the pre-command and the job are successful, but the post-command returns a non-zero code, the job log reports both the job and the post-command as failed.</p>

Run post-command even if job fails	<p>Runs the post-command regardless of whether the job is successful or not.</p> <p>If you also select Let Backup Exec check the exit codes of the commands to determine if the commands completed successfully and the post-command returns a non-zero code, the job log reports the post-command as failed.</p>
Let Backup Exec check the exit codes of the commands to determine if the commands completed successfully	<p>Lets Backup Exec check the return codes of the pre- and post-commands to determine if they completed successfully.</p> <p>Backup Exec interprets an exit code of zero from either the pre- or post-command to mean that the command completed successfully. Backup Exec interprets a non-zero code to mean that the command ended with an error.</p> <p>After Backup Exec checks the return codes, it continues processing according to the selections you made for running the pre- and post-commands.</p> <p>If this option is not selected, the success of the pre- and post-commands is not determined based on the return code.</p>
On this Backup Exec server	Runs the pre- and post-commands on this Backup Exec server only.
On each server backed up	<p>Runs the pre- and post-commands one time on each server that is backed up.</p> <p>The pre- and post-command selections apply to each server independently. If you select this option, the pre- and post-commands are run and completed for each server before Backup Exec begins processing on the next selected server.</p>
Cancel command if not completed within x minutes	Designates the number of minutes Backup Exec should wait before it cancels a pre- or post-command that did not complete. The default timeout is 30 minutes.

4 Click **OK**.

See [“Changing default backup job settings”](#) on page 515.

See [“Backing up data”](#) on page 158.

Configuring file and folder options for backup jobs

You can configure options for how Backup Exec processes file system attributes such as junction points and symbolic links.

File and folder options can be configured as default settings for all backup jobs. If the default settings are not appropriate for a particular job, you can override them when you create the job. You do not have to create default settings for files and folders, however. If you want to configure file and folder settings only for specific jobs, you can configure the settings when you create those jobs.

To configure file and folder options for backup jobs

1 Do one of the following:

To configure default file and folder options for all backup jobs

Complete the following steps:

- Click the Backup Exec button, and then select **Configuration and Settings**.
- Select **Job Defaults**, and then select the type of backup for which you want to configure default file and folder settings.

To configure file and folder settings for specific backup jobs

Complete the following steps:

- Create a new backup definition or edit an existing backup definition.
- In the **Backup** box, click **Edit**.

2 In the left pane, click **Files and Folders**.

3 Complete the following options:

Backup method for files

Select one of the following backup methods:

- **By modified time**

When Backup Exec runs a full backup job or an incremental backup job, the time that the backup job starts is recorded in the Backup Exec Database. Backup Exec adds the time of the backup job to the Backup Exec Database only if the full backup job completes successfully. The next time that you run an incremental backup job or a differential backup job, Backup Exec compares the file system time to the backup time. If the file system time is later than the time that is recorded in the database, the file is backed up. If the file's modified time is older than the previous backup's modified time, that file is not backed up. If the job does not complete successfully, subsequent differential or incremental backup jobs back up all of the data instead of only the data that has changed.

Note: A file's last modified date and timestamp do not change when the file is copied or moved. To ensure that the files are protected, run a full backup after you copy or move files.

When you run an incremental backup job, Backup Exec records a new time in the Backup Exec Database. The database time is not updated for differential backup jobs.

When you select the modified time method, Backup Exec uses the Windows change journal to determine if a file has changed since the last time it was backed up. If the change journal is not available, Backup Exec compares the file information to the previous backup time to determine if the file has changed.

Using modified time lets Backup Exec run more accurate incremental backups or differential backups even if other processes have modified files' archive bits.

- **Using archive bit**

Backup Exec uses the archive bit from the file system to determine if a file has changed since the last time it was backed up.

When you use the archive bit, Backup Exec turns the archive bit off when a file is backed up. Turning off the archive bit indicates to Backup Exec that the file has been backed up. If the file changes again before the next backup job, the bit is turned on again.

Backup Exec backs up the file in the next backup. If the next backup job is a full backup job or an incremental backup job, the bit is turned off when the backup job completes. If the next backup job is a differential backup job, the archive bit is left intact.

■ **Using catalogs**

Backup Exec compares path names, modified time, deleted and renamed files and folders, and other attributes. When you select the catalog method, Backup Exec uses the Windows change journal to determine if a file has changed since the last time it was backed up. If the change journal is not available, Backup Exec compares the file information to previous catalogs to determine if it has changed.

Note: You cannot use Backup Exec catalogs to determine if files were backed up for any differential backups.

The catalog method is only available if the Advanced Disk-based Backup Option (ADBO) is installed.

Note: The off-host backup feature of ADBO does not support the catalog method.

See [“How Backup Exec determines if a file has been backed up”](#) on page 199.

Enable single instance backup for NTFS volumes

Select this option if you want Backup Exec to check the NTFS volume for identical files. If Backup Exec finds multiple copies of a file, it backs up only one instance of that file.

Single instance backup can considerably reduce the storage space that is required for your backups. Many applications automatically generate some files that have identical content. The actual amount of space that you save depends on the number of duplicate files on the volume.

This option displays only if you use the Microsoft Windows Single Instance Store (SIS) feature.

Warning: If the backup job does not run to completion, the file data may not be included in the backup set. Rerun the backup job until it is successfully completed. If it was an incremental backup, running the job again does not back up the same files. You must run a full or duplicate backup job to ensure that all files are backed up completely.

Back up files and directories by following junction points and mount points

Select this option to back up the information for the junction points and the files and directories to which they are linked. If this check box is not selected, then only the information for the junction points is backed up. The files and directories to which the junction points are linked are not backed up.

Backup Exec does not follow junction points that are automatically created by Microsoft Windows because it can cause the data to be backed up repeatedly.

For more information, see the following Symantec knowledge base article:

<http://entsupport.symantec.com/umi/V-269-9>

You cannot select any mounted drives that do not have a drive letter assigned to them. The files and directories to which they are linked are backed up regardless of whether this option is selected.

If the files and directories to which the junction points are linked are also included in the backup selections, then they are backed up twice. They are backed up once during the full file and directory backup job, and again by the junction point.

Warning: If a junction point is linked to a location that encompasses it, then recursion (a situation where data is backed up repeatedly) occurs. Recursion results in an error and a job failure. For example, if c:\junctionpoint is linked to c:\, recursion occurs when Backup Exec attempts to back up c:\junctionpoint, and the backup job fails.

Back up files and directories by following symbolic links

Select this option to back up the information for any symbolic links and the files and directories to which they are linked.

If you do not select this option, only the information for the symbolic links is backed up. The files and directories to which they are linked are not backed up.

If the symbolic link points to files and directories on a remote computer, the files and directories on the remote computer are not backed up.

Back up data in Remote Storage

Select this option to back up any data that has been migrated from primary storage to secondary storage. The data is not recalled to its original location. It is backed up directly to the backup media.

If this option is selected, you should not run a backup of your entire system. Backup Exec has to load the data that has been migrated to secondary storage and additional time is required for migrated data.

If this check box is cleared, only the placeholder that stores the location of the data on secondary storage is backed up, not the data itself.

This option should not be selected if the device used for secondary storage and backups contains only one drive. If there is only one drive, Remote Storage and Backup Exec compete for use of the drive.

Back up open files

Select one of the following options to determine how Backup Exec processes any open files for the backup job.

The options are as follows:

- **Never**
 Backup Exec skips any open files that are encountered during the backup job. A list of any files that were skipped appears in the job log.
- **If closed within X seconds**
 Backup Exec waits the specified time interval for files to close before it skips them and continues the backup job.
 If the file does not close during the specified interval, it is skipped. A list of skipped files appears in the job log.
 If multiple files are open, Backup Exec waits the specified time interval for each file. Depending on the number of open files, the wait may significantly increase the backup time.
- **With a lock**
 Backup Exec attempts to open any files that are in use. If Backup Exec is able to open a file, the file is locked while it is backed up. Locking the file prevents other processes from writing to it.
 Backing up open files is not as effective as closing applications and allowing the files to be backed up in a consistent state.
- **Without a lock**
 Backup Exec attempts to open any files that are in use. If Backup Exec is able to open the file, the file is not locked while it is backed up. Other applications can write data to the file during the backup operation.
Warning: This option allows some files that contain inconsistent data and possibly corrupt data to be backed up.

Backup method

Select the backup method that you want to use to back up files and folders for each backup job in the backup definition.

See [“Backup methods in Backup Exec”](#) on page 188.

Delete selected files and folders after successful backup

Select this option if you want Backup Exec to delete the data you selected to back up after the backup job completes successfully.

Backup Exec backs up the selected data, verifies the backup sets, and then deletes the data from the server. The logon account credentials that you use to run the job must also have the rights to delete a file. Otherwise, the data is backed up, but it is not deleted.

Note: This option is only available when you configure backup jobs. You cannot configure it as a default setting for all backup jobs.

See [“Configuring Backup Exec to automatically delete files after a backup”](#) on page 202.

Preserve tree on back up and delete

Select this option to retain the file system's directory structure for the files that are backed up in a full backup job. This option is available only when you select the **Delete selected files and folders after successful backup** option.

4 Click **OK**.

See [“Changing default backup job settings”](#) on page 515.

See [“Backing up data”](#) on page 158.

Setting default schedule options for rule-based jobs and run now jobs

You can configure default schedule options for rule-based and run now jobs. Backup Exec applies the schedule options whenever you change a run now job or a rule-based job into a recurring scheduled job. A rule-based job is a job that is linked to another job. The rule-based job runs when the job to which it is linked is finished. An example of a rule-based job would be a duplicate stage that is configured to run when a full backup job completes. If you change the duplicate stage's schedule settings, Backup Exec uses the default schedule settings for the duplicate stage. You can override the default settings when you edit the newly scheduled job.

To set default schedule options for rule-based jobs and run now jobs

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Job Defaults**.
- 2 Select **Schedule**.

3 In the **Recurrence Pattern** group box, select the default frequency for backup jobs:

To run jobs every X hours or minutes

Click **Hours**, and then enter the frequency in the **Every X hour/minute** field.

Choose between the following options:

- **From**
Designate the starting time for a job to run.
- **Between**
Restrict the job to certain hours and days. For example, if you only want the job to run during business hours, you can select 9:00 AM to 5:00 PM on Monday, Tuesday, Wednesday, Thursday, and Friday.

To run jobs every X days

Click **Days**, and then choose between the following options:

- **Every X day**
Indicate the number of days between the start time of a job and the start time of the next job instance.
- **Every weekday**
Specifies that the job should run on Mondays, Tuesdays, Wednesdays, Thursdays, and Fridays.

To run jobs every X weeks

Click **Weeks**, and then specify the number of weeks between the start time of a job and the start time of the next job instance in the **Every X week on** field.

Select the days and time at which jobs should run.

- To run jobs every X months Click **Months**, and then choose between the following options:
- **Day X of every X month**
Specify the specific day on which jobs should run and indicate the number of months between the start time of a job and the start time of the next job instance.
 - **Every X X of every X month**
Specify the day on which jobs should run and indicate the number of months between the start time of a job and the start time of the next job instance.
 - **Selected days of the month**
Specify the days of the month on which Backup Exec should run jobs. The recurrence pattern that you select repeats itself every month.
The default setting is for the job to run every month on the current week and day of the month. For example, if you create the job on the third Monday of the month, the default setting is for the job to run once a month on the third Monday.
You can select additional days on which the job should run. Any additional days that you select are added to the monthly recurrence pattern.
 - **Selected dates of the month**
Specify the dates of the month on which Backup Exec should run jobs. The recurrence pattern that you select repeats itself every month.
The default setting is for the job to run every month on the current date of the month. For example, if you create the job on the 15th, the default setting is for the job to run once a month on the 15th.
You can select additional days on which the job should run. Any additional days that you select are added to the monthly recurrence pattern.
If you select the 31st, the job runs on the last day of the month in months that do not have 31 days. For example, if you configure the job to run on the 31st, in September the job runs on the 30th instead.

To run jobs every X years

Click **Years**, and then enter the frequency in the **Every X year** field.

Choose between the following options:

- **On X**
Specify the date on which Backup Exec should run jobs.
- **On the X of X**
Specify the day and month of the year on which Backup Exec should run jobs.

- 4 (Optional) Click **Calendar** to view all scheduled backup jobs on a calendar to check for scheduling conflicts.
- 5 In the **Reschedule the job if it does not start X hours after its scheduled start time** field, specify the amount of time past the job's scheduled start time at which Backup Exec changes the job completion status to Missed. The job is rescheduled to run based on the time window that you configured.
- 6 In the **Cancel the job if it is still running X hours after its scheduled start time** field, specify the amount of time after the job's scheduled start time at which you want to cancel the job if it is still running. Backup Exec changes the job completion status to Canceled, timed out.
- 7 Click **OK**.

See [“How job scheduling works in Backup Exec”](#) on page 215.

See [“List of job statuses in Backup Exec”](#) on page 277.

Excluding dates from the backup schedule for all backups

You can exclude specific dates, such as holidays, from your backup schedule. You may want to exclude holidays, for example, to ensure that Backup Exec does not run jobs on those days. You can exclude dates from the schedule for a specific backup job or you can exclude dates for all backup jobs.

When you exclude dates for all backups, any scheduled backup jobs do not run on those dates. All jobs resume running on their normal schedules after the exclude date. You can still create and run backup jobs and restore jobs on excluded dates, as long as they are not scheduled.

You can exclude dates in Backup Exec by selecting or typing dates on the **Exclude Dates** dialog box. Or you can create a text file with a list of dates to exclude and then import the text file.

After you create a list of dates to exclude, you can export a new text file with those dates. Exporting the text file can be useful if you want to copy your exclude dates from one Backup Exec server to another.

See [“Exporting a list of dates that are excluded from all backups to another server”](#) on page 565.

You can exclude dates from all backups in the job default settings or by using the backup calendar. Both features let you exclude dates from all backups. You can only import dates using the default settings. However, you may prefer to use the backup calendar because it gives you a visual representation of all of your scheduled jobs.

This topic includes the following procedures:

[To exclude dates from the backup schedule for all backups](#)

[To exclude dates from all backups using the backup calendar](#)

To exclude dates from the backup schedule for all backups

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Job Defaults**.
- 2 Select **Exclude Dates**.
- 3 Do any of the following:

To manually enter the date

Complete the following steps:

- In the **Select Date** field, type the date that you want to exclude from the backup schedule.
- Click **Add**.

Note: You can add only one date at a time.

To select the date from the calendar

Click the date that you want to exclude.

The calendar displays 3 months at a time. You can navigate forward and backward to view additional months by clicking the arrows.

Note: You can select only one date at a time.

To import a list of dates

Complete the following steps:

- Click **Browse**.
- Select the text file that contains the exclude dates.
- Click **Open**.
- Click **Import**.

To delete a date from the list of dates to exclude Complete the following steps:

- Select the date or dates that you want to remove from the list.
- Click **Delete**.

4 When you are finished selecting dates, click **OK**.

To exclude dates from all backups using the backup calendar

- 1 On the **Backup and Restore** tab, in the **Backups** group, click **Backup Calendar**.
- 2 Select the date that you want to exclude from the backup schedule.
- 3 Click **Exclude Dates**.
- 4 Click **Add exclude date for all backups**.

Note: To remove the exclusion from a selected date, click **Remove exclude date for all backups**.

5 Click **Close**.

Removing dates from the list of excluded dates

If you no longer want to exclude a date from your backup schedule, you can remove it from the list of excluded dates. When you remove a date from the list of excluded dates, the date becomes part of your regular backup schedule. Any recurring jobs that normally fall on that day are now scheduled to run rather than being skipped.

To remove dates from the list of excluded dates

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Job Defaults**.
- 2 Select **Exclude Dates**.
- 3 Select the date or dates that you want to remove from the list of excluded dates.

You can also remove dates from the list of excluded dates by clicking the excluded dates on the calendar.
- 4 Click **Delete**.
- 5 When you are finished removing dates from the list, click **OK**.

See [“Excluding dates from the backup schedule for all backups”](#) on page 562.

Exporting a list of dates that are excluded from all backups to another server

You can exclude specific dates, such as holidays, from your backup schedule. When you exclude dates, any regularly scheduled backups do not run on those dates. You can create a list of dates to exclude in Backup Exec.

See [“Excluding dates from the backup schedule for all backups”](#) on page 562.

You can import or export a list of dates to exclude as a text file. This may be useful if you want to copy a list of exclude dates from one Backup Exec server to another.

To export a list of exclude dates

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Job Defaults**.
- 2 Select **Exclude Dates**.
- 3 Click **Export**.
- 4 Browse to the location where you want to save the text file.
- 5 Click **Save**.

Changing the default preferences

You can change the settings for how you prefer Backup Exec to display various screens, indicators, and alerts.

To change the default preferences

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Preferences**.

3 Configure any of the following preferences:

Show splash screen at startup	Clear this option to show the Backup Exec Administration Console instead of the splash screen when you start Backup Exec.
Display progress indicators for backup jobs. This requires additional time to pre-scan resource.	<p>Select this option to display the percentage complete number while a backup job runs. These indicators appear in the Job Activity dialog box, and they let you monitor the progress of the active job. Backups might take longer to complete when this option is selected because the backup sources must be scanned to determine the amount of data to be backed up.</p> <p>Due to the amount of time that is required to scan the backup sources, you should not select this option when you back up remote resources.</p>
Allow Backup Exec to report anonymous usage information (No personally identifiable information will be returned to Symantec).	<p>Select this option to participate in the Symantec Backup Exec product improvement program.</p> <p>General Backup Exec usage and statistical information is periodically collected and sent anonymously to Symantec. Symantec uses the information to help improve the Backup Exec customer experience.</p> <p>Note: Although usage and statistical information is collected, Symantec never collects specific user information.</p>
Use alternating row colors	Select this option to display alternating row colors for all lists in Backup Exec. Alternating row colors can make it easier to distinguish between rows.
Re-enable	Select this option to enable any messages that you have disabled.

4 Click **OK**.

Configuring the default setting for backing up multiple servers or applications

You can select to back up multiple servers or applications at once using Backup Exec. You can back them up as part of one backup definition or you can back them up individually in separate backup definitions. It may be easier for you to manage

backing up multiple servers as part of one backup definition. However, it is easier to troubleshoot job failures if each server has its own backup definition.

Each time you create a backup definition that contains multiple servers or applications, you can select whether you want to create one backup definition or separate backup definitions. You can configure a default scenario for backing up multiple servers or applications so that Backup Exec automatically creates either one backup definition or separate backup definitions.

To configure the default setting for backing up multiple servers or applications

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, select **Backups**.
- 3 Select one of the following:
 - **Create one backup that includes all servers or applications**
 - **Create separate backups for each server or application**
- 4 If you want Backup Exec to prompt you each time you select to back up multiple servers or applications, select **Prompt each time I create a backup for multiple servers**.

If you disable the prompt, Backup Exec automatically uses the preference that you selected in the **When creating backups for multiple servers** field. You can enable the prompt at any time.

See [“Backing up data”](#) on page 158.

Configuring database maintenance

The Database Maintenance option lets you manage the Backup Exec Database. Each database maintenance operation is performed independently on each database. The Backup Exec Database maintains a record of the files and data that you have configured.

Database maintenance lets you perform the following:

- Optimize database size.
- Delete expired data.
- Save the contents of the database files.
- Perform a database consistency check.

Backup Exec generates informational alerts at the beginning and at the end of the database maintenance process each time database maintenance is performed.

The alerts provide details about the type of maintenance that was performed on each database and the amount of time that the maintenance took to complete. If the database maintenance process fails, the alert indicates where the failure occurred and the reason for the failure.

You do not have to select all the options; however, each one performs a different process that enables you to protect and maintain your database. Selecting all the options enables you to recover the database quickly and maintain optimal performance.

To configure database maintenance

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, click **Database Maintenance**.

3 Configure any of the following options:

Enable Backup Exec database maintenance	Select this option to enable the database maintenance process.
Perform database maintenance daily at	<p>Select the time that you want to perform database maintenance.</p> <p>All the maintenance occurs once a day at the time you specify.</p>
Delete aged data	Select this option to delete expired job history, job logs, alert history, and reports from the Backup Exec Database after the specified number of days have passed.
Keep job history for data on media that have current overwrite protection periods	<p>Select this option to keep all job history data for any media to which an overwrite protection period is currently assigned.</p> <p>After a media's overwrite protection period expires, the media's job history data can be deleted.</p>
Keep job history for specified number of days	<p>Select this option to indicate the number of days to keep job history data in the database before it is deleted.</p> <p>Job history data includes summary statistics for a job and details about media, devices, and any backup sets that were used to process the job.</p>
Job logs	<p>Indicate the number of days to keep job logs in the database before they are deleted.</p> <p>Job logs include detailed information about the job.</p>
Alert history	<p>Indicate the number of days to keep alert history data in the database before it is deleted.</p> <p>Alert history data includes property and response information for the alert.</p>
Reports	<p>Indicate the number of days to keep report data in the database before it is deleted.</p> <p>Report data includes property information about any report jobs that were generated. The report itself is not deleted.</p>

Audit logs

Indicate the number of days to keep audit log data in the database before it is deleted.

The audit log includes information about any operations that are performed in Backup Exec.

See [“Configuring audit logs”](#) on page 624.

Perform database consistency check

Select this option to check the logical consistency and physical consistency of the data in the database.

The option is not checked by default. Symantec recommends that you run a consistency check periodically at a time when there is minimal activity from Backup Exec.

Save contents of database to the Backup Exec data directory

Select this option to save the data that is contained in the database to the Backup Exec data directory so that the database backup file (BEDB.bak) can be backed up.

The dump file is maintained in the data directory until the next database maintenance process is performed and then this file is overwritten. Selecting this option enables you to recover the database in the event of failure.

Optimize database size

Select this option to organize fragmented pages and decrease the size of the physical database to 10 percent above what is actually used.

4 Click **OK**.

Scheduling Backup Exec to check logon accounts

You can schedule Backup Exec to check that the backup sources in jobs can be accessed with the logon accounts that you selected. Checking whether your logon accounts have access to backup sources lets you diagnose and fix any access issues before you run backup jobs. If Backup Exec discovers any backup sources that cannot be accessed with the logon accounts that you selected, it reports the error in an alert.

By default, Backup Exec is scheduled to check logon accounts every day at 1:00 am. You can disable the test if you do not want Backup Exec to regularly check logon accounts. You can also reschedule the test so that it occurs less frequently.

To schedule Backup Exec to check logon accounts

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Logon Accounts**.
- 3 Configure any of the following options:

Check that the backup sources in jobs can be accessed with the logon accounts you selected

Select this option to test whether Backup Exec can access the backup sources in your backup jobs with the logon accounts that you selected.

Perform check every X days at X

Enter the interval for Backup Exec to test logon accounts. Backup Exec automatically tests the logon accounts at the interval that you select in this field. You can select the time at which you want Backup Exec to run the test and the number of days between tests.

Check logon accounts only at the server level

Select this option to limit the logon account test to the server level. If you select this option, Backup Exec tests only whether the logon accounts can access backup sources at the server level. Any resources that reside on the server are not tested. The logon account test takes less time if you only check accounts at the server level, however it is less thorough.

- 4 Click **OK**.

See [“Backup Exec logon accounts”](#) on page 610.

Configuring Backup Exec to discover data to back up

Backup Exec’s **Discover Data to Back Up** option detects new backup content within a Windows or Active Directory domain. This option lets you configure a job that searches for new server volumes, databases, or application data.

By default, the data discovery operation runs at noon every day. It also runs each time the Backup Exec services are restarted. Backup Exec cancels the operation if it is still running after four hours. You can disable the operation or change the default settings in the global Backup Exec settings.

The **Discover Data to Back Up** option performs three main tasks:

- Discovers any top-level computers or computer contents
When the data discovery operation discovers top-level computers or computer contents, it adds them to the **Credentials** pane on the **Backup and Restore**

tab. The operation updates any information about the computers or computer contents and their backup status. You can view information about backup sources on the **Credentials** pane.

- Discovers any servers that do not have an Agent for Windows installed on them
 If the operation discovers any servers that do not have an Agent for Windows installed on them, Backup Exec sends you an alert. You can add the servers to the list of servers by using the **Add a Server** Wizard. After you add the servers to the list of servers, you can back them up and monitor them.
 See [“Adding discovered servers to the list of servers in Backup Exec”](#) on page 573.
- Discovers and validates instances of the Agent for Windows
 The data discovery operation searches for any instances of the Agent for Windows on your network. When it finds an Agent for Windows, the operation checks the version to make sure that it is up to date. If an Agent for Windows is not up to date with the most recent version, Backup Exec sends you an alert.

The data discovery operation only discovers the servers that meet the following criteria:

- Belongs to the same domain as the Backup Exec server
- Has the Windows Management Instrumentation (WMI) service enabled and running
- Allows WMI access for the same user that the Backup Exec Management Service runs under
 Members of the server's "Administrators" group have this level of access.
- Has firewalls that are configured to allow WMI network traffic

To configure Backup Exec to discover data to back up

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Discover Data to Back Up**.

3 Configure any of the following options:

Discover servers that have data that has not been backed up	<p>Select this option to allow Backup Exec to discover any data that needs to be backed up.</p> <p>When this option is selected, Backup Exec automatically checks your network for any data that has not been backed up.</p>
Frequency	<p>Lets you configure the frequency with which Backup Exec searches for any data that needs to be backed up.</p> <p>You can select to let Backup Exec search for any data that needs to be backed up daily, weekly, or monthly.</p>
Interval	<p>Lets you configure the interval at which Backup Exec searches for any data that needs to be backed up.</p> <p>You can select different intervals based on the frequency you selected.</p>
Cancel data discovery if not completed within	<p>Lets you select the number of hours after which the data discovery process is canceled if it is not finished.</p> <p>Canceling the data discovery process can help prevent it from affecting your system resources.</p>

4 Click **OK**.

Adding discovered servers to the list of servers in Backup Exec

You can use the **Add Server** Wizard to install the Agent for Windows on any servers that Backup Exec discovers by the **Discover Data to Back Up** option. After you install the Agent for Windows, the server is added to the list of servers in Backup Exec.

To add discovered servers to the list of servers in Backup Exec

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Select **Microsoft Windows computers and servers**, and then click **Next**.
- 3 Select **Allow Backup Exec to establish a trust with the servers**, and then click **Next**.
- 4 Click **Browse**.

- 5 Expand **Servers without an Agent for Windows installed** to see the names of the servers that the data discovery operation discovered.
- 6 Select the servers on which you want to install the Agent for Windows, and then click **OK**.
- 7 In the **Logon Account** field, select the logon account that you want to use to access each server.
- 8 Click **Next**.
- 9 Select any of the following options and then click **Next**.

Upgrade the Backup Exec Agent for Windows to the current version automatically	Upgrades the Agent for Windows on the server that you are adding to the list of servers, if necessary.
---	--

Restart the remote computer automatically after installing the Backup Exec Agent for Windows when a restart is required	Restarts the remote computer after the Agent for Windows is installed.
--	--

- 10 Click **Install**.

See [“Configuring Backup Exec to discover data to back up”](#) on page 571.

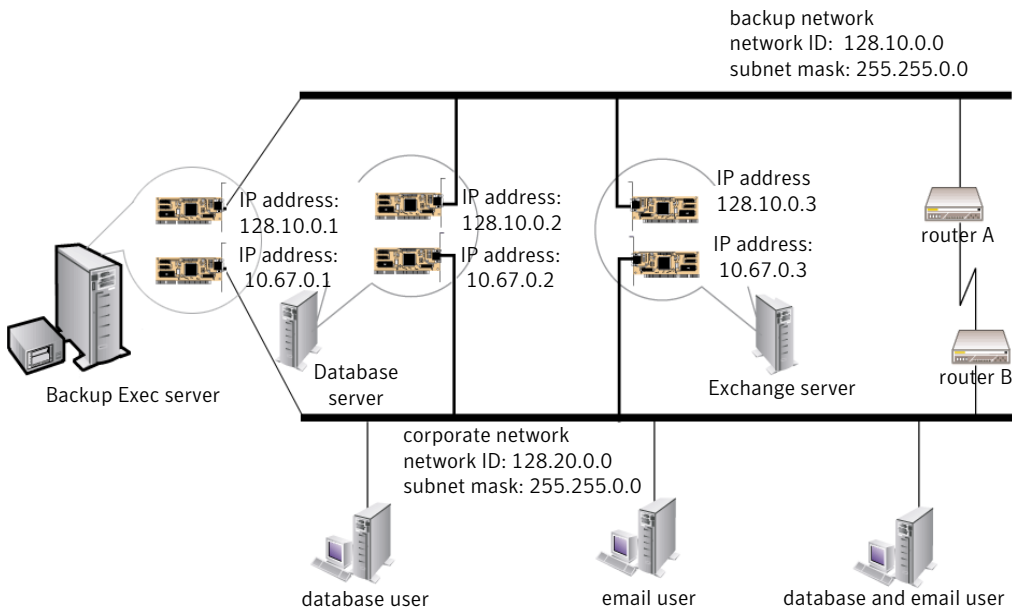
Backup networks

The backup network feature lets you direct any primary backup traffic that Backup Exec generates to a specific local network. Directing backup jobs to a specified local network isolates the backup data so that other connected networks are not affected when backup operations are performed. You also can use a backup network when you restore data. The feature is enabled on the Backup Exec server and lets you protect all the remote computers that reside on the specified local network.

When you specify a backup network and submit a job, Backup Exec verifies that the remote computer is on the same subnet as the selected interface on the Backup Exec server. If the remote computer is on the selected subnet, then the backup operation is performed.

If the remote computer is not on the selected subnet, then the job fails. However, you can set up Backup Exec to use any available network to back up remote computers.

The following diagram shows an example of a backup network configuration.

Figure 14-1 Example of backup network

In the example, the database server and mail server are connected to both the backup network and the corporate network.

When the Backup Exec server performs backup operations, the backup data uses either the backup network or the corporate network to back up the database server. If the backup data goes through the corporate network, the amount of time it takes to back up the database server increases. The amount of time increases because the network route between the two computers is longer. Users may experience network latencies when they access the mail server since there is an increase in network traffic.

In contrast, if you specify a backup network and you back up the database server, the backup data traffic is isolated to the backup network. Any users accessing the mail server are not affected. The backup network is used to perform all backup operations, unless the remote computer is not connected to the backup network.

To back up any remote computers that are not connected to the backup network, choose to use any available network route. Choosing any available network lets you back up the remote computer even though it does not reside on the backup network.

You can configure global network settings for all backup jobs on the **Network and Security** pane in the Backup Exec settings. If you want to override the global settings

for a particular backup job, you can configure network settings for individual jobs on the **Network** pane when you create backups.

See [“Changing network and security options for Backup Exec”](#) on page 576.

See [“Configuring network options for backup jobs”](#) on page 204.

See [“Using Backup Exec with firewalls”](#) on page 578.

Using IPv4 and IPv6 in Backup Exec

Backup Exec supports versions 4 and 6 of the Internet Protocol (IP), which are commonly referred to as IPv4 and IPv6. You can use IPv4 and IPv6 in backup and restore networks. Support for IPv6 is dependent upon operating system support for the protocol, as well as proper network configuration.

You can use Backup Exec in a mixed IPv4/IPv6 environment or an IPv4-only environment.

Enter an IPv4 or IPv6 address for a computer anywhere that you can enter a computer name in Backup Exec, except in the following locations:

- Clusters. Microsoft Windows does not support an IPv6 address as a clustered computer.
- The **Connect to Backup Exec Server** dialog box.

A Backup Exec agent that supports IPv6 can be backed up or restored using IPv6 only from a Backup Exec server that is IPv6-compliant.

Changing network and security options for Backup Exec

You can configure how Backup Exec works with your network configuration and security. The network and security options are global options that affect all Backup Exec jobs.

If the global network and security settings that you configure do not apply for a specific backup job, you can change the network settings when you create the backup job.

See [“Configuring network options for backup jobs”](#) on page 204.

To edit network and security options

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Network and Security**.
- 3 Configure any of the following options:

Network interface	Select the name of the network interface card that connects the Backup Exec server to the default network that you want to use for backup jobs. The list includes all available network interfaces on the Backup Exec server.
Protocol	<p>Select the default protocol you want to use for backup jobs.</p> <p>The options are as follows:</p> <ul style="list-style-type: none"> ■ Use any available protocol ■ IPv4 ■ IPv6
Subnet	Select the 32-bit number that determines the subnet to which the network interface card belongs.
Allow use of any available network interface, subnet, or protocol for Backup Exec agents not bound to the above network interface, subnet, or protocol	<p>Select this option to let Backup Exec use any available network if the remote system that you selected for backup or restore is not part of the specified backup network.</p> <p>If you do not select this option and the remote system is not part of the specified backup network, the job fails. Backup Exec cannot access the data from the remote system.</p>
Interface Details	Click this option to view the Media Access Control (MAC) address, adapter type, description, IP addresses, and subnet prefixes for the interface that you selected for the backup network.
Enable selection of user shares	<p>Select this option to include user-defined shares in jobs.</p> <p>If you do not select this option, you cannot select user-defined shares when you create jobs.</p>
Enable TCP dynamic port range	<p>Lets Backup Exec agents use a range of ports for communication.</p> <p>You enter the port range. If the first port that Backup Exec attempts to use is not available, Backup Exec attempts to use one of the other ports in the range. If none of the ports in the range is available, Backup Exec uses any available dynamic port. Default port ranges are 1025 to 65535. Symantec recommends using a range of 25 allocated ports for the remote system if you use Backup Exec with a firewall.</p> <p>See “Using Backup Exec with firewalls” on page 578.</p>

Use a custom port to receive operation requests from the Oracle server

Lets you specify the port that Backup Exec uses for communication between the Backup Exec server and the remote computer for both DBA and Backup Exec server-initiated operations. By default, Backup Exec uses port 5633.

If you change the port number on the remote Windows or Linux computer, you must also change it on the Backup Exec server. Then you must restart the Backup Exec Job Engine service on the Backup Exec server.

See [“About Oracle instance information changes ”](#) on page 1008.

Use FIPS 140-2 compliant software encryption

Lets you enable software encryption that complies with FIPS 140-2 standards. If you select this option, you must use a 256-bit AES encryption key. This option is available only for Windows computers.

You must stop and restart the Backup Exec services for this change to take effect.

Manage Keys

Lets you create a new encryption key or manage existing encryption keys.

4 Click **OK**.

See [“Backup networks”](#) on page 574.

Using Backup Exec with firewalls

In firewall environments, Backup Exec provides the following advantages:

- The number of ports that are used for backup network connections is kept to a minimum.
- Open ports on the Backup Exec server and remote systems are dynamic and offer high levels of flexibility during browsing, backup, and restore operations.
- You can set specific firewall port ranges and specify backup and restore networks within these ranges. You can use specific ranges to isolate data traffic and provide high levels of reliability.

Note: The Agent for Windows is required to perform remote backups and restores.

Firewalls affect system communication between a Backup Exec server and any remote systems that reside outside the firewall environment. You should consider special port requirements for your firewall when you configure Backup Exec.

Symantec recommends that you open port 10000 and make sure that it is available on the Backup Exec server and any remote systems. In addition, you must open the dynamic port ranges that Backup Exec uses for communications between the Backup Exec server and Backup Exec agents.

When a Backup Exec server connects to a remote system, it initially uses port 10000. The agent listens for connections on this predefined port. The Backup Exec server is bound to an available port, but additional connections to the agent are initiated on any available port.

When you back up data, up to two ports may be required on the computer on which the agent is installed. To support simultaneous jobs, you must configure your firewall to allow a range of ports large enough to support the number of simultaneous operations desired.

If there is a conflict, you can change the default port to an alternate port number by modifying the `%systemroot%\System32\drivers\etc\services` file. You can use a text editor such as Notepad to modify your NDMP entry or add an NDMP entry with a new port number. You should format the entry as follows:

```
ndmp      9999/tcp      #Network Data Management Protocol
```

Note: If you change the default port, you must change it on the Backup Exec server and all remote systems that are backed up through the firewall.

When you set up TCP dynamic port ranges, Symantec recommends that you use a range of 25 allocated ports for the remote computer. The number of ports that remote computers require depends on the number of devices you protect and the number of tape devices you use. You may need to increase these port ranges to maintain the highest level of performance.

Unless you specify a range, Backup Exec uses the full range of dynamic ports available. When performing remote backups through a firewall, you should select a specific range on the **Network and Security** settings dialog box.

To browse systems through a firewall

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, select **Network and Security**.

- 3 Verify that a dynamic range of ports has been set for the Backup Exec server and the Backup Exec agent and that the firewall is configured to pass these port ranges and the 10000 port (which is used for the initial connection from the Backup Exec server to the Backup Exec agent).

Port 6101 must be open to browse Windows systems in the backup selections tree.

- 4 Click **OK**.

See [“Backup Exec ports”](#) on page 580.

See [“Backup Exec listening ports”](#) on page 582.

See [“About enabling a SQL instance behind a firewall”](#) on page 583.

Backup Exec ports

You may have special port requirements for Backup Exec if you use a firewall. Firewalls sometimes affect system communications between a Backup Exec server and remote systems that reside outside the firewall environment.

See [“Using Backup Exec with firewalls”](#) on page 578.

The following table provides more information about which ports Backup Exec and its agents and options use:

Table 14-1 Backup Exec ports

Service or Process	Port	Port Type
Backup Exec Agent Browser (process=benetns.exe)	6101	TCP
Backup Exec Agent for Windows (process=beremote.exe)	10000	TCP
Backup Exec server (process=beserver.exe)	3527, 6106	TCP

Table 14-1 Backup Exec ports (*continued*)

Service or Process	Port	Port Type
MSSQL\$BKUPEXEC (process=sqlservr.exe)	A port number between 1024 and 5000 (Windows 2003) or between 49152 and 65535 (Windows 2008). See Microsoft's documentation for more information on assigning the TCP/IP port numbers that are required to communicate with SQL through a firewall.	TCP UDP
Agent for Oracle on Windows or Linux Servers	Random port unless configured otherwise	
Agent for Linux	Default NDMP port, typically 10000	TCP
Backup Exec deduplication engine (process=spoold.exe)	10082	TCP
Backup Exec deduplication manager (process=spad.exe)	10102	TCP
Kerberos	88	UDP
NETBIOS	135	TCP, UDP
NETBIOS Name Service	137	UDP
NETBIOS Datagram Service	138	UDP
NETBIOS Session Service	139	TCP
NETBIOS	445	TCP
DCOM/RPC	3106	TCP
Agent for Windows	6103	TCP
Push Install - Check for conflicts in message queue for CASO, which is part of beserver.exe	103x	TCP
Push Install	441	TCP

Table 14-1 Backup Exec ports (*continued*)

Service or Process	Port	Port Type
SMTP email notification	25 outbound from Backup Exec server	TCP
SNMP	162 outbound from Backup Exec server	TCP

Backup Exec listening ports

You may have special port requirements for Backup Exec if you use a firewall. Firewalls sometimes affect system communications between a Backup Exec server and remote systems that reside outside the firewall environment.

See [“Using Backup Exec with firewalls”](#) on page 578.

When Backup Exec is not running operations, it listens to ports for incoming communication from other services and agents. Backup Exec initially communicates with the agent using a static listening port to begin an operation. The agent and the Backup Exec server then use dynamic ports to pass data back and forth.

Backup Exec uses the following listening ports:

Table 14-2 Backup Exec listening ports

Service	Port	Port Type
Backup Exec Agent Browser (benetns.exe)	6101	TCP
Backup Exec Agent for Windows (beremote.exe)	10000	TCP
Backup Exec server (beserver.exe)	3527, 6106	TCP
MSSQL\$BKUPEXEC (sqlsevr.exe)	1125 1434	TCP UDP
Agent for Linux (RALUS)	10000	TCP
DBA-initiated backups for Oracle	5633	TCP

About enabling a SQL instance behind a firewall

If you want to connect to a SQL instance behind a firewall, you must enable the SQL instance for communication. To enable the SQL instance for communication, you must make the SQL port static and configure the Windows Firewall.

The Backup Exec SQL instance is configured to use a dynamic port by default. Each time SQL Server is started, the port number can change.

You also must configure the Windows Firewall to allow connections to the SQL instance. There may be multiple ways to configure the Windows Firewall based on your system configuration. You can add sqlsvr.exe and sqlbrowser.exe to the Windows Firewall Exceptions list or you can open a port in the Windows Firewall for TCP access. Refer to the Microsoft Knowledge Base for more information or to determine which configuration is best for your network.

See [“Installing a managed Backup Exec server across a firewall”](#) on page 1107.

See [“Using Backup Exec with firewalls”](#) on page 578.

Using encryption with Backup Exec

Backup Exec provides you with the ability to encrypt data. When you encrypt data, you protect it from unauthorized access. Anyone that tries to access the data has to have an encryption key that you create. Backup Exec provides software encryption, but it also supports some devices that provide hardware encryption with the T10 standard. Backup Exec configures encryption when you specify which storage devices that you want to use for a backup job.

Backup Exec supports two security levels of encryption: 128-bit Advanced Encryption Standard (AES) and 256-bit AES. The 256-bit AES encryption provides a stronger level of security because the key is longer for 256-bit AES than for 128-bit AES. However, 128-bit AES encryption enables backup jobs to process more quickly. Hardware encryption using the T10 standard requires 256-bit AES.

When you run a duplicate backup job, any backup sets that are already encrypted are not re-encrypted. However, you can encrypt any unencrypted backup sets.

This topic includes the following information:

[Software encryption](#)

[Hardware encryption](#)

[Encryption keys](#)

[Restricted keys and common keys](#)

[Pass phrases](#)

Software encryption

When you install Backup Exec, the installation program installs encryption software on the Backup Exec server and on any remote computers that use a Backup Exec agent. Backup Exec can encrypt data at a computer that uses a Backup Exec agent, and then transfer the encrypted data to the Backup Exec server. Backup Exec then writes the encrypted data on a set-by-set basis to tape or to disk storage.

Backup Exec encrypts the following types of data:

- User data, such as files and Microsoft Exchange databases.
- Metadata, such as file names, attributes, and operating system information.
- On-tape catalog file and directory information.

Backup Exec does not encrypt Backup Exec metadata or on-disk catalog file and directory information.

You can use software compression with encryption for a backup job. First Backup Exec compresses the files, and then encrypts them. However, backup jobs take longer to complete when you use both encryption compression and software compression.

Symantec recommends that you avoid using hardware compression with software encryption. Hardware compression is performed after encryption. Data becomes randomized during the encryption process. Compression does not work effectively on data that is randomized.

Hardware encryption

Backup Exec supports hardware encryption for any storage devices that use the T10 encryption standard. When you use hardware encryption, the data is transmitted from the host computer to the storage device and then encrypted on the device. Backup Exec manages the encryption keys that are used to access the encrypted data.

Backup Exec only supports approved devices for T10 encryption.

You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/v-269-2>

Note: Hardware encryption that uses the T10 standard requires 256-bit AES. Backup Exec does not let you enable hardware encryption for a job unless it uses at least a 16-character pass phrase.

Encryption keys

You must create encryption keys to use encryption in Backup Exec. When a user creates an encryption key, Backup Exec marks that key with an identifier based on the logged-on user's security identifier. The person who creates the key becomes the owner of the key.

If you use encryption for synthetic backups, all of the associated backups must use the same encryption key. Do not change the encryption key after the baseline is created. The encryption key that you select for the baseline backup is automatically applied to all associated backups.

When you select encrypted data for restore, Backup Exec verifies that encryption keys for the data are available in the database. If any of the keys are not available, Backup Exec prompts you to recreate the missing keys. If you delete the key after you schedule the job to run, the job fails.

If Backup Exec cannot locate an encryption key while a catalog job is running, Backup Exec sends an alert. You can then recreate the missing encryption key if you know the pass phrase.

Simplified Disaster Recovery supports the recovery of computers with previously encrypted backup sets. If you have Simplified Disaster Recovery backups that are encrypted during backup, the **Recover This Computer** wizard prompts you for the pass phrase of each encrypted backup set that is required to complete the recovery.

See [“Encryption key management”](#) on page 586.

Restricted keys and common keys

Backup Exec has the following types of encryption keys:

Table 14-3 Types of encryption keys

Key type	Description
Common	Anyone can use the key to encrypt data during a backup job and to restore encrypted data.
Restricted	Anyone can use the key to encrypt data during a backup job, but users other than the key owner must know the pass phrase. If a user other than the key owner tries to restore the encrypted data, Backup Exec prompts the user for the pass phrase. If you cannot supply the correct pass phrase for the key, you cannot restore the data.

Pass phrases

Encryption keys require a pass phrase, which is similar to a password. Pass phrases are usually longer than passwords and are comprised of several words or groups

of text. A good pass phrase is between 8 and 128 characters. The minimum number of characters for 128-bit AES encryption is eight. The minimum number of characters for 256-bit AES encryption is 16. Symantec recommends that you use more than the minimum number of characters.

Note: Hardware encryption that uses the T10 standard requires 256-bit AES. Backup Exec does not let you enable hardware encryption for a job unless it uses at least a 16-character pass phrase.

Also, a good pass phrase contains a combination of upper and lower case letters, numbers, and special characters. You should avoid using literary quotations in pass phrases.

A pass phrase can include only printable ASCII characters, which are characters 32 through 126. ASCII character 32 is the space character, which is entered using the space bar on the keyboard. ASCII characters 33 through 126 include the following:

!"#\$%&'()*+,-./0123456789;<=>?@ABCDEFGHIJKLMN O PQRSTUVWXYZ

[] ^ _ ' a b c d e f g h i j k l m n o p q r s t u v w x y z { } ~

See [“Encryption key management”](#) on page 586.

Encryption key management

When a user creates an encryption key, Backup Exec marks that key with an identifier based on the logged-on user's security identifier. The person who creates the key becomes the owner of the key.

See [“Using encryption with Backup Exec”](#) on page 583.

Backup Exec stores the keys in the Backup Exec database. However, Backup Exec does not store the pass phrases for the keys. The owner of each key is responsible for remembering the pass phrase for the key.

To protect your keys, Symantec recommends the following:

- Maintain a written log of the pass phrases. Keep the log in a safe place in a separate physical location from the encrypted backup sets.
- Back up the Backup Exec database. The database keeps a record of the keys.

Caution: If you do not have a backup of the Backup Exec database and do not remember your pass phrases, you cannot restore data from the encrypted media. In addition, Symantec cannot restore encrypted data in this situation.

A key that is created on a Backup Exec server is specific to that Backup Exec server. You cannot move keys between Backup Exec servers. However, you can create new keys on a different Backup Exec server by using existing pass phrases. A pass phrase always generates the same key. In addition, if you delete a key accidentally, you can recreate it by using the pass phrase.

If a Backup Exec database becomes corrupted on a Backup Exec server and is replaced by a new database, you must manually recreate all of the encryption keys that were stored on the original database.

If you move a database from one Backup Exec server to another Backup Exec server, the encryption keys remain intact as long as the new Backup Exec server meets the following criteria:

- Has the same user accounts as the original Backup Exec server.
- Is in the same domain as the original Backup Exec server.

See [“Creating encryption keys”](#) on page 587.

See [“Replacing an encryption key”](#) on page 588.

See [“Deleting encryption keys”](#) on page 590.

Creating encryption keys

When you create an encryption key, you select the type of encryption to use.

To create an encryption key

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Network and Security**.
- 3 Click **Manage Keys**.
- 4 Click **New**.
- 5 In the **Key name** field, type a unique name for this key. The name can include up to 256 characters.
- 6 In the **Encryption type** field, select the encryption type to use for this key. Your choices are 128-bit AES or 256-bit AES.

The default type is 256-bit AES. The 256-bit AES encryption provides a stronger level of security than 128-bit AES encryption. However, backup jobs may process more slowly with 256-bit AES encryption than with 128-bit AES encryption. Hardware encryption that uses the T10 standard requires 256-bit AES.

- 7 In the **Pass phrase** field, type a pass phrase for this key. You can use only printable ASCII characters.

For 128-bit AES encryption, the pass phrase must be at least eight characters.
For 256-bit AES encryption, the pass phrase must be at least 16 characters.

Symantec recommends that you use more than the minimum number of characters.

Warning: If an encryption key that is used in a backup is no longer available, you must provide the pass phrase during restore. Without the pass phrase, the data cannot be accessed.

- 8 In the **Confirm pass phrase** field, type the pass phrase again to confirm it.
- 9 In the **Encryption key type** group box, select whether you want to create a common or restricted encryption key.

If a key is common, any user of this installation of Backup Exec can use the key to back up and restore data. If a key is restricted, anyone can use the key to back up data. But only the key owner or a user who knows the pass phrase can use the restricted key to restore the encrypted data.
- 10 Click **OK**.

See [“Encryption key management”](#) on page 586.
See [“Using encryption with Backup Exec”](#) on page 583.

Replacing an encryption key

You can replace one encryption key with another for all backup jobs and duplicate backup set jobs.

Note: You cannot replace an encryption key if it is used in a restore job.

To replace an encryption key

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Network and Security**.
- 3 Click **Manage Keys**.
- 4 Select the key that you want to replace.
- 5 Click **Replace**.

- 6 In the **Select an encryption key to replace <key name>** field, do one of the following:

To use an existing key Select the key from the list.

To create a new key Complete the following steps:

- Click **New**.
- In the **Key name** field, type a unique name for the key.
The name can include up to 256 characters.
- In the **Encryption type** field, select the type of encryption to use for this key.
You can select 128-bit AES or 256-bit AES. The default type is 256-bit AES.
The 256-bit AES encryption provides a stronger level of security than 128-bit AES encryption. However, backup jobs may process more slowly with 256-bit AES encryption than with 128-bit AES encryption.
Hardware encryption that uses the T10 standard requires 256-bit AES.
- In the **Pass phrase** field, type a pass phrase for this key.
For 128-bit AES encryption, the pass phrase must be at least eight characters. For 256-bit AES encryption, the pass phrase must be at least 16 characters. Symantec recommends that you use more than the minimum number of characters.
You can use only printable ASCII characters.
Warning: If an encryption key that is used in a backup is no longer available, you must provide the pass phrase during restore. Without the pass phrase, the data cannot be accessed.
- In the **Confirm pass phrase** field, type the pass phrase again to confirm it.
- In the **Encryption key type** group box, select whether you want to create a common or restricted encryption key.
- Click **OK**.

- 7 Click **OK**.

See [“Encryption key management”](#) on page 586.

See [“Using encryption with Backup Exec”](#) on page 583.

Deleting encryption keys

You should be cautious when you delete encryption keys. When you delete an encryption key, you cannot restore the backup sets that you encrypted with that key unless you create a new key that uses the same encryption key and pass phrase as the original key.

You can delete encryption keys in the following situations:

- The encrypted data on the tape has expired or the tape is retired.
- The encryption key is not the default key.
- The encryption key is not being used in a job. If the key is being used, you must select a new key for the job.

If you delete an encryption key that is being used in a scheduled restore job, you cannot replace the key. Therefore, any scheduled restore job in which you delete an encryption key fails.

To delete an encryption key

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **Network and Security**.
- 3 Click **Manage Keys**.
- 4 Select the key that you want to delete.
- 5 Click **Delete**.
- 6 Click **Yes**.
- 7 If the key is used in a job, do the following:
 - In the **Select an encryption key to replace "key name"** box, select the new key for the job or click **New** to create a new key.
 - Click **OK**.

See [“Encryption key management”](#) on page 586.

See [“Using encryption with Backup Exec”](#) on page 583.

Granular Recovery Technology

You can use Granular Recovery Technology (GRT) to restore certain individual items from backup sets. For example, you can use the Agent for Microsoft Exchange Server to restore an email from a backup without having to restore the entire mailbox.

Or, you can use the Agent for Microsoft SharePoint to restore a list without restoring the entire site.

To restore individual items, the Granular Recovery Technology feature must be enabled when you create a backup job.

GRT is enabled by default for backups for the following agents:

- Agent for Microsoft Active Directory
- Agent for Microsoft Exchange Server
- Agent for Microsoft SharePoint
- Agent for VMware and Hyper-V

You can restore either full backup sets or individual items from GRT-enabled backups.

By default, the Agent for VMware and Hyper-V uses Granular Recovery Technology to protect files and folders at a granular level. You can also enable the granular recovery of Microsoft Exchange, SQL, SharePoint, and Active Directory application data that resides on virtual machines.

When you back up data, Backup Exec creates a catalog that contains information about the backup sets and about the storage device on which the backup sets are stored. GRT-enabled backup jobs require more time to catalog because of the amount of granular information that they contain. For GRT-enabled backup jobs, the catalog operation is delayed and run as a separate operation to have less of an effect on your backup window. Because the catalog operation runs separately from the backup job, it does not prevent another scheduled backup job from starting on time.

When you enable GRT for Microsoft Exchange, Microsoft SharePoint, Microsoft Hyper-V, or VMware backups, the delayed catalog operation runs after the backup job by default. For Exchange and SharePoint backups, the delayed catalog operation runs immediately after all full backups. It runs once every 24-hours for all incremental backups and differential backups, even if you schedule more than one GRT-enabled job to run in the 24 hour period.

For Hyper-V and VMware backups, the delayed catalog operation runs immediately after all full, incremental, and differential backups by default.

You can configure the delayed catalog operation to run on a schedule if you do not want it to run immediately after the backup job. You can also choose to not run a delayed catalog operation.

See [“Configuring delayed catalog operations to improve backup performance”](#) on page 536.

The following table lists the individual items you can restore for each agent.

Table 14-4 Individual items that can be recovered for each agent

Agent	Individual items
Agent for Microsoft Active Directory	<p>You can restore the following individual items:</p> <ul style="list-style-type: none"> ■ Active Directory objects and attributes ■ Active Directory Application Mode (ADAM) and Active Directory Lightweight Directory Services (AD LDS) objects and attributes
Agent for Microsoft Exchange Server	<p>You can restore the following individual items:</p> <ul style="list-style-type: none"> ■ Mailboxes ■ Mail messages and their attachments ■ Public folders ■ Calendar items ■ Contacts ■ Notes ■ Tasks
Agent for Microsoft SharePoint	<p>The following are examples of the individual items that can be restored:</p> <ul style="list-style-type: none"> ■ Site collections ■ Sites or subsites ■ Document or picture libraries ■ Lists ■ Individual list items ■ Documents, pictures, or other files that are stored in libraries
Agent for VMware and Hyper-V	<p>You can restore drives, folders, and files from virtual machines that run a Windows operating system.</p> <p>You can also enable the granular recovery of Microsoft Exchange, SQL, SharePoint, and Active Directory application data that resides on virtual machines.</p> <p>See “Using Granular Recovery Technology (GRT) with the Agent for VMware” on page 847.</p>

When you run a GRT-enabled backup job, Backup Exec creates media with an IMG prefix (for example, IMG00001). IMG media is a specific media type that Backup Exec creates only for GRT-enabled backup operations. When you run a GRT-enabled backup job, the IMG media stores the backup data.

Note: Symantec recommends that you do not use software compression or encryption for GRT-enabled backup jobs. The compression and encryption processes are resource-intensive. Enabling either software compression or encryption can result in degraded performance for GRT-enabled backup jobs.

You should consider which device you use for GRT-enabled backups before you begin. You should also consider any special requirements for the type of data you back up.

Recommended devices for backups that use Granular Recovery Technology

Symantec recommends that you select a disk storage device for any backups that are enabled for Granular Recovery Technology (GRT). The disk storage device should be on a volume that does not have file size limitations. An NTFS drive is an example of a volume without file size limitations. Some examples of volumes that have file size limitations include FAT and FAT32 volumes.

If you must use a disk storage device on a volume with file size limitations, Backup Exec requires a staging location. Backup Exec temporarily stores a small amount of metadata in the staging location during the backup job. It deletes the data from the staging location when the backup is finished. The staging location is not necessary, however, if you use a disk storage device on a volume without file size limitations as the destination.

The staging location's default path is C:\temp.

The volume that is used for a staging location for backup jobs should meet the following requirements:

- It is local to the Backup Exec server
- It does not have any file size limitations

Additionally, Symantec recommends the following to avoid disk space problems:

- It should not be a system volume
- It should have at least 1 GB of available space

Backup Exec also uses a staging location to restore GRT-enabled data from a tape or from a disk storage device on volumes with file size limitations. The staging location must be on a volume that does not have file size limitations and is local to the Backup Exec server. The staging location is not necessary if you restore GRT-enabled data from disk storage on a volume without file size limitations, such as NTFS.

Backup Exec uses the staging area differently for the following types of restores:

Table 14-5 Staging processes

Location of data to be restored	Staging process
Tape	<p>Backup Exec copies the entire backup set or sets to the staging area. The staging area must have enough disk space for the entire backup set or sets from which you want to restore an individual item.</p> <p>Before you use a tape device for a GRT-enabled backup, ensure that sufficient disk space is available to perform a restore.</p> <p>Backup Exec deletes the data from the staging area when the restore job is complete.</p>
Disk storage device that is on a volume with file size limitations (such as FAT or FAT32)	<p>Backup Exec must copy a small amount of metadata that is associated with the backup set to the staging area to complete the restore.</p> <p>Backup Exec deletes the data from the staging area when the restore job is complete.</p>

The staging location's default path is C:\temp. You can change the default backup and restore staging locations in the Backup Exec settings.

Requirements for jobs that use Granular Recovery Technology

Keep in mind the following requirements when you use Granular Recovery Technology (GRT) with the agents listed:

Table 14-6 Granular Recovery Technology requirements

Agent	Restrictions
Files and Folders	<p>GRT is not supported for the following scenarios:</p> <ul style="list-style-type: none">■ Using a Windows Server 2003 Backup Exec server to back up System State on a Windows Server 2012 or 2012 R2 computer■ Using any 32-bit Backup Exec server to back up System State on a 64-bit computer <p>In both of these scenarios the backup completes successfully, but it is not GRT-enabled.</p>

Table 14-6 Granular Recovery Technology requirements (*continued*)

Agent	Restrictions
Agent for Microsoft Active Directory	<p>You must use a Backup Exec server that is running Windows Server 2012 R2 to back up an Active Directory server that is running Windows Server 2012 R2.</p> <p>In a CASO environment, Backup Exec runs the backup job on the central administration server if it is running Windows Server 2012 R2 and the storage is configured locally. Otherwise, Backup Exec attempts to find a managed Backup Exec server that is running Windows Server 2012 R2 to run the job. If it cannot find one, the job completes with a status of success with exceptions and the backup sets are not GRT-enabled.</p> <p>If the Active Directory server is a virtual machine, Backup Exec may not be able to detect that the server is running Windows Server 2012 R2. The resulting backup may not be GRT-enabled.</p>
Agent for Microsoft Exchange Server	<p>Backup Exec must have access to a uniquely named mailbox within the Exchange organization for backup and restore of the Information Store.</p> <p>See "Requirements for accessing Exchange mailboxes " on page 944.</p> <p>For Microsoft Exchange 2007, if the log files that you back up reside on a large-sector drive, you must use the same type of large-sector drive as the staging area. For example, if the log files are on a 512e disk, the staging area must also be on a 512e disk.</p> <p>For Microsoft Exchange 2010, you must use a large-sector drive as the staging area if the Exchange data resides on a large-sector drive.</p> <p>You can create a vhd file or vhdx file to use as the staging area instead of a physical volume, if you want to minimize the effect on your resources. Configure the vhd or vhdx file as a large-sector volume and make sure that it is at least 1GB in size. Then mount the vhd file or vhdx file to a directory. Backup Exec automatically finds the virtual volume when it needs a staging area to back up any data that resides on a large-sector drive.</p>
Agent for Microsoft SharePoint	<p>You must have a current version of the Agent for Windows installed on all of the servers that participate in the SharePoint farm.</p>

Table 14-6 Granular Recovery Technology requirements (*continued*)

Agent	Restrictions
Agent for VMware and Hyper-V	<p>You can recover only individual items to virtual machines that run a Windows operating system.</p> <p>By default, the Agent for VMware and Hyper-V uses Granular Recovery Technology to protect files and folders at a granular level. You can also enable the granular recovery of Microsoft Exchange, SQL, SharePoint, and Active Directory application data that resides on virtual machines.</p> <p>If you back up a virtual machine to tape, the Use storage-based catalogs option must be enabled in the Catalog settings to ensure that you can recover individual items from the backup sets.</p> <p>See “Configuring default options for catalogs” on page 249.</p>

See [“Setting default Granular Recovery Technology \(GRT\) options”](#) on page 596.

Setting default Granular Recovery Technology (GRT) options

Backup Exec's Granular Recovery Technology (GRT) feature lets you restore certain individual items from backup sets. For example, you can use the Agent for Microsoft Exchange Server to restore an email from a backup without having to restore the entire mailbox. Or, you can use the Agent for Microsoft SharePoint to restore a list without restoring the entire site.

See [“Granular Recovery Technology”](#) on page 590.

You can configure default settings for Granular Recovery Technology. Backup Exec applies the default settings to any backup jobs that you configure to use GRT.

To set default Granular Recovery Technology options

- Click the Backup Exec button, and then select **Configuration and Settings**, and then click **Backup Exec Settings**.
- In the left pane, select **Granular Recovery Technology**.
- Configure any of the following options:

If Granular Recovery Technology (GRT) is enabled for backups, enter the path to an NTFS volume of the local Backup Exec server where Backup Exec can stage temporary data

Lets you designate a location where Backup Exec can stage temporary data during GRT-enabled backup jobs.

Ensure that the location is an NTFS volume and that it is not a system volume. If the default path of C:\TEMP does not meet these requirements, type a different path on the Backup Exec server where Backup Exec can stage temporary data.

Backup Exec deletes the data when the backup job is completed.

At least 1 GB of disk space is required.

Enter the path to an NTFS volume of the local Backup Exec server where Backup Exec can store temporary data (Microsoft Hyper-V, Microsoft Exchange, Microsoft SharePoint, Microsoft Active Directory, and VMware)

Lets you designate a location where Backup Exec can stage temporary data during GRT restore jobs.

This option is applicable only when you restore individual items under the following conditions:

- The backup of Microsoft Hyper-V, Microsoft Exchange, Microsoft SharePoint, Microsoft Active Directory, or VMware Virtual Infrastructure was enabled for Backup Exec GRT.
- The backup is on a tape.
- The backup is on disk storage on a volume that has size limitations. FAT and FAT32 are examples of types of volumes that have file size limitations.

Type the path to a folder on an NTFS volume on this Backup Exec server. Restore data and metadata for this job are stored here temporarily before the individual items are restored. The staged data is automatically deleted when the restore job is completed.

Symantec recommends that you avoid using system volumes for temporary staging locations.

4 Click **OK**.

DBA-initiated job templates

When you create a DBA-initiated backup operation, you can specify the default job template in Backup Exec. You can also specify a new job template that you create in Backup Exec. The job template contains the settings that Backup Exec applies to DBA-initiated jobs.

Make sure that the name of the job template that you want to use is also configured in the instance information on the Windows computer.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 994.

See [“About performing a DBA-initiated backup job for Oracle”](#) on page 1010.

Note the following about DBA-initiated jobs:

- DBA-initiated jobs fail when the related job template is deleted. To stop DBA-initiated jobs from running, delete the related DBA-initiated job template.
- All DBA-initiated backup and restore jobs are deleted after the jobs have completed.
- You cannot set minimum device requirements for DBA-initiated jobs.

See [“Creating DBA-initiated job templates”](#) on page 598.

See [“Editing DBA-initiated job templates”](#) on page 599.

See [“Deleting DBA-initiated job templates”](#) on page 599.

Creating DBA-initiated job templates

You can create a new job template that Backup Exec applies to DBA-initiated jobs.

To create DBA-initiated job templates

- 1 Click the Backup Exec button, and then select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **DBA-initiated Job Settings**.
- 3 Click **New**.
- 4 In the left pane, select the type of options that you want to set. You determine the options that need to be set based on the needs of your environment.
- 5 Click **OK**.

See [“DBA-initiated job templates”](#) on page 597.

See [“Storage options for DBA-initiated jobs”](#) on page 600.

See [“General options for DBA-initiated jobs”](#) on page 605.

See [“Network options for DBA-initiated jobs”](#) on page 605.

See [“Migrator for Enterprise Vault options”](#) on page 1053.

See [“Notification options for jobs”](#) on page 305.

See [“Duplicate job settings for DBA-initiated jobs”](#) on page 606.

Editing DBA-initiated job templates

You can edit the job template settings that Backup Exec applies to DBA-initiated jobs.

To edit DBA-initiated job templates

- 1 Click the Backup Exec button, and then select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **DBA-initiated Job Settings**.
- 3 Select the job template that you want to edit.
- 4 Click **Edit**.
- 5 In the left pane, select the type of options that you want to edit. You determine the options that need to be set based on the needs of your environment.
- 6 Click **OK**.

See [“DBA-initiated job templates”](#) on page 597.

See [“Storage options for DBA-initiated jobs”](#) on page 600.

See [“General options for DBA-initiated jobs”](#) on page 605.

See [“Network options for DBA-initiated jobs”](#) on page 605.

See [“Migrator for Enterprise Vault options”](#) on page 1053.

See [“Notification options for jobs”](#) on page 305.

See [“Duplicate job settings for DBA-initiated jobs”](#) on page 606.

Deleting DBA-initiated job templates

You can delete the templates that Backup Exec applies to DBA-initiated jobs if you no longer use them.

To delete a job template for DBA-initiated jobs

- 1 Click the Backup Exec button, and then select **Configuration and Settings**, and then click **Backup Exec Settings**.
- 2 In the left pane, select **DBA-initiated Job Settings**.
- 3 Select the job template that you want to delete.
- 4 Click **Delete**.
- 5 Click **Yes**.

See [“DBA-initiated job templates”](#) on page 597.

Storage options for DBA-initiated jobs

You can configure storage settings for DBA-initiated jobs.

See [“DBA-initiated job templates”](#) on page 597.

Table 14-7 Storage options for DBA-initiated jobs

Item	Description
Backup Exec server or Backup Exec server pool	<p>Specifies if you want a job to run on devices on a specific managed Backup Exec server or on devices that are on a group of managed Backup Exec servers.</p> <p>This option displays only if you have the Central Admin Server Option installed. This option is an additional filter that lets you control where certain jobs are delegated. For example, to always run backups of Exchange databases only on the devices that are attached to managed Backup Exec servers in a pool named Exchange Backups, select this option. Then select the Exchange Backups Backup Exec server pool.</p>
Storage	<p>Specifies the storage device to which you want to send backup data for the DBA-initiated jobs.</p> <p>See “Creating storage device pools” on page 416.</p> <p>See “About the Any Virtual Disk Storage device pool in the Storage Provisioning Option” on page 1278.</p> <p>See “About the Remote Media Agent for Linux ” on page 1241.</p> <p>See “Features and types of disk-based storage and network-based storage” on page 313.</p>
Enable the remote computer to directly access the storage device and to perform client-side deduplication, if it is supported	<p>Enables a remote computer to send data directly to an OpenStorage device or a deduplication disk storage device, and to perform client-side deduplication if the device supports it. The Backup Exec server is bypassed, which leaves the Backup Exec server free to perform other operations. If client-side deduplication cannot be performed, then either Backup Exec server deduplication or Appliance deduplication is performed.</p> <p>This option appears if the Deduplication Option is installed and an OpenStorage device or a deduplication disk storage device is selected in the Storage field.</p> <p>See “How to use client-side deduplication” on page 821.</p>

Table 14-7 Storage options for DBA-initiated jobs (*continued*)

Item	Description
Enable the remote computer to access the storage device through the Backup Exec server and to perform Backup Exec server-side deduplication, if it is supported	<p>Enables a remote computer to send data through the Backup Exec server to an OpenStorage device or a deduplication disk storage device, and to perform Backup Exec server-side deduplication if it is supported. If the Backup Exec server does not support deduplication, the data is deduplicated on an intelligent disk device, such as Symantec PureDisk or a device from a third-party vendor.</p> <p>This option appears if the Deduplication Option is installed and an OpenStorage device or a deduplication disk storage device is selected in the Storage field.</p> <p>See “About the Deduplication Option” on page 795.</p>
Keep for	Designates the amount of time for which you want to keep the backup sets or job history from the DBA-initiated jobs.
Media set	<p>Indicates the media set to use for the DBA-initiated jobs. The media set specifies the overwrite protection period and the append period for the backup data on the media.</p> <p>If you want to create a new media set for this backup job, click the icon to the right of the media set drop-down menu.</p> <p>This option is available only if you selected a tape device in the Storage field.</p> <p>See “Default media sets” on page 379.</p>

Table 14-7 Storage options for DBA-initiated jobs (*continued*)

Item	Description
Overwrite media	<p>Indicates that the backup job is placed on an overwritable media. Ensure that appropriate media is in the storage device that you select.</p> <p>Appropriate media for an overwrite job includes the following:</p> <ul style="list-style-type: none">■ Scratch media■ Media for which the overwrite protection period has expired <p>Allocated or imported media may also be overwritten depending on the media overwrite protection level that is set.</p> <p>Depending on your configuration, overwritable media is selected from scratch media or recyclable media.</p> <p>If the media in the storage device is not overwritable, an alert appears to prompt you to insert overwritable media.</p> <p>This option is available only if you selected a tape device in the Storage field.</p> <p>See “Managing tapes” on page 378.</p> <p>See “Media overwrite protection levels for tape media” on page 394.</p> <p>See “How Backup Exec searches for overwritable media in tape drives ” on page 395.</p>
Append to media, overwrite if no appendable media is available	<p>Appends this backup job to the specified media set if an appendable media is available. Otherwise, Backup Exec searches for an overwritable media and adds it to the media set.</p> <p>If an append job fills a media, the backup job continues on an overwritable media. If the media in the storage device is not overwritable, an alert appears to prompt you to insert overwritable media.</p> <p>This option is available only if you selected a tape device in the Storage field.</p>
Append to media, terminate job if no appendable media is available	<p>Appends this backup job to the specified media set if an appendable media is available. Otherwise, Backup Exec terminates the job.</p> <p>This option is available only if you selected a tape device in the Storage field.</p>

Table 14-7 Storage options for DBA-initiated jobs (*continued*)

Item	Description
Eject media after job completes	<p>Ejects the media from the drive or slot when the operation completes. You can also schedule a job to eject media.</p> <p>This option is available only if you selected a tape device in the Storage field.</p> <p>See “Ejecting media from a disk cartridge or tape drive” on page 458.</p>
Retension media before backup	<p>Runs the tape in the drive from beginning to end at a fast speed. Retensioning helps the tape wind evenly and run more smoothly past the tape drive heads. This option is available only if you select a tape drive that supports retensioning.</p>
Use Write once, read many (WORM) media	<p>Specifies the use of WORM (write once, read many) media as the default for DBA-initiated jobs. Backup Exec confirms that the destination device is or contains a WORM-compatible drive, and that the WORM media is available in the drive. If WORM media or a WORM-compatible drive is not found, an alert is sent.</p> <p>See “How WORM media is used in Backup Exec” on page 401.</p>

Table 14-7 Storage options for DBA-initiated jobs (*continued*)

Item	Description
Compression	<p>Provides the following compression options:</p> <ul style="list-style-type: none"> ■ None Copies the data to the media in its original form (uncompressed). Using some form of data compression can help expedite backups and preserve storage space. Hardware data compression should not be used in environments where storage devices that support hardware compression are used interchangeably with devices that do not have that functionality. In this situation, hardware compression is automatically disabled. You can manually turn on hardware compression on the drives that support it, but this results in media inconsistency. If the drive that supports hardware compression fails, the compressed media cannot be restored with the non-compression drive. ■ Software Uses STAC software data compression, which compresses the data before it is sent to the storage device. ■ Hardware (if available, otherwise none) Uses hardware data compression if the storage device supports it. If the drive does not feature data compression, the data is backed up uncompressed. ■ Hardware (if available, otherwise software) Uses hardware data compression if the storage device supports it. If the drive does not feature hardware data compression, STAC software compression is used.
Encryption type	<p>Specifies the type of encryption that you want to use, if any.</p> <p>See “Using encryption with Backup Exec” on page 583.</p>
Encryption key	<p>Specifies the encryption key that you want to use, if you selected to use encryption.</p> <p>See “Using encryption with Backup Exec” on page 583.</p>
Manage Keys	<p>Lets you manage your encryption keys.</p> <p>You can delete or replace existing encryption keys. You can also create a new encryption key.</p> <p>This option is available only if you select an encryption type.</p> <p>See “Encryption key management” on page 586.</p>

General options for DBA-initiated jobs

You can configure general options for DBA-initiated jobs.

See [“DBA-initiated job templates”](#) on page 597.

Table 14-8 General options for DBA-initiated jobs

Item	Description
Job name	Specifies the name for this backup template. You can accept the default name that appears or enter a name. The name must be unique.
Backup set description	Describes the information in the backup set for future reference.
Verify after backup completes	Performs a verify operation automatically to make sure that the media can be read after the backup has been completed. Verifying all backups is recommended.

Network options for DBA-initiated jobs

You can configure network options for DBA-initiated jobs.

See [“DBA-initiated job templates”](#) on page 597.

Note: Some of these options may not display in a CASO environment.

Table 14-9 Network options for DBA-initiated jobs

Item	Description
Network interface	Specifies the name of the network interface card that connects the Backup Exec server to the network that you want to use for this backup job. The list includes all available network interfaces on the Backup Exec server.
Protocol	Specifies the protocol you want to use for this backup job. The options are as follows: <ul style="list-style-type: none">■ Use any available protocol■ Use IPv4■ Use IPv6
Subnet	Displays the 32-bit number that determines the subnet to which the network interface card belongs.

Table 14-9 Network options for DBA-initiated jobs (*continued*)

Item	Description
Allow use of any available network interface, subnet, or protocol for Backup Exec agents not bound to the above network interface, subnet, or protocol	<p>Lets Backup Exec use any available network if the remote system that you selected for backup or restore is not part of the specified backup network.</p> <p>If you do not select this option and the remote system is not part of the specified backup network, the job fails. Backup Exec cannot access the data from the remote system.</p>
Interface Details	Displays the Media Access Control (MAC) address, adapter type, description, IP addresses, and subnet prefixes for the interface that you selected for the backup network.
Allow managed Backup Exec server to use any network interface to access Backup Exec agents	<p>Lets a job use any network interface to access Backup Exec agents if the selected network interface is unavailable.</p> <p>Enabling this option lets the managed Backup Exec server use an alternate network interface to run any important backup jobs that would otherwise fail.</p> <p>This option is available only if the Central Admin Server Option (CASO) is installed.</p> <p>See “About the Central Admin Server Option” on page 1096.</p>

Duplicate job settings for DBA-initiated jobs

You can configure duplicate job template settings for DBA-initiated jobs.

See [“DBA-initiated job templates”](#) on page 597.

Table 14-10 Duplicate job settings for DBA-initiated jobs

Item	Description
Enable settings to duplicate backup sets for this job	Enables the settings for a duplicate backup set template.
Storage	Specifies the storage device to which you want to send backup data for the duplicate DBA-initiated job.
Keep for	Designates the amount of time for which you want to keep the backup sets or job history from the duplicate DBA-initiated job.

Table 14-10 Duplicate job settings for DBA-initiated jobs (*continued*)

Item	Description
Media set	<p>Indicates the media set to use for the duplicate DBA-initiated job. The media set specifies the overwrite protection period and the append period for the backup data on the media.</p> <p>If you want to create a new media set for this backup job, click the icon to the right of the media set drop-down menu.</p> <p>See “Default media sets” on page 379.</p>
Overwrite media	<p>Indicates that the backup job is placed on an overwritable media. Ensure that appropriate media is in the storage device that you select.</p> <p>Appropriate media for an overwrite job include the following:</p> <ul style="list-style-type: none">■ Scratch media■ Media for which the overwrite protection period has expired <p>Depending on your configuration, overwritable media is selected from scratch media or recyclable media.</p> <p>If the media in the storage device is not overwritable, an alert prompts you to insert overwritable media.</p> <p>See “Managing tapes” on page 378.</p> <p>See “Media overwrite protection levels for tape media” on page 394.</p> <p>See “How Backup Exec searches for overwritable media in tape drives” on page 395.</p>
Append to media, overwrite if no appendable media is available	<p>Appends this backup job to the specified media set if an appendable media is available. Otherwise, Backup Exec searches for an overwritable media and adds it to the media set.</p> <p>If an append job fills a media, the backup job continues on an overwritable media. If the media in the storage device is not overwritable, an alert appears to prompt you to insert overwritable media.</p>
Append to media, terminate job if no appendable media is available	<p>Appends this backup job to the specified media set if an appendable media is available. Otherwise, Backup Exec terminates the job.</p>

Table 14-10 Duplicate job settings for DBA-initiated jobs (*continued*)

Item	Description
Eject media after job completes	<p>Ejects the media from the drive or slot when the operation completes. You can also schedule a job to eject media.</p> <p>See “Ejecting media from a disk cartridge or tape drive” on page 458.</p>
Retension media before backup	<p>Runs the tape in the drive from beginning to end at a fast speed. Retensioning helps the tape wind evenly and run more smoothly past the tape drive heads. This option is available only if you select a tape drive that supports retensioning.</p>
Use Write once, read many (WORM) media	<p>Specifies the use of WORM (write once, read many) media as the default for DBA-initiated jobs. Backup Exec confirms that the destination device is or contains a WORM-compatible drive, and that the WORM media is available in the drive. If WORM media or a WORM-compatible drive is not found, an alert is sent.</p> <p>See “How WORM media is used in Backup Exec” on page 401.</p>
Enable DirectCopy to tape	<p>Enables Backup Exec to coordinate the movement of data from virtual storage directly to a physical storage device.</p> <p>The Backup Exec server records information about the data in the catalog. Therefore, you can restore data from either the virtual storage or the physical storage.</p> <p>See “Copying data from a virtual tape library to a physical tape device using DirectCopy to tape” on page 229.</p>

Table 14-10 Duplicate job settings for DBA-initiated jobs (*continued*)

Item	Description
Compression	<p>Provides the following compression options:</p> <ul style="list-style-type: none">■ None Copies the data to the media in its original form (uncompressed). Using some form of data compression can help expedite backups and preserve storage space. Hardware data compression should not be used in environments where storage devices that support hardware compression are used interchangeably with devices that do not have that functionality. In this situation, hardware compression is automatically disabled. You can manually turn on hardware compression on the drives that support it, but this results in media inconsistency. If the drive that supports hardware compression fails, the compressed media cannot be restored with the non-compression drive.■ Software Uses STAC software data compression, which compresses the data before it is sent to the storage device.■ Hardware (if available, otherwise none) Uses hardware data compression if the storage device supports it. If the drive does not feature data compression, the data is backed up uncompressed.■ Hardware (if available, otherwise software) Uses hardware data compression if the storage device supports it. If the drive does not feature hardware data compression, STAC software compression is used.
Encryption type	<p>Specifies the encryption key that you want to use, if any.</p> <p>See “Using encryption with Backup Exec” on page 583.</p>
Encryption key	<p>Specifies the encryption key that you want to use, if you selected to use encryption.</p> <p>See “Using encryption with Backup Exec” on page 583.</p>
Manage Keys	<p>Lets you manage your encryption keys.</p> <p>You can delete or replace existing encryption keys. You can also create a new encryption key.</p> <p>This option is available only if you select an encryption type.</p> <p>See “Encryption key management” on page 586.</p>

Table 14-10 Duplicate job settings for DBA-initiated jobs (*continued*)

Item	Description
Preferred source device	Specifies the preferred source device that you want to use as the storage for the duplicate job.
Verify after backup completes	Performs a verify operation automatically to make sure that the data can be read after the backup has been completed. Verifying all backups is recommended.

Backup Exec logon accounts

A Backup Exec logon account stores the credentials of a user account that you use to access a computer. Backup Exec logon accounts enable Backup Exec to manage user names and passwords and can be used to browse computers or process jobs. Using Backup Exec logon accounts enables you to apply credential changes to the jobs that use them.

Backup Exec logon accounts are used to browse local and remote computers. Whenever the Backup Exec logon credentials are passed between the Backup Exec server and the remote computer, the credentials are encrypted.

Backup Exec logon accounts can also be associated with backup data at the device level such as shares, databases, etc. If you need to edit the credentials, you can edit the Backup Exec logon account. Any changes are applied to the selected computers that use the Backup Exec logon account.

Backup Exec logon accounts are not user accounts. When you create a Backup Exec logon account, an entry for the account is entered into the Backup Exec database; no operating system accounts are created. If your user account credentials change, you must update the Backup Exec logon account with the new information. Backup Exec does not maintain a connection with the user account.

You can view, create, edit, replace, and delete Backup Exec logon accounts.

The following types of logon accounts are included in Backup Exec:

[Default Backup Exec logon account](#)

[Backup Exec System Logon Account](#)

[Restricted logon accounts](#)

Default Backup Exec logon account

The default Backup Exec logon account enables you to browse, make selections, or restore data. The first time you start Backup Exec, you must specify a default

Backup Exec logon account using the Logon Account Wizard. You can select an existing Backup Exec logon account or create a new one.

You can create multiple Backup Exec logon accounts; however, each Backup Exec user can have only one default Backup Exec logon account.

Your default Backup Exec logon account enables you to perform the following:

- Browse data. Your default Backup Exec logon account enables you to browse local and remote computers when you create backup jobs. To browse computers, each user must have a default Backup Exec logon account that is associated with their user account. The Backup Exec logon account does not have to be the same user name as the user that is used to log on to Backup Exec. For example, you are logged on to a Backup Exec server named BACKUPSERVER as the local Windows administrator. When you start Backup Exec, you are prompted to create a default Backup Exec logon account for the local administrator because one does not exist. You can create a Backup Exec logon account for the local administrator that has the credentials for a domain administrator. The Backup Exec logon account has the following properties:
User name: DOMAIN\Administrator
Description: BACKUPSERVER\Administrator Default Account
Owner: BACKUPSERVER\Administrator
When you change your default Backup Exec logon account, you can use your new default Backup Exec logon account to browse computers immediately. You do not have to restart your system in order for the changes take effect.
See [“Changing your default Backup Exec logon account”](#) on page 618.
- Make backup selections. You can select a different Backup Exec logon account when you make selections for backup. If your default logon account does not have rights, the **Logon Account Selection** dialog box appears and lets you create or select a different Backup Exec logon account.
See [“Requirements for using the SQL Agent”](#) on page 892.
See [“Requirements for accessing Exchange mailboxes ”](#) on page 944.
- Restore. You can assign Backup Exec logon accounts to computers when you create restore jobs. The default Backup Exec logon account is used unless you choose a different Backup Exec logon account when you create the restore job.

Backup Exec System Logon Account

The Backup Exec System Logon Account (SLA) is created when you install Backup Exec. When the SLA is created, the user name and password match the credentials that were provided during install for the Backup Exec Services credentials. The owner of the SLA is the user that installed Backup Exec. It is a common account, by default. Common accounts are the shared accounts that all users can access.

See [“Creating a new Backup Exec System Logon Account”](#) on page 619.

The Backup Exec System Logon Account may have access to most or all of your data since it contains the Backup Exec Services credentials. If you want to make Backup Exec more secure, you can change the SLA to be a restricted account. You can also delete it after making another logon account the default. However, if you delete the SLA, the jobs in which it is used may fail. If the SLA is deleted, you can re-create it using the **Logon Account Management** dialog box.

The SLA is used for the following tasks and jobs:

- Jobs that were migrated from a previous version of Backup Exec
- Duplicate backup data jobs
- Command Line Applet (bemcli.exe)

Restricted logon accounts

Backup Exec logon accounts can be common or restricted. When you create a Backup Exec logon account, you can designate it as a restricted account. To use a restricted logon account, you must be the owner of the logon account or you must know the password for the logon account. The person who created the logon account is the owner. If you authorize only a few people to back up or restore data, you can make the logon account a restricted logon account.

The main reasons to restrict a logon account are as follows:

- To help you limit access to the computers available for backup.
- To help you limit the computers to which you can restore.

When you use a restricted logon account to select the data for a job, the logon account information is saved with the selection list. Anyone who tries to edit the job must provide the password to the restricted logon account. Backup Exec loads the selections for that job only when the password for the restricted logon account is provided.

See [“Creating a Backup Exec logon account”](#) on page 612.

See [“Editing a Backup Exec logon account”](#) on page 614.

See [“Changing the password for a Backup Exec logon account”](#) on page 616.

See [“Replacing a Backup Exec logon account”](#) on page 616.

See [“Deleting a Backup Exec logon account”](#) on page 617.

See [“Copying logon account information to another Backup Exec server”](#) on page 620.

Creating a Backup Exec logon account

You can create Backup Exec logon accounts using the Logon Account Wizard, which guides you through the creation of a Backup Exec logon account, or by using

the **Logon Account Management** dialog box. You can enter Backup Exec logon account property information when you create the Backup Exec logon account. However, Backup Exec assigns the Backup Exec logon account owner to the user name you used to log on to Backup Exec. The owner of the Backup Exec logon account cannot be modified.

This topic includes the following information:

[To create a Backup Exec logon account using the Logon Account Wizard](#)

[To create a Backup Exec logon account manually](#)

To create a Backup Exec logon account using the Logon Account Wizard

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Logon Accounts**, and then select **Logon Account Wizard**.
- 3 Click **Add a new logon account**, and then click **Next**.
- 4 Type a user name and password.
- 5 Click **Next**.
- 6 In the **Logon account name** field, type the unique name for the Backup Exec logon account.
- 7 Under **Make this account**, select whether you want the account to be a common logon account or a restricted logon account.

Common logon accounts are the shared accounts that all users can access. Restricted logon accounts can only be used by the owner of the logon account or by those who know the password.
- 8 If you want to make this the default logon account that is used to browse, make selections, and restore data on your local computers and remote computers, select **The default logon account**.
- 9 Click **Next**.
- 10 Review the options that you selected, and then click **Finish** to create the logon account.

To create a Backup Exec logon account manually

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Logon Accounts**, and then select **Manage Logon Accounts**.
- 3 Click **Add**.

- 4 In the **User name** field, type the fully qualified user name for the Backup Exec logon account.

For example, type "DOMAIN\Administrator".

The user name is provided when you attempt to connect to a computer. The user name is not case-sensitive for the computers that are accessed.

- 5 In the **Password** field, type the password for the account.

The password you enter is encrypted for security. You can leave this field blank if this Backup Exec logon account does not need a password.

- 6 In the **Confirm password** field, type the password again to verify it.

- 7 In the **Account name** field, type the unique name for the Backup Exec logon account.

- 8 In the **Notes** field, type any optional notes to explain how the Backup Exec logon account is used.

- 9 Select **This is a restricted logon account** if you want the Backup Exec logon account to be used only by the owner of the logon account and those who know the password.

If this option is not selected, the Backup Exec logon account is created as a common account. Common accounts are the shared accounts that all users can access.

- 10 Select **This is my default account** to make this account your default Backup Exec logon account, which is used to browse, make selections, or restore data on your local computers and remote computers.

See ["Backup Exec logon accounts"](#) on page 610.

Editing a Backup Exec logon account

When you edit a Backup Exec logon account, the changes are automatically applied to all the content that uses the Backup Exec logon account. Changes made to a Backup Exec logon account are applied immediately. You do not have to restart your system for the changes to take effect.

You can edit the following properties for a Backup Exec logon account:

- Type (restricted, common, or default)
- Account name
- Password
- User name
- Notes

To edit a Backup Exec logon account

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Logon Accounts**.
- 2 Select **Manage Logon Accounts**.
- 3 Select the Backup Exec logon account you want to change, and then click **Edit**.

If you are not logged on to Backup Exec with the same user name as the Backup Exec logon account owner, you must provide the password before you can edit the account.

- 4 Edit any of the following properties:

User name	Type the fully qualified user name for the Backup Exec logon account. For example, DOMAIN\Administrator. The user name is provided when you attempt to connect to a computer. The user name you enter is not case sensitive for the computers that are accessed.
Change Password	Click this option to change the password for the account. The password you enter is encrypted for security.
Account name	Type a unique name for the Backup Exec logon account. The user name is automatically added if you do not enter information into the field.
Notes	Type any optional notes indicating how the Backup Exec logon account is used.
This is a restricted logon account	Select this option to make this Backup Exec logon account a restricted logon account. Restricted logon accounts can be used only by the owner of the logon account and those who know the password. If this option is not selected, the Backup Exec logon account is a common account. Common accounts are the shared accounts that all users can access.
This is my default account	Select this option to make this account your default Backup Exec logon account. Your default account is used to browse, make selections, or restore data on your local computers and remote computers.

- 5 On the **Edit Logon Credentials** dialog box, click **OK**.
- 6 On the **Logon Account Management** dialog box, click **OK**.

See [“Backup Exec logon accounts”](#) on page 610.

Changing the password for a Backup Exec logon account

You can change a Backup Exec logon account password using the following steps. Changes made to a Backup Exec logon account password are applied immediately.

To change the password for a Backup Exec logon account

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Logon Accounts**.
- 2 Select **Manage Logon Accounts**.
- 3 Select the Backup Exec logon account that you want to change, and then click **Edit**.

If you are not logged on to Backup Exec with the same user name as the Backup Exec logon account owner, you must provide the password before you can edit the account.

- 4 Click **Change Password**.
- 5 In the **Password** field, type a new password.
- 6 In the **Confirm** field, re-type the password, and then click **OK**.
- 7 On the **Edit Logon Credentials** dialog box, click **OK**.
- 8 On the **Logon Account Management** dialog box, click **OK**.

See [“Backup Exec logon accounts”](#) on page 610.

Replacing a Backup Exec logon account

You can replace a Backup Exec logon account within all existing jobs. The data in existing jobs that use the Backup Exec logon account will be updated to use the new Backup Exec logon account. If the new Backup Exec logon account is restricted, you must provide the password.

To replace a Backup Exec logon account

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Logon Accounts**.
- 2 Select **Manage Logon Accounts**.
- 3 Select the Backup Exec logon account you want to replace, and then click **Replace**.

- 4 On the **Replace Logon Account** dialog box, select the Backup Exec logon account with which you want to replace the selected Backup Exec logon account.

If the Backup Exec logon account is restricted and you are not logged on to Backup Exec with the same user name as the Backup Exec logon account owner, you must provide the password before you can select the account.

- 5 Click **OK**.

See [“Backup Exec logon accounts”](#) on page 610.

Deleting a Backup Exec logon account

If you no longer need a Backup Exec logon account, you can delete it.

You cannot delete a Backup Exec logon account in the following situations:

- It is referenced by a job.
- It is owned by a user who is logged on to the Backup Exec server.
- It is set as the default Backup Exec logon account of a user who is logged on to the Backup Exec server.

If a logon account is used in any of these situations, you must replace it with a different logon account before you can delete it.

To delete a Backup Exec logon account

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Logon Accounts**.
- 2 Select **Manage Logon Accounts**.
- 3 Select the Backup Exec logon account you want to delete, and then click **Delete**.
- 4 Do either of the following:

If the logon account is not referenced by any backup jobs

Click **Yes** to confirm the deletion.

If the logon account is referenced by backup jobs

Complete the following steps to replace the logon account with a different logon account in any referenced backup jobs.

- On the **Delete Logon Account** dialog box, click **Replace**.
- On the **Replace Logon Account** dialog box, select the Backup Exec logon account with which you want to replace the selected Backup Exec logon account. If the Backup Exec logon account is restricted and you are not logged on to Backup Exec with the same user name as the Backup Exec logon account owner, you must provide the password before you can select the account.

The logon account that you select here replaces the previous logon account in all existing jobs and selection lists.

5 Click **OK**.

See [“Backup Exec logon accounts”](#) on page 610.

See [“Replacing a Backup Exec logon account”](#) on page 616.

Changing your default Backup Exec logon account

You can change your default Backup Exec logon account that enables you to browse, make selections, or restore data.

To change your default Backup Exec logon account

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Logon Accounts**.
- 2 Select **Manage Logon Accounts**.
- 3 Select the Backup Exec logon account that you want to use as your default Backup Exec logon account, and then do one of the following:
 - Click **Set as Default**.
 - Click **Edit**, select **This is my default account**, and then click **OK**.
- 4 Click **OK**.

See [“Backup Exec logon accounts”](#) on page 610.

Creating a new Backup Exec System Logon Account

The Backup Exec System Logon Account enables you to perform several operations. If you delete the Backup Exec System Logon Account, you should create a new one that enables you to perform the specified operations.

To create a new Backup Exec System Logon Account

- 1
- Click the Backup Exec button, select **Configuration and Settings**, and then select **Logon Accounts**.
- 2
- Select **Manage Logon Accounts**.
- 3
- Click **System Account**.
- 4
- Complete the following options:

User name	Type the fully qualified user name for the Backup Exec logon account. For example, DOMAIN\Administrator. The user name is provided when you attempt to connect to a computer. The user name you enter is not case-sensitive for the computers that are accessed.
Change Password	Click this option to change the password for the account. The password you enter is encrypted for security.
Account name	Type a unique name for the Backup Exec logon account. The user name is automatically added if you do not enter information into the field.
Notes	Type any optional notes indicating how the Backup Exec logon account is used.
This is a restricted logon account	Select this option to make this Backup Exec logon account a restricted logon account. Restricted logon accounts can be used only by the owner of the logon account and those who know the password. If this option is not selected, the Backup Exec logon account is a common account. Common accounts are the shared accounts that all users can access.
This is my default account	Select this option to make this account your default Backup Exec logon account. Your default account is used to browse, make selections, or restore data on your local computers and remote computers.

- 5
- Click **OK** to create the system logon account.
- See [“Backup Exec logon accounts”](#) on page 610.

Copying logon account information to another Backup Exec server

You can copy logon account information from one Backup Exec server to a different Backup Exec server.

To copy logon account information to another Backup Exec server

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Logon Accounts**.
- 2 Select **Manage Logon Accounts**.
- 3 Select the logon account that you want to copy, and then click **Copy to Servers**.
- 4 If prompted, enter the password for the logon account that you selected.
- 5 Do one of the following:
 - To add individual servers manually, in the **Server Name** field, enter the name of the Backup Exec server that you want to copy the logon account information to, and then click **Add**.
 - To add several servers from a list, click **Import List**, and then browse to the list of server names.
- 6 If you want to overwrite a logon account with the same name on the destination Backup Exec server, check **Overwrite logon account if one with this description already exists on the destination server**.
- 7 Click **OK**.

See [“Backup Exec logon accounts”](#) on page 610.

Testing logon accounts

You can test Backup Exec logon accounts to ensure that they can access your backup sources before you run jobs. Testing your logon accounts before you attempt to run jobs can help prevent failures and save you time.

To test logon accounts

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Logon Accounts**.
- 2 Select **Manage Logon Accounts**.
- 3 Click **Test**.
- 4 In the **Logon Account** field, select the logon account that you want to test.
- 5 In the **Server** field, select the server that you want to access with the logon account.

- 6 Click **Test**.

Backup Exec checks to make sure whether the logon account can access the server that you selected.

- 7 When you are finished testing logon accounts, click **Close**.

- 8 Click **OK**.

See [“Backup Exec logon accounts”](#) on page 610.

Starting and stopping Backup Exec services

You can use the Backup Exec Services Manager to start, stop, and restart Backup Exec services.

To start or stop Backup Exec services

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Services**.
- 2 Select the server for which you want to start or stop services.
- 3 Do any of the following:

To start all services for the selected server Click **Start all services**.

To stop all services for the selected server Click **Stop all services**.

To restart all services for the selected server Click **Restart all services**.

- 4 Click **OK**.

See [“Changing the credentials for a service account”](#) on page 621.

See [“Changing startup options for Backup Exec services”](#) on page 623.

Changing the credentials for a service account

On the Backup Exec server, all Backup Exec services run in the context of a user account that is configured for the Backup Exec system services.

Note: The Backup Exec service account and the Backup Exec system logon account are set to the same user name when Backup Exec is installed. If you need to change the user name for the service account or if the service account is no longer used, then you should also change the Backup Exec system logon account to use new credentials.

If this computer is in a domain, enter a Domain Administrators account, or an equivalent account that is part of the Domain Admins group. In the Domain list, select or enter the Domain name.

If this computer is in a workgroup, enter an Administrators account, or an equivalent account that is part of the Administrators group on the computer. In the Domain list, select or enter the computer name.

The account that you designate for Backup Exec services, whether it is a new account or an existing user account, is assigned the following rights:

- Authenticate as any user and gain access to resources under any user identity.
- Create a token object, which can then be used to access any local resources.
- Log on as a service.
- Administrative rights (provides complete and unrestricted rights to the computer).
- Backup operator rights (provides rights to restore files and directories).
- Manage auditing and security log.

See [“Required user rights for backup jobs”](#) on page 151.

Due to security implementations in Microsoft Small Business Server, the service account must be Administrator.

To change the credentials for a service account

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Services**.
- 2 On the **Backup Exec Services Manager** dialog box, select the appropriate server, and then select the service for which you want to change the service account.
- 3 Click **Edit credentials**.
- 4 Check the **Change service account credentials** check box.

5 Complete the following fields:

Old user name	Type the current user name for the service account that you want to change.
Old password	Type the current password for the service account that you want to change.
New user name	Type the new user name for the service account that you want to change.
New password	Type the new password for the service account that you want to change.
Confirm password	Type the new password again to confirm it.
Grant required rights to the service account	Select this option to grant the service account the proper system service rights.

6 Click **OK**.

7 Click **Close**.

See [“Starting and stopping Backup Exec services”](#) on page 621.

See [“Changing startup options for Backup Exec services”](#) on page 623.

Changing startup options for Backup Exec services

You can change startup options for Backup Exec services. Each individual service can be configured to start automatically or manually. Or you can disable a service entirely.

Services that are configured for automatic startup automatically start when the server starts. Services that are configured for manual startup do not start automatically. You must manually start services that are configured for manual startup. You can start, stop, or restart services in the Backup Exec Services Manager.

To change service startup options

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Services**.
- 2 On the **Backup Exec Services Manager** dialog box, select the appropriate server, and then select the service for which you want to change startup options.
- 3 Click **Edit credentials**.
- 4 Check the **Change startup options** check box.

5 Select from the following options:

Automatic	Select this option to automatically start the service account at system startup.
Manual	Select this option to prevent the service account from automatically starting at system startup. If you select this option, you must start the service account manually.
Disabled	Select this option to disable the service account at system startup.

6 Click **OK**.

7 Click **Close**.

See [“Starting and stopping Backup Exec services”](#) on page 621.

See [“Changing the credentials for a service account”](#) on page 621.

Configuring audit logs

You can use audit logs to examine and review information about the operations that have been performed in Backup Exec. The audit log displays the date and time of the activity, who performed it, what the activity was, and a description of the activity.

Audit logs can be configured to display information about the activities that occur for all or any of the following:

- Alerts
- Archive
- Audit logs
- Backup set retention
- Devices and media
- Encryption keys
- Error-handling rules
- Install
- Jobs
- Logon accounts
- Server configuration

You can delete the audit logs as part of the Backup Exec database maintenance, and you can save the audit log to a text file. Any changes that are made to the audit log, such as when database maintenance occurs, can also be displayed in the audit log.

To configure audit logs

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Audit Log**.
- 2 Click **Configure Logging**.
- 3 On the **Audit Log Configuration** dialog box, select the check box of the category that you want to display in the audit log.

Expand the category by clicking the arrow to the left of the category. Select the operations that you want to display for the category.

Clear the check box of any item or operation that you do not want to display.
- 4 Click **OK**.

See [“Viewing the audit log”](#) on page 625.

See [“Removing entries from the audit log”](#) on page 625.

See [“Saving an audit log to a text file”](#) on page 626.

Viewing the audit log

You can view audit logs to see when changes were made in Backup Exec and which users made the changes.

To view the audit log

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Audit Log**.
- 2 In the **Select category to view** field, select the category for which you want to view audit information.
- 3 Use the scroll bar at the bottom of the Audit Logs window to view the whole entry, or double-click the entry to display the same information in an easy-to-read Audit Log Record.

See [“Configuring audit logs”](#) on page 624.

Removing entries from the audit log

You can remove the entries for all categories or for a selected category.

To remove entries from the audit log

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Audit Log**.
- 2 In the **Select category to view** field, select the category for which you want to view audit information.
- 3 Click **Clear Category Log** to remove all entries from an audit log category.

If you select specific categories, only the logs that are generated for the selected categories are cleared when you click **Clear Category Log**.

See [“Configuring audit logs”](#) on page 624.

Saving an audit log to a text file

You can save the audit log as a text file.

To save the audit log to a text file

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Audit Log**.
- 2 Click **Save Log to File** to specify a file name and location to save the audit log entries.

See [“Configuring audit logs”](#) on page 624.

Copying configuration settings to another Backup Exec server

If you have the Central Admin Server Option (CASO), you can copy configuration settings and logon information from one Backup Exec server to another. This copy ability lets you quickly set up a group of Backup Exec servers with the same configuration or logon settings.

Note: To copy configuration settings and logon information to other Backup Exec servers, you must install the **Copy Server Configurations** feature.

See [“Installing additional Backup Exec options to the local Backup Exec server”](#) on page 76.

To copy configuration settings to another Backup Exec server

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then click **Copy Settings to Other Servers**.
- 2 Select any of the following options:

Default job options	Select this option to copy default job options from this Backup Exec server to another Backup Exec server.
Default schedule	Select this option to copy the default schedule settings from this Backup Exec server to another Backup Exec server.
Error-handling rules	Select this option to copy error-handling rules from this Backup Exec server to another Backup Exec server.
Alert configuration	Select this option to copy the alert configuration from this Backup Exec server to another Backup Exec server.

- 3 Select the Backup Exec server or servers to which you want to copy the selected settings.

If the destination server is not in the list, do one of the following:

- To add a server manually, click **Add**, and then select the server or servers that you want to add to the list.
- To import a list of servers, click **Import List**, browse to select the list, and then click **Open**.

- 4 Click **OK**.

See [“About the Central Admin Server Option”](#) on page 1096.

Viewing server properties

You can view properties for the local Backup Exec server or any other server that you monitor with Backup Exec.

You can view the following properties for the local Backup Exec server:

- Server name
- Server description
- Server status
- Version and license information
- Date and time zone information

- Operating system information
- Memory and page file information

Additionally, you can view the following properties for any server that you monitor with Backup Exec:

- Server name
- Server description
- Operating system information
- Backup Exec version and license information

If you have the Central Admin Server Option (CASO), you can also view information about the Backup Exec database, device and media database, and catalog database.

See [“Viewing the settings for a central administration server”](#) on page 1148.

To view server properties

- ◆ Complete either of the following as necessary:

To view the local Backup Exec server's properties

- Click the Backup Exec button, select **Configuration and Settings**, and then click **Local server properties**.
- When you are finished viewing the local server properties, click **OK**.

To view any other server's properties

- On the **Backup and Restore** tab, double-click the server whose properties you want to view.
- In the left pane, click **Properties**.

Reports

This chapter includes the following topics:

- [Reports in Backup Exec](#)
- [Running a report now](#)
- [Scheduling a report](#)
- [Creating a custom report](#)
- [Saving a report](#)
- [Printing a report from the Backup Exec Report Viewer](#)
- [Viewing completed reports](#)
- [Editing a report](#)
- [Re-running a completed report](#)
- [Deleting a report](#)
- [Setting defaults for standard and custom reports](#)
- [Viewing report properties](#)
- [List of Backup Exec standard reports](#)

Reports in Backup Exec

Backup Exec includes more than 40 standard reports that can provide detailed information about the alerts, devices, media, and jobs in your Backup Exec environment. In addition, Backup Exec provides the ability to create custom reports so you can create reports to fit your specific environment and needs. Both standard reports and custom reports can be run on demand at any time, can be scheduled

to run at a specific time, or can be scheduled to run on a recurring schedule. When you schedule a report job, you can set up email notification so that you or anyone else can be notified when the report job completes, and you can include a copy of the report in the email

Reports can be viewed and printed in the following formats:

- PDF
- HTML
- XML
- Microsoft Excel (XLS)
- Comma-separated Value (CSV)

Reports are grouped on the **Reports** tab by category. For example, reports relating to devices are grouped in the **Devices** report group. To see the names and descriptions of the standard reports for a group, click the name of the appropriate report group.

Note: The **See Completed** report group only includes the scheduled reports that have completed. Reports that were run immediately are deleted when you close the report viewer.

Requirements for reports

Before you attempt to run either standard reports or custom reports, review the following requirements:

- To properly format integrated Backup Exec reports, you must configure a default printer using the Windows Control Panel Printers applet. This is required even if you do not have a printer attached to your system. For information on configuring a printer by using the Windows Control Panel Printers applet, see your Microsoft Windows documentation.
- To run reports across multiple Backup Exec servers, you must install the Backup Exec Enterprise Server Option, even if you do not have a shared storage environment.

Reports and the Job Monitor

Reports cannot be monitored from the **Job Monitor**. Reports that are running, are scheduled to run, or are completed do not appear on the **Job Monitor**. All report operations are displayed on the **Reports** tab.

See [“Running a report now”](#) on page 631.

See [“List of Backup Exec standard reports”](#) on page 647.

Running a report now

When you run a report, you can specify the criteria that is used to determine the items that will be included in the report. The settings that are available for you to select depend on the type of data that can be included in the report. After the report is generated, only the items that match the criteria appear in the report.

If you do not want to run the report immediately, you should use the **New Scheduled Report** option instead.

See [“Scheduling a report”](#) on page 631.

To run a report

- 1 On the **Reports** tab, under **Report Groups**, click the report group that contains the report that you want to run now.
- 2 Right-click the report you want to run, and then click **Run Now**.
- 3 If the **Run Report Now - <report_name>** properties page appears, select the appropriate settings for the data that you want to include in the report, and then click **OK**.
- 4 After you have finished viewing the report, click **Close**.

Backup Exec automatically deletes the report when you close the Report Viewer.

Scheduling a report

You can schedule a report to run once at a specific time or multiple times on a recurring schedule. Scheduled reports are listed in the **See Upcoming** report group on the **Reports** tab. You can edit or delete scheduled reports.

To schedule a report

- 1 On the **Reports** tab, under **Report Groups**, click the report group that contains the report that you want to schedule.
- 2 Right-click the report name, and then click **Schedule Report**.
- 3 Type a name for the report.
- 4 If any of the following options appear in the left pane, click the option name, and then enter the criteria for the data that you want to include on the report. Note that some reports do not include any of the following options.
 - **Backed up servers**
 - **Ranges**
 - **Audit log category**

- **Job status**
 - **Media sets**
 - **Vault**
- 5 (Optional) If you want to send an email to yourself or someone else when the report is completed, do the following:
- In the left pane, click **Notification**.
 - Select an existing recipient or click **Manage Recipients** to add a new recipient.
 - If you want to include a copy of the report in the email notification, check **Include the report in email notifications**.
- 6 In the left pane, click **Schedule**, and then select one of the following scheduling options:

Recurrence	Lets you schedule the job to run multiple times. You can set the recurrence pattern so that the job runs
Run now with no recurring schedule	Lets you run the job immediately.
Run on	Lets you set the date and time to run the report.
Create without a schedule	Lets you create and save the report, but not schedule or run the report at this time. If you select this option, you can then use an external scheduling tool to schedule the job to run at a specific time, or you can manually run the job when you are ready. Even though the job is not scheduled, the report is listed in the See Upcoming report group on the Report tab.

- 7 Click **OK**.
- The scheduled report is saved in the **See Upcoming** report group. After the report runs, it is saved in the **See Completed** report group.
- See [“Viewing completed reports”](#) on page 643.

Creating a custom report

You can create reports that contain information to meet the specific requirements of your organization. You choose the data to include in the report, and then determine how the data is filtered, sorted, and grouped. In addition, you can set up a pie graph or a bar graph to graphically represent the report data.

Filters let you customize reports to include only the information that meets specific criteria. For example, you can use filters to find the jobs that contain a specific word, the alerts that occurred on a specific day, or the media that are in a specific location. You use filter criteria to create filter expressions. You can use one or multiple filter expressions. A filter expression consists of a field name, an operator, and a value.

The following example filter expression finds all alerts for errors:

Table 15-1 Filter expression for finding alerts for errors

Filter type	Data
Field name	Alert Type
Operator	= (Equal)
Value	Errors

If you want the report to include only the alerts for errors that occurred on a specific day, add another filter expression for the date and time, as shown in the example below:

Table 15-2 Filter expression for finding alerts on a specific day

Filter type	Data
Field name	Date Entered
Operator	=(Equal)
Value	06/03/2014 <time>

Grouping fields creates sections on the report. For example, if you group by Backup Exec server, Backup Exec creates a section for each Backup Exec server that matches the filter criteria. Under each Backup Exec server's section, the report displays the data that corresponds to the remaining fields that you selected for the report.

You can sort a custom report by up to three of the fields that you have chosen for the report. When you sort on fields, Backup Exec arranges all of the data that matches the sort criteria together in the report. For example, if you sort on the

Backup Exec server field in ascending order, all data for Backup Exec server A displays first, followed by all data for Backup Exec server B, and so on.

To create a custom report

- 1 On the **Reports** tab, click **New Custom Report**.
- 2 On the **Custom Report** dialog box, type a name and description for the report.
- 3 If you do not want this report to include the default header and footer settings, uncheck **Use header and footer settings specified in Backup Exec Settings**.

The default header and footer settings can include a customized logo, a custom color for the banner, and text for the footer. These items are set in the default Backup Exec settings.

See [“Setting defaults for standard and custom reports”](#) on page 645.

- 4 In the left pane, click **Field Selection**.
- 5 In the **Category** box, select a group for which you want to create a report.
- 6 For additional field selections, click **Show advanced fields**.
- 7 In the **Available fields** list, select the fields that you want to include on the report, and then click the Right arrow (>>) button to move the fields to the **Fields selected for the report** list.
- 8 After you have moved all of the fields that you want to include on the report to the **Fields selected for the report** list, arrange the order in which you want the fields to appear.

The fields appear on the report in the order in which they are listed in the **Fields selected for the report** list, with the first field appearing as the first field on the left of the report. To move a field, select it from the **Fields selected for the report** list, and then click **Move Up** or **Move Down** to move it to the appropriate location in the list.

- 9 To adjust the width of the column for a field, do the following in the order listed:
 - Click the field name in the **Fields selected for the report** list.
 - In the **Column width** field, type the new width.
 - Click **Set**.
- 10 (Optional) If you want to filter the data on the report, do the following:
 - In the left pane, click **Filters**.
 - In the **Field name** list, select the field on which you want to filter data.
 - In the **Operator** list, select the operator that you want to use for this filter.

- In the **Value** field, type or select the specific data that you want to include on the report.
- Click **Add**.
- To combine sets of filter expressions, do any of the following:

To combine two filter expressions so that both expressions must be true for the result to be true

Click **AND**.

For example, to find all backup jobs that failed, add the following expressions:

- Status = Failed
- Type = Backup

After you set up the expressions, do the following:

- Click **AND** to combine the two expressions.

The combined expression is:

Status = Failed AND Type = Backup

To combine two filter expressions so that one of the expressions must be true for the result to be true

Click **OR**.

For example, to find jobs that either failed or were canceled, add the following expressions:

- Status = Failed
- Status = Canceled

After you set up the expressions, do the following:

- Click **OR** to combine Status = Failed with Status = Canceled.

The combined expression is:

Status = Failed OR Status = Canceled

To combine two filter expressions into a single expression

Click () +

For example, to find backup jobs and restore jobs that failed, add the following expressions:

- Status = Failed
- Type = Backup
- Type = Restore

After you set up the expressions, do the following:

- Use OR to combine Type = Backup with Type = Restore.
- Press and hold Ctrl while you click Type = Backup and Type = Restore.
- Click () + to combine Type = Backup with Type = Restore.
- Use AND to combine Status = Failed with (Type = Backup OR Type = Restore).

The combined expression is:

Status = Failed AND (Type = Backup OR Type = Restore)

To separate two filter expressions that were combined into a single expression

Click () -

For example, if you used () + to combine Type = Backup with Type = Restore, it is displayed on the **Filters** dialog box as follows:

(Type = Backup OR Type = Restore)

To make the combined expression into two individual expressions, do the following:

- Press and hold Ctrl while you click both Type = Backup and Type = Restore.
- Click () -

After you separate the expressions, they are displayed without the parentheses.

11 (Optional) If you want to organize the report into sections, do the following:

- In the left pane, click **Grouping**.
- In the **Group by** list, select the field that you want to use as a group.
- Click **Ascending** or **Descending**.
Ascending order lists numbers from smallest to largest and lists letters in alphabetical order. Descending order lists numbers from largest to smallest and lists letters in reverse alphabetical order.
- If you want to further group the data, select the fields in the **Then group by** lists, and then click **Ascending** or **Descending** for those fields.

A report must have at least one field that is not grouped. For example, if you select three fields to include on the report, you can group only two of the fields. If you group all of the fields, no data appears on the report because all of the data is listed in the group section titles. In addition, you must have at least four fields on the report to use all three grouping fields.

- 12 (Optional) If you want to sort the data on the report, do the following:
 - In the left pane, click **Sort**.
 - In the **Sort on** list, select the field on which you want to sort the data in the report.
 - Select **Ascending** or **Descending**.
Ascending order lists numbers from smallest to largest and lists letters in alphabetical order. Descending order lists numbers from largest to smallest and lists letters in reverse alphabetical order.
 - If you want to further sort the data, select the fields on which you want to sort in the **Then sort on** lists, and then click **Ascending** or **Descending** for those fields.
- 13 (Optional) If you want to add a pie graph or bar graph to the report, do the following:
 - In the left pane, click **Graph Options**.
 - In the **Graph type** list, select either **Bar** or **Pie**.
 - Select the types of data that you want to include on the graph.
- 14 (Optional) If you want to see what the report will look like before you save it, in the left pane, click **Preview**.
- 15 Click **OK** to save the custom report.

Adding or removing fields on a custom report

To change the data that appears on a custom report, add new fields or remove existing fields.

To add or remove fields on a custom report

- 1 On the **Reports** tab, under **Report Groups**, click **Custom**.
- 2 Right-click the report that contains the fields you want to change, and then click **Edit**.
- 3 In the left pane, click **Field Selection**.
- 4 Do any of the following:

To add new fields to the report	<ul style="list-style-type: none"> ■ Select a category. ■ Under Available fields, select the fields you want to add, and then click the right arrow (>>) button.
To remove fields from the report	Under Fields selected for the report , click the field you want to remove, and then click the left arrow (<<) button.

Changing filters for a custom report

Use the following steps to change the existing filters on a custom report.

To change filters for a custom report

- 1 On the **Reports** tab, under **Report groups**, click **Custom Reports**.
- 2 In the reports list, right-click the report that you want to change and then click **Edit**.
- 3 In the left pane, click **Filters**.
- 4 Create a filter by defining one or more filter expressions.

To add a new filter expression	Select a field name and operator, and then enter a value. Click Add .
To edit an existing filter expression	Do the following in the order listed. <ul style="list-style-type: none"> ■ Under Filter criteria, select the filter expression that you want to edit, and then click Edit. ■ Under Filter expression, edit the expression's values. ■ Under Filter expression, click Update.
To delete a filter expression	Under Filter criteria , select the filter expression that you want to delete, and then click Remove .

- 5 To combine sets of filter expressions, do any of the following:

To combine two filter expressions so that both expressions must be true for the result to be true

Click **AND**.

For example, to find all backup jobs that failed, add the following expressions:

- Status = Failed
- Type = Backup

After you set up the expressions, do the following:

- Click AND to combine the two expressions.

The combined expression is:

Status = Failed AND Type = Backup

To combine two filter expressions so that one of the expressions must be true for the result to be true

Click **OR**.

For example, to find jobs that either failed or were canceled, add the following expressions:

- Status = Failed
- Status = Canceled

After you set up the expressions, do the following:

- Click OR to combine Status = Failed with Status = Canceled.

The combined expression is:

Status = Failed OR Status = Canceled

To combine two filter expressions into a single expression

Click () +

For example, to find backup jobs and restore jobs that failed, add the following expressions:

- Status = Failed
- Type = Backup
- Type = Restore

After you set up the expressions, do the following:

- Use OR to combine Type = Backup with Type = Restore.
- Press and hold Ctrl while you click Type = Backup and Type = Restore.
- Click () + to combine Type = Backup with Type = Restore.
- Use AND to combine Status = Failed with (Type = Backup OR Type = Restore).

The combined expression is:

Status = Failed AND (Type = Backup OR Type = Restore)

To separate two filter expressions that were combined into a single expression

Click () -

For example, if you used () + to combine Type = Backup with Type = Restore, it is displayed on the **Filters** dialog box as follows:

(Type = Backup OR Type = Restore)

To make the combined expression into two individual expressions, do the following:

- Press and hold Ctrl while you click both Type = Backup and Type = Restore.
- Click () -

After you separate the expressions, they are displayed without the parentheses.

6 Click **OK**.

Changing the way data is grouped or sorted in a custom report

Use the following steps to change the way data is grouped or sorted in a custom report.

To change the way data is grouped or sorted in a custom report

- 1 On the **Reports** tab, under **Report Groups**, click **Custom**.
- 2 In the list of custom reports, right-click the report that you want to change, and then select **Edit**.
- 3 Do any of the following:

To change the field that is used as a group

- In the left pane, click **Grouping**.
- On the **Group by** or **Then group by** list that contains the field you want to change, click the down arrow, and then select the new field to use as a group.

To remove a group

- In the left pane, click **Grouping**.
- On the **Group by** or **Then group by** list that contains the group you want to remove, click the down arrow, and then click **<None>**.

To change the order in which data is grouped

Click **Ascending** or **Descending**.

- To change the field that is used to sort the data
 - In the left pane, click **Sorting**.
 - On the **Sort on** or **Then sort on** list that contains the sorting option you want to change, click the down arrow, and then select the new field to use to sort the data.
- To disable sorting
 - In the left pane, click **Sorting**.
 - On the **Sort on** or **Then sort on** list that contains the sorting option you want to remove, click the down arrow, and then click **<None>**.
- To change the order in which data is sorted Click **Ascending** or **Descending**.

4 Click **OK**.

Changing graph options in custom reports

You can include a pie graph or a bar graph in custom reports.

At least two fields must be selected on the **Field Selection** dialog box to create a pie graph, and at least three fields must be selected to create a bar graph.

To change graph options in custom reports

- 1 On the **Reports** tab, under **Report Groups**, click **Custom**.
- 2 Right-click the report you want to edit, and select **Edit**.
- 3 In the **Custom Report** window, click **Graph Options**.
- 4 Change the graph title or select new fields to populate the graph.
- 5 Click **OK**.

Previewing custom reports

Use the preview feature to verify that you created a custom report correctly.

To preview custom reports

- 1 On the **Reports** tab, under **Report Groups**, click **Custom**.
- 2 Right-click the report that you want to preview, and then click **Edit**.
- 3 In the left pane, click **Preview**.

Copying a custom report

You can make one or more copies of a custom report. Each copy of the custom report resides in the Custom report group, along with the original custom report.

To copy a custom report

- 1 On the **Reports** tab, under **Report Groups**, click **Custom**.
- 2 Right-click a custom report that you want to copy, and then click **Copy**.
- 3 Type a name for the report, and then click **OK**.

The copy of the custom report appears in the **Custom** report group.

Saving a report

Reports can be saved to any location that you choose on your hard drive or network in any of the following formats:

- HTML file (.htm)
- Adobe PDF file (.pdf)
- XML file (.xml)
- Comma-separated value file (.csv)
- Microsoft Office Excel workbook (.xls)

You can save a report that is currently displayed on the screen or that is in the **See Completed** report group.

To save a report

- 1 On the report, in the Report Viewer, click **Save As**.

To save a report that is currently displayed on the screen

On the report, in the Report Viewer, click **Save As**.

To save a report that is in the See Completed report group

- On the **Reports** tab, under **Report Groups**, click **See Completed**.
- Click **Save As**.
- Double-click a completed report that you want to save to a new location.

- 2 Enter the file name and location where you want to save the report.

- 3 In the **Save as type** box, select a format in which to save the report.
When you save a report in HTML format, both the HTML file and a .GIF image file are saved.
- 4 Click **Save**.

Printing a report from the Backup Exec Report Viewer

You can print reports from a locally-attached printer or a network printer. To print a report, the printer must be configured to print in the landscape mode.

The following printer settings ensure that the report prints correctly:

- On the **Layout** tab of the **Print** dialog box, under **Orientation**, **Landscape** should be selected. Note that you may need to select **Preferences** to access the **Layout** tab.
- To print all of the pages in a multiple page report, on the **Options** tab of the **Print** dialog box, the option **Print all linked documents** should be selected.
- To print all of the pages in a multiple-page report in the correct order, the first page of the report must be displayed on the screen before the print job is initiated.
- If the report does not print correctly, you may need to change the page setup options in Internet Explorer to remove the header and footer and reduce the margins.

To print a report from the Backup Exec Report Viewer

- 1 Run a report.
See [“Running a report now”](#) on page 631.
- 2 On the **Report Viewer**, click **Print**.
- 3 Select a printer from the Windows **Print** dialog box.
- 4 Click **Print**.

Viewing completed reports

After a scheduled report runs, it is saved in the **See Completed** report group.

To view completed reports

- 1 On the **Reports** tab, under **Report Groups**, click **See Completed**.
- 2 Double-click the report that you want to view.

Editing a report

Use the following steps to edit the properties of a scheduled standard report or a custom report before it runs. If the report that you want to edit has been run in a previous report job, the changes you make now may affect the appearance of the reports in job history. Symantec recommends that you copy the report and then edit the copy.

To edit a scheduled standard report or a custom report

- 1 On the **Reports** tab, under **Report Groups**, click **See Upcoming**.

To edit a scheduled standard report

On the **Reports** tab, under **Report Groups**, click **See Upcoming**.

To edit a custom report

On the **Reports** tab, under **Report Groups**, click **Custom**.

- 2 Right-click a report you want to edit, and then click **Edit**.
- 3 Edit the report properties and then click **OK**.

Re-running a completed report

You can run the reports that appear in the **See Completed** report group multiple times.

To re-run a completed report

- 1 On the **Reports** tab, under **Report Groups**, click **See Completed**.
- 2 Right-click a report, and then click **Retry Report Now**.

Backup Exec creates and runs another iteration of the report.

- 3 To view the report again, double-click the new report.

Deleting a report

Reports that you create using the **Run now** option are automatically deleted after you view the report. However, custom reports, completed reports, and scheduled reports can be deleted at your convenience.

Note: Standard Backup Exec reports cannot be deleted.

To delete a report

- 1 On the **Reports** tab, under **Report Groups**, do one of the following:
 - To delete a custom report, click **Custom**.
 - To delete a scheduled report, click **See Upcoming**.
 - To delete a completed report, click **See Completed**.
- 2 Right-click the report that you want to delete, and then click **Delete**.
- 3 Click **Yes** to confirm that you want to delete the report.

Setting defaults for standard and custom reports

You can set Backup Exec to display all reports in either HTML or Adobe Portable Document Format (PDF). The default setting is HTML. The format that you select does not affect the format of the reports that are sent to users with the notification feature.

For custom reports, you can do the following:

- Include a logo in the header.
- Choose a color for the banner in the header.
When you choose a color for the banner, you can type the numbers that correspond to the colors (RGB values), or you can select the color from a chart.
- Include text in the footer.
- Include the time in the footer.

To set defaults for standard and custom reports

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, click **Reports**.

3 Complete the appropriate options.

To select a default report format for standard and custom reports

Under **Report format**, select either **HTML** or **PDF**.

To limit the number of rows that can be included in standard and custom reports

Under **Report content**, in **Maximum number of rows to include in a report**, type the appropriate number.

To enable standard and custom reports to show all data that is generated, even if some of the data is duplicated

Under **Report content**, click **Show all rows**.

To enable standard and custom reports to show only unique data

Under **Report content**, click **Show distinct rows**.

To add a logo to the header section of custom reports

Under **Header**, check **Use custom image file**, and then enter the path to the image that you want to use in the Image file path field.

To customize the colors in the header section of custom reports

Under **Banner color**, either enter the numbers that correspond to the colors you want to use, or click **Colors** to select a color from a chart.

To include default text or the time in the footer of custom reports

Under **Footer**, in the **Text** field, type the default text that you want to appear on every custom report. If you want the time of the report to be included in the footer, verify that **Include time** is checked.

4 Click **OK**.

Viewing report properties

Report properties provide detailed information about each report, such as the file name, file size, and report creation date. The properties can be viewed, but not edited.

To view report properties

- 1 On the **Reports** tab, under **Report Groups**, select a report group.
- 2 Right-click a report for which you want to view properties, and then click **Properties**.
- 3 Click **OK** after you have finished viewing the properties.

List of Backup Exec standard reports

This section provides detailed information about each standard report that is available in Backup Exec. The data that appears on each report varies depending on the criteria you selected to include in the report.

The following standard reports are included in Backup Exec:

Table 15-3 Backup Exec standard reports

Report Name	Description
Alert History	Lists all alerts in the alert history chronologically, displaying the most recent alerts first See “Alert History report” on page 652.
Alert History by Backup Exec server	Lists all alerts in the alert history, grouped and filtered by Backup Exec server, displaying the most recent alerts first. See “Alert History By Backup Exec Server report” on page 653.
Archive Job Success Rate	Displays the number of archive jobs for the backed up servers that successfully ran. See “Archive Job Success Rate report” on page 653.
Archive Selections by Archive Rules and Retention Categories	Displays the archive rules and the retention categories that are applied to each archive selection. See “Archive Selections by Archive Rules and Retention Categories report” on page 654.
Audit Log	Lists the contents of the audit logs for selected servers for the specified time period. See “Audit Log report” on page 655.
Backup Job Success Rate	Lists the success rate for backup jobs run to back up selected servers. See “Backup Job Success Rate report” on page 655.
Backup Recommendations	Lists any recommendations that can help you to better manage your backups. See “Backup Recommendations report” on page 655.

Table 15-3 Backup Exec standard reports (*continued*)

Report Name	Description
Backup Resource Success Rate	Lists the success rate for backup jobs for specified past number of days for resources on selected servers. See “Backup Resource Success Rate report” on page 656.
Backup Sets by Media Set	Lists all backup sets by media set. See “Backup Sets by Media Set report” on page 656.
Backup Size by Resource	Lists the backup size for each resource job for up to seven previous runs and then computes the trailing average for up to seven previous runs for each job run. See “Backup Size By Resource report” on page 657.
Daily Device Utilization	Lists the percentage of the storage devices' capacity that the Backup Exec server uses. See “Daily Device Utilization report” on page 658.
Deduplication device summary	Displays a summary of the deduplication operations for local deduplication disk storage and shared deduplication disk storage. See “Deduplication Device Summary report” on page 659.
Deduplication summary	Displays a deduplication summary for all of the deduplication jobs that run on the Backup Exec server. See “Deduplication Summary report” on page 660.
Device Summary	Lists the device usage and error summary for each selected Backup Exec server. See “Device Summary report” on page 660.
Disk Storage Summary	Displays disk-based usage statistics for Backup Exec server disk storage. See “Disk Storage Summary report” on page 661.
Error-Handling Rules	Lists all the defined error-handling rules. See “Error-Handling Rules report” on page 662.

Table 15-3 Backup Exec standard reports *(continued)*

Report Name	Description
Event Recipients	Lists all events that were received by each notification recipient. See “Event Recipients report” on page 663.
Exchange Mailbox Group Archive Settings	Displays the archive settings that are applied to mailbox groups in each domain. See “Exchange Mailbox Group Archive Settings report” on page 664.
Failed Archive Jobs	Displays what archive jobs failed recently. See “Failed Archive Jobs report” on page 664.
Failed Backup Jobs	Lists all the failed backup jobs, sorted by the resource server and time frame. See “Failed Backup Jobs report” on page 665.
File System Archive Settings	Displays the archive settings that are applied to archive selections for each server. See “File System Archive Settings report” on page 666.
Jobs Summary	Lists all the jobs that ran within the last 72 hours in chronological order. See “Jobs Summary report” on page 666.
Managed Backup Exec Servers	Lists the status and configuration for all Backup Exec servers that are managed by a central administration server. See “Managed Backup Exec Servers report” on page 667.
Media Audit	Lists the recent media configuration changes. See “Media Audit report” on page 669.
Media Errors	Lists the number of errors that occur on all media. See “Media Errors report” on page 670.

Table 15-3 Backup Exec standard reports (*continued*)

Report Name	Description
Media Required for Recovery	<p>Lists the media that contain the backup sets for each system that is backed up on selected servers for the specified time period. This report can be inaccurate if media overwrite settings allow the media to be overwritten.</p> <p>See “Media Required for Recovery report” on page 670.</p>
Media Summary	<p>Lists all the media sets and media that are used by Backup Exec servers. The current location is given for each media. Also lists usage statistics for media and the location of media within Backup Exec media sets.</p> <p>See “Media Summary report” on page 671.</p>
Media Vault Contents	<p>Lists the media that are located in each media vault.</p> <p>See “Media Vault Contents report” on page 671.</p>
Move Media to Vault	<p>Lists all media that can be moved to a media vault. The listed media are not currently in a media vault and the media’s append period has expired.</p> <p>See “Move Media to Vault report” on page 672.</p>
Operations Overview	<p>Lists past and future operations data for user-set period.</p> <p>See “Operations Overview report” on page 673.</p>
Overnight Archive Summary	<p>Displays the status of the archive jobs that ran in the last 24 hours.</p> <p>See “Overnight Archive Summary report” on page 675.</p>
Overnight Summary	<p>Lists the results of backup jobs for each resource during the last 24 hours. This report includes backup jobs that were scheduled to run but did not run. Jobs are given a grace period of 24 hours before they are marked as past due.</p> <p>See “Overnight Summary report” on page 676.</p>
Problem Files	<p>Lists all the problem files that are reported for jobs. The files are grouped by day and resource.</p> <p>See “Problem Files report” on page 677.</p>

Table 15-3 Backup Exec standard reports *(continued)*

Report Name	Description
Recently Written Media	Lists all media that have been modified in the last 24 hours. See “Recently Written Media report” on page 678.
Resource Protected Recently	Lists all job detail statistics and exceptions that occurred on a Backup Exec server for which you run this report. See “Resource Protected Recently report” on page 678.
Resource Risk Assessment	Lists job information for resources on which the last backup job run on the resource failed. The data is filtered by resource server. See “Resource Risk Assessment report” on page 679.
Restore Set Details by Resource	Lists all restore sets that ran within the last 72 hours. The sets are grouped by the server and resource. See “Restore Set Details by Resource report” on page 680.
Retrieve Media from Vault	Lists all reusable media currently in the specified vault. See “Retrieve Media from Vault report” on page 680.
Robotic Library Inventory	Lists the contents of slots in robotic libraries that are attached to Backup Exec servers. Usage statistics are provided for each piece of media. See “Robotic Library Inventory report” on page 681.
Scheduled Server Workload	Lists the estimated scheduled workload for the next 24-hour period by server. See “Scheduled Server Workload report” on page 682.
Scratch Media Availability	Lists the aging distribution of media. Shows how many media are available for overwrite and when other media will become available for overwrite. See “Scratch Media Availability report” on page 682.
Test Run Results	Lists the results for the test run jobs that are set for the selected time period and Backup Exec servers. See “Test Run Results report” on page 683.

Table 15-3 Backup Exec standard reports (*continued*)

Report Name	Description
Vault Store Usage Details	Displays the archives that are in each store and the size of each archive. See “Vault Store Usage Details report” on page 684.
Vault Store Usage Summary	Displays the archived items that are in each vault store and the total size of the vault store. See “Vault Store Usage Summary report” on page 685.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Alert History report

The Alert History report lists all the alerts in the Alert History chronologically, displaying the most recent alerts first.

Table 15-4 Alert History report

Item	Description
Time	Date and time the alert occurred.
Received	Time the alert occurred.
Responded	Time when the user responded to the alert.
Responding User	User that responded to the alert.
Job Name	The name of the job that is associated with the alert.
Backup Exec server	Name of the Backup Exec server on which the alert occurred.
Category	Title of the alert, such as Service Start or Job Failed.
Message	A description of the event that caused the alert.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Alert History By Backup Exec Server report

The Alert History by Backup Exec server report lists all alerts in the alert history, grouped and filtered by Backup Exec server, displaying the most recent alerts first.

Table 15-5 Alert History by Backup Exec Server report

Item	Description
Backup Exec server	Name of the Backup Exec server on which the alert occurred.
Time	Date and time the alert occurred.
Received	Time the alert occurred.
Responded	Time when the user responded to the alert.
Responding User	User that responded to the alert.
Job Name	Name of the job that is associated with the alert.
Category	Title of the alert, such as Service Start or Job Failed.
Message	Describes the event that caused the alert.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Archive Job Success Rate report

The Archive Job Success Rate report displays the number of archive jobs for the backed up servers that successfully ran.

Table 15-6 Archive Job Success Rate report

Item	Description
Date	Displays the date on which archive jobs ran.
Total Jobs	Displays the total number of archive jobs that have run.
Successful	Displays the total number of archive jobs that are successful.
Success Rate	Displays the success rate of the archive jobs in percentages.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Archive Selections by Archive Rules and Retention Categories report

The Archive Selections by Archive Rules and Retention Categories report displays the archive rules and the retention categories that are applied to each archive selection.

Table 15-7 Archive Selections by Archive Rules and Retention Categories report

Item	Description
Archive Rule	Displays the archive rule that you specify to identify the files and mail messages that are eligible for archiving.
Archive Selection	Displays only the archive selection. In the case of an NTFS archive, the network path appears. In the case of an Exchange mailbox archive, the mailbox group appears along with information about the mailbox group selections.
Archive Type	Displays the type of data that is being archived Archive types include: <ul style="list-style-type: none">■ File System Archive■ Mailbox Archive
Windows Domain	Displays the Windows domain in which the archived selection resides.
Retention Category	Displays the retention category that applies to the file system selections in the archive job. A retention category specifies the time period for which you want to keep archived items.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Audit Log report

The Audit Log report lists the contents of the audit logs for the selected servers for the selected time period.

Table 15-8 Audit Log report

Item	Description
Category	Category in which the change occurred, such as Logon Account, Alerts, or Job.
Date Entered	Time and date the change occurred.
Message	Description of the change that was made in Backup Exec.
User Name	User that made the change.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Backup Job Success Rate report

The Backup Job Success Rate report lists the success rate for backup jobs.

Table 15-9 Backup Success Rate report

Item	Description
Server	Name of the server being backed up.
Date	Date the backup job was processed.
Total Jobs	Total number of jobs that were processed by the Backup Exec server.
Successful	Total number of jobs successfully performed by the Backup Exec server.
Success Rate	Percentage of successful jobs that were processed by the Backup Exec server.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Backup Recommendations report

The Backup Recommendations report lists any recommendations that can help you to better manage your backups. The recommendations may include better

ways to back up specific types of data or suggestions for how to use Symantec products.

Table 15-10 Backup Recommendations report

Item	Description
Backup Exec server	Name of the Backup Exec server for which the recommendation applies.
Job Name	Name of the job that is associated with the recommendation.
Start Time	Date and time when the job that is associated with the recommendation ran.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Backup Resource Success Rate report

The Backup Resource Success Rate report lists the success rate for backup jobs for a specific number of days for resources on selected servers.

Table 15-11 Backup Resource Success Rate report

Item	Description
Resource	Name of the system being backed up.
Date	Date the backup job was processed.
Total Jobs	Total number of jobs that were processed by the Backup Exec server.
Successful	Total number of jobs successfully performed by the Backup Exec server.
Success Rate	Percentage of successful jobs that were processed by the Backup Exec server.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Backup Sets by Media Set report

The Backup Sets by Media Set report lists all the backup sets by media set.

Table 15-12 Backup Sets by Media Sets report

Item	Description
Media Set	Name of the media set on which the job ran.
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned barcode label.
Set	Sequential number for backup sets on the media.
Method	Specific type of backup.
Date / Time	Date and time the data was backed up.
Backup Set Description / Source	Describes the data that was backed up and the location of the data.
Directories	Number of directories that were backed up.
Files	Number of files that were backed up.
Size	Amount of data that was backed up.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Backup Size By Resource report

The Backup Size By Resource report lists the backup size for each resource job for up to seven previous jobs. It also computes the trailing average, which is the average of the amount of data that was backed up in the seven previous jobs.

Table 15-13 Backup Size by Resource Job report

Item	Description
Server	Name of the Backup Exec server where the data for the backup job was located.
Resource	Name of the resource that was backed up.
Job	Name of the backup job.
Job Date and Time Run	Date and time the backup job was processed.
Backup Size	Amount of data that was backed up.

Table 15-13 Backup Size by Resource Job report (*continued*)

Item	Description
Trailing Avg	Average amount of data that was backed up during the seven previous runs.
Difference %	Amount by which the data that was backed up in the current job differs from the data that was backed up in the previous backup jobs.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Daily Device Utilization report

The Daily Device Utilization report lists the percentage of the storage devices' capacity that the Backup Exec server uses.

Table 15-14 Daily Device Utilization report

Item	Description
Drive Name	Name of the storage device and the Backup Exec server where the device is located.
Status	<p>Status of the storage device</p> <p>Statuses for storage devices are as follows:</p> <ul style="list-style-type: none"> ■ Pause The storage device is temporarily stopped. ■ Enable The storage device is available for use with Backup Exec. If the storage device is disabled, it is available for use with other applications. ■ Online The storage device is available for use. ■ Offline Backup Exec cannot access the storage device.
Date	Date the storage device was used.
Jobs	Number of jobs that were processed by the Backup Exec server's storage device.
Size	The amount of data that was processed by the Backup Exec server's storage device.

Table 15-14 Daily Device Utilization report (*continued*)

Item	Description
Utilization (%)	Percentage of device utilization.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Deduplication Device Summary report

The Deduplication device summary report displays a summary of the deduplication operations for local deduplication disk storage and shared deduplication disk storage.

Table 15-15 Deduplication device summary report

Item	Description
State	Device state, such as online and enabled.
Created	Date media was created.
Total Capacity	Total capacity of the deduplication disk storage.
Used Capacity	Capacity presently used by the deduplication disk storage.
Available Capacity	Remaining capacity of the deduplication disk storage.
Percent Full	Percentage of storage space that is available in the deduplication disk storage.
Protected Bytes	Total amount of data that is selected for backup in all jobs using the device before deduplication occurs.
Deduplication Ratio	Ratio of the amount of data before deduplication to the amount of data after deduplication.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Deduplication Summary report

The Deduplication Summary report displays a deduplication summary for all of the deduplication jobs that run on the Backup Exec server.

Table 15-16 Deduplication Summary report

Item	Description
Job Name	Name of the job.
Start Time	Time of day that Backup Exec attempted to start the job.
Duration	Length of time the operation took to process.
Size	The amount of data processed.
Size/Minute	Number of kilobytes, megabytes, or gigabytes processed per minute.
Scanned Byte Count	Total amount of data that is selected for backup before deduplication occurs.
Stored Byte Count	The amount of unique data is stored after deduplication occurs.
Deduplication Ratio	Ratio of the amount of data before deduplication to the amount of data after deduplication.
Status	Status of the operation, such as Completed (Success), Failed, or Canceled.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Device Summary report

The Device Summary report lists all the devices for each selected Backup Exec server.

Table 15-17 Device Summary report

Item	Description
Server	Name of the server where the device is located.
Drive Name	Name of the drive in the robotic library.

Table 15-17 Device Summary report (*continued*)

Item	Description
Target	Address of the storage device that is connected to the Backup Exec server.
State	Device state, such as online.
Created	Date on which the media was created.
Cleaned	Date on which the last cleaning job was run on the drive.
Hours	Hours the device has been in use since the last cleaning job.
Errors	Number of errors occurring since the last cleaning job.
Size	Amount of data that was read and written since the last cleaning job.
Mounts	Number of mounts occurring since the last cleaning job.
Hours	Total number of hours the device has been in use.
Errors	Total number of errors occurring on the device.
Size	Amount of data that was read and written to the device.
Mounts	Total number of mounts occurring to the device.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Disk Storage Summary report

The Disk Storage Summary report displays disk usage statistics for Backup Exec server disk storage.

Table 15-18 Disk storage summary report

Item	Description
Device Name	Name of the disk storage device.

Table 15-18 Disk storage summary report (*continued*)

Item	Description
State	<p>State of the device.</p> <p>Device states include the following:</p> <ul style="list-style-type: none"> ■ Online ■ Enabled ■ Offline ■ Paused ■ Disabled
Local Access Path	Path on the disk where backup data is stored.
Total Capacity	Total capacity of the disk.
Used Space	Amount of disk space being used as storage.
Free Space	Amount of disk space remaining.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Error-Handling Rules report

The Error-Handling Rules report lists all error-handling rules and provides details about each rule.

Table 15-19 Error Handling Rules report

Item	Description
Rule Name	Name of the Error-Handling rule.
Notes	Information that was entered in the Notes section when the error-handling rule was created.
Job Status	<p>Final job status that activates the rule.</p> <p>Possible statuses are as follows:</p> <ul style="list-style-type: none"> ■ Error ■ Canceled

Table 15-19 Error Handling Rules report (*continued*)

Item	Description
Error Category	<p>Category of error for which the rule will be applied.</p> <p>Available error categories include the following:</p> <ul style="list-style-type: none"> ■ Device ■ Job ■ Media ■ Network ■ Other ■ Resource ■ Security ■ Server ■ System
Enabled	Displays if the rule is enabled or disabled.
Cancel Job	Displays an X if this option is selected for the error-handling rule. The option cancels all jobs after the maximum number of retries have been attempted.
Pause Job	Displays an X if this option is selected for the error-handling rule. The option enables Backup Exec to pause the job until you can manually clear the error.
Retry Job	Displays an X if this option is selected for the error-handling rule. The option enables Backup Exec to retry the job.
Maximum Retries	Number of times the job is to be retried.
Retry Interval (minutes)	Number of minutes Backup Exec waits before trying to run the job again.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Event Recipients report

The Event Recipient report lists events that were registered for each notification recipient.

Table 15-20 Event Recipients report

Item	Description
Recipient Name	Name of the recipient.
Recipient type	Designates to whom the Event Recipients report is sent, such as an individual recipient or a group of recipients.
Event Type	Alert category or ad hoc job.
Event Name	Detail for the alert category or ad hoc job.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Exchange Mailbox Group Archive Settings report

The Exchange Mailbox Group Archive Settings report displays the archive settings that are applied to mailbox groups in each domain.

Table 15-21 Exchange Mailbox Group Archive Settings

Item	Description
Windows Domain	Displays the name of the Windows domain in which the Exchange server belongs.
Mailbox Group	Displays the name of the mailbox group to be archived.
Archive Rules	Displays the archive rule that is used to archive the mailbox group.
Retention Category	<p>Displays the retention category that applies to the mailbox group selections in the archive job.</p> <p>A retention category specifies the time period for which you want to keep archived items.</p>

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Failed Archive Jobs report

The Failed Archive Jobs report displays what archive jobs failed recently.

Table 15-22 Failed Archive Jobs report

Item	Description
Start Time	Displays the time that the archive job started.
Duration	Displays the amount of time that the archive job took to run.
Job Name	Displays the name of the archive job.
Category	Displays the status of the failed archive job.
Error Code	Displays the error code for the error that caused the archive job to fail.
Description	Displays the description of the error that caused the archive job to fail.
Status	Displays the category for the error that may be generated due to system, job, media, or device issues
Device Name	Displays the name of the storage device that processed the archive job.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Failed Backup Jobs report

The Failed Backup Jobs report lists all the failed backup jobs. The jobs are sorted by the server and specified time frame.

Table 15-23 Failed Jobs report

Item	Description
Resource	Name of the system being backed up.
Start Time	Date and time the backup job started.
Duration	Length of time the operation took to process.
Job Name	Name of job that failed.
Category	Category for the failed job that may be generated by a system, job, media, or device error.

Table 15-23 Failed Jobs report (*continued*)

Item	Description
Error Code	Displays the error code that corresponds to the failure.
Description	Describes the event that caused the error.
Status	Status of the operation, such as Error.
Device Name	Name of the device on which the job ran.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

File System Archive Settings report

The File System Archive Settings report displays the archive settings that are applied to archive selections for each server.

Table 15-24 NTFS Archive Settings report

Item	Description
Server	Displays the name of the Windows server from where the data was archived.
Resource	Displays the path of the resource.
Archive Rules	Displays the archive rule that is used to archive the files.
Vault Store	Displays the name of the vault store where the archived files reside.
Retention Category	<p>Displays the retention category that applies to the file selections in the archive job.</p> <p>A retention category specifies the time period for which you want to keep archived items.</p>

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Jobs Summary report

The Jobs Summary report lists all jobs that have run within the specified time range. The jobs are listed in chronological order.

Table 15-25 Jobs Summary report

Item	Description
Start Time	Date and time the operation started.
Job Name	Name of the completed job.
Duration	Length of time that the operation processed.
Size	The amount of data processed.
Files	Number of files processed.
Directories	Number of directories processed.
Size/Minute	Number of kilobytes, megabytes, or gigabytes processed per minute.
Skipped	Number of files that were skipped during the operation.
Corrupt Files	Number of corrupt files that were encountered during the operation.
Files in Use	Number of files in use during the operation.
Status	Status of the operation, such as Completed (Success), Failed, or Canceled.
Type	Lists the type of job that Backup Exec ran within the specified time range.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Managed Backup Exec Servers report

The Managed Backup Exec Servers report lists status and configuration information for the managed Backup Exec servers in a Central Admin Server Option environment.

Table 15-26 Managed Backup Exec Servers report

Item	Description
CAS Servers	Names of the central administration servers.
Managed Backup Exec server	Names of the managed Backup Exec servers.

Table 15-26 Managed Backup Exec Servers report (*continued*)

Item	Description
Status	<p>Status of the server.</p> <p>Possible statuses include the following:</p> <ul style="list-style-type: none"> ■ Online - available for use. ■ Stalled - not responding immediately to messages ■ No Comm - communications to the server have been lost for some period of time.
Stalled	Time limit that was used for determining Stalled communications status.
No Comm	Time limit that was used for determining No Comm communications status.
Catalog Location	<p>Location of the catalog information.</p> <p>Possible locations are as follows:</p> <ul style="list-style-type: none"> ■ Local - the catalog information is located on the managed Backup Exec server. ■ CASO - the catalog information is located on the central administration server.
Logs	<p>When job logs are uploaded from the managed server to the CASO database.</p> <p>Possible upload times are as follows:</p> <ul style="list-style-type: none"> ■ Timed basis in seconds ■ Scheduled time ■ Completion of job ■ Never
History	<p>When job history is uploaded from the managed server to the CASO database.</p> <p>Possible upload times are as follows:</p> <ul style="list-style-type: none"> ■ Timed basis in seconds ■ Scheduled time ■ Completion of job ■ Never

Table 15-26 Managed Backup Exec Servers report (*continued*)

Item	Description
Status	<p>When status is uploaded from the managed server to the CASO database.</p> <p>Possible upload times are as follows:</p> <ul style="list-style-type: none"> ■ Timed basis in seconds ■ Scheduled time ■ Completion of job ■ Never
Display Alert	Displays Yes if you have configured an alert to be set if the time difference between the central administration server's clock and a managed Backup Exec server's clock exceeds a preset value (maximum time difference tolerance).
Sec	Maximum time difference tolerance in seconds that is set for server.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Media Audit report

The Media Audit report lists the recent configuration changes that you made to your media.

In a Central Admin Server Option (CASO) environment, if you run this report from the central administration server the report provides data only for the media for the central administration server; it does not provide any data for any of the managed Backup Exec servers. To obtain the media audit data for a managed Backup Exec server, you can do one of the following:

- Log on locally to the managed Backup Exec server and run the report from it.
- Use the Remote Administrator feature to log on to the managed Backup Exec server from a remote Windows server or workstation, and then run the report.

Table 15-27 Media Audit report

Item	Description
Date Entered	Time and date the change occurred.
Message	Description of the change that was made to the media.
User Name	User that made the change.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Media Errors report

The Media Errors report lists the number of errors that occur on all media.

Table 15-28 Media Errors report

Item	Description
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned barcode label.
Total Mounts	Total number of times this media has been mounted.
Total In Use Hours	Total number of hours that this media has been in use.
Total Errors	Total number of error alerts for the system, jobs, media, and devices.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Media Required for Recovery report

The Media Required for Recovery report lists the media that contain the backup sets for each system that was backed during the specified time period. However, this report may be inaccurate if media overwrite settings allow the media to be overwritten.

Table 15-29 Media Required for Recovery report

Item	Description
Date	Date and time the backup job set was created.
Media Location Name	Name of the storage device where the media that was used for the backup job is stored.
Media Label	Media label that is assigned to the media.
Backup Method	Specific type of backup. See “Backup methods in Backup Exec” on page 188.
Recycle Time	Displays the date and the time after which the media can be overwritten.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Media Summary report

The Media Summary report lists all media sets and media that are used by Backup Exec servers. Usage statistics are given for each piece of media.

Table 15-30 Media Summary report

Item	Description
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned barcode label.
Media Type	Type of media cartridge, such as 4mm.
Allocated	Date media was allocated to a media set as a result of an overwrite operation.
Modified	Date data was last written to the media.
Hours	Total number of hours that the media has been in use.
Mounts	Total number of times the media has been mounted.
Soft Errors	Number of recoverable read errors encountered.
Hard Errors	Number of unrecoverable read errors encountered.
Write Size	Amount of data that has been written to the media.
Current Size	Estimate of the amount of data currently on the media.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Media Vault Contents report

The Media Vault Contents report lists all the media in a specified media vault.

Table 15-31 Media Vault Contents report

Item	Description
Vault Name	Location of the media.

Table 15-31 Media Vault Contents report (*continued*)

Item	Description
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned barcode label.
Overwrite Protection End Date	Date that data on the media may be overwritten.
Vault Media Rule Move Date	Date media can be moved to vault.
Media Set	Name of media set to which the media belongs.
Vault Media Rule Name	Name of vault media rule.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Move Media to Vault report

Lists all media that you can move to a media vault.

The media listed are not currently in a media vault and meet one of the following criteria:

- The media has met or exceeded the vault move date that was specified for the media containing the media.
- The append period has expired, but the overwrite protection period is still current (allocated).

Table 15-32 Move Media to Vault report

Item	Description
Backup Exec server	Name of the Backup Exec server where the data for the backup job was located.
Media Set	Name of the media set.
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned barcode label.
Location	Location of the media.

Table 15-32 Move Media to Vault report (*continued*)

Item	Description
Append Period End Date	Last date that data may be added to the media.
Overwrite Protection End Date	Date that data on the media may be overwritten.
Vault Media Rule Move Date	Date media can be moved to vault.
Vault Name	Name of vault to which media is to be moved.
Vault Media Rule Name	Name of vault media rule.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Operations Overview report

The Operations Overview report lists details for past and future Backup Exec operations.

Table 15-33 Operations Overview report

Item	Description
Job summary for jobs completed in the past x Hours	Details Backup Exec job activity for the specified time period.
Errors	Total number of error alerts for the system, jobs, media, and devices.
Warnings	Total number of warning alerts for of jobs, media, and devices.
Information	Total number of informational alerts for the system, jobs, media, and devices.
Attention Required	Total number of alerts that require a response from the user.
Completed (Failed)	Total number of jobs that failed.

Table 15-33 Operations Overview report (*continued*)

Item	Description
Completed (Canceled)	Total number of canceled jobs.
Completed (Success)	Total number of jobs that completed successfully.
Exceptions	Total number of jobs that completed successfully, but may contain one or more skipped files, corrupt files, virus infected files or files in use.
Total Data Backed Up	Total amount of data that was backed up in kilobytes, megabytes, or gigabytes.
Total Media Used	Total number of media that were used to back up the completed jobs.
Missed	Total number of missed jobs.
Recovered	Total number of recovered jobs.
Active Jobs	Total number of active jobs.
Scheduled Jobs	Displays those jobs whose scheduled start times begin within 72 hours of the job being created. Jobs with recurring schedules also appear if their start times begin within 72 hours of the job's last start time.
Jobs On Hold	Total number of jobs on hold.
Job Status	The status of the jobs.
Scratch Media	Total number of scratch media available.
Recyclable	Total number of recyclable media available.
Imported	Number of imported media. Imported media is media that was created by a product other than this installation of Backup Exec.
Allocated	Number of allocated media (media belonging to a user media set).
Total Overwritable Media	Total number of overwritable media available.
Total Appendable Media	Total number of appendable media available.
Media Overwrite Protection Level	Displays the level of overwrite protection (Full, Partial, None) that is assigned to the media.
Online Devices	Total number of online devices.

Table 15-33 Operations Overview report (*continued*)

Item	Description
Offline Devices	Total number of offline devices.
Disabled Devices	Total number of disabled devices.
Paused Devices	Total number of paused devices.
Disabled	Lists the name of the devices that are disabled.
Paused	Name of the paused devices.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Overnight Archive Summary report

The Overnight Archive Summary report displays the status of the archive jobs that ran in the last 24 hours.

Table 15-34 Overnight Archive Summary report

Item	Description
Resource	Displays the name of the server that you are protecting
Type	Displays the type of job that ran in the last 24 hours.
Start Time	Displays the date and the time that the archive operation started.
Status	Displays the status of the archive operation.
Error Category	Displays the category for the error that may be generated due to system, job, media, or device issues.
Backup Exec server	Displays the name of the Backup Exec server on which the job ran.
Device Name	Displays the name of the device on which the job ran.
Total Tasks	Displays the total number of archive jobs that have run during the preceding 24 hours.

Table 15-34 Overnight Archive Summary report (*continued*)

Item	Description
Uncorrected Exceptions	Displays the number of archive jobs that failed because the error condition was never corrected and were not run again with successful completion.
Service Level	Displays the percentage of jobs that ran successfully.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Overnight Summary report

The Overnight Summary report lists the results of backup jobs for each resource during the last 24 hours. This report includes backup jobs that were due to run but did not run. Jobs are given a grace period of 24 hours before being marked as past due.

Table 15-35 Overnight Summary report

Item	Description
Resource	System being backed up.
Type	Displays the type of job that Backup Exec runs to produce the Overnight Summary report. Because the Overnight Summary report lists the results of backup jobs for each resource during the past 24 hours, Backup is always the type of job that appears
Start time	Date and time the operation started.
Status	Status of the operation.
Error Category	Category for the job that may be generated by a system, job, media, or device error.
Backup Exec server	Name of the Backup Exec server on which the job ran.
Device Name	Name of the device on which the job ran.
Total Tasks	Total number of jobs that ran within the last 24 hours.

Table 15-35 Overnight Summary report (*continued*)

Item	Description
Uncorrected Exceptions	<p>Number of the jobs that fail and were not run again with successful completion.</p> <p>Some of the archive jobs that ran in the past 24 hour period encountered exceptions. You must resolve the exceptions. Otherwise, the jobs that fail because of the exceptions continue to appear during subsequent 24 hour periods until the exceptions are resolved.</p>
Service Level	Percentage of jobs that ran successfully.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Problem Files report

The Problem Files report lists all the problem files that were reported for jobs. The files are grouped by day and resource.

Table 15-36 Problem Files report

Item	Description
Date	Date the problem file was encountered.
Resource	System on which the problem file is located.
Time	Time the problem file was encountered.
Reason	Error code that is listed in the job log summary.
File Name	Name of the problem file.
Type	Lists the type of job that Backup Exec ran when problematic files were detected.
Backup Exec Server	Name of the Backup Exec server on which the file is located.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Recently Written Media report

The Recently Written Media report lists all the media that has been modified within the specified period.

Table 15-37 Recently Written Media report

Item	Description
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned barcode label.
Location	Location of the media, such as the storage vault name or drive name.
Set	Name of backup set.
Date and Time Modified	Date and time media was last modified.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Resource Protected Recently report

The Resource Protected Recently report lists all job detail statistics and exceptions that occurred on a Backup Exec server for which you run this report.

Table 15-38 Resource Protected Recently report

Item	Description
Start Time	Date and time the backup job started.
Duration	Amount of time that was required for the job to complete.
Size	Amount of data that was backed up.
Files	Number of files that were backed up.
Directories	Number of directories that were backed up.
Size per Min	Amount of data that was backed up per minute.
Skipped	Number of files that were skipped during the backup.

Table 15-38 Resource Protected Recently report (*continued*)

Item	Description
Corrupt Files	Number of corrupt files that were detected during the backup.
Files in Use	Number of files that were in use during the backup.
Status	Status of the backup job.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Resource Risk Assessment report

The Resource Risk Assessment report shows job information for resources on which the last backup job that was run on the resource failed.

Table 15-39 Resource Risk Assessment report

Item	Description
Resource	System on which the job ran.
Error Text	Describes the event that caused the job to fail.
Start Time	Time the operation started.
Job	Name of the job that failed.
Error Category	The category for the failed job that may be generated by a system, job, media, or device error.
Backup Exec Server	Name of the Backup Exec server on which the job ran.
Device Name	Name of the device on which the job ran.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Restore Set Details by Resource report

The Restore Set Details by Resource report lists all restore jobs that ran within the specified time range on a selected server. The jobs are grouped by the server and resource.

Table 15-40 Resource Set Details by Resource report

Item	Description
Resource	Name of the system being backed up.
Start Time	Date and time the operation started.
Duration	Length of time the operation took to process.
Size	The amount of data processed.
Files	Number of files processed.
Directories	Number of directories processed.
Data/Minute	Amount of data that was processed per minute.
Skipped	Number of files that were skipped during the operation.
Corrupt Files	Number of corrupt files that were encountered during the operation.
Files in Use	Number of files in use during the operation.
Status	Status of the operation, such as Completed.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Retrieve Media from Vault report

The Retrieve Media from Vault report lists all reusable media currently in a specified media vault.

Table 15-41 Retrieve Media from Vault report

Item	Description
Cartridge Label	<p>Displays the name of the disk cartridge. Disk cartridge names cannot exceed 128 characters.</p> <p>You can rename the disk cartridge.</p> <p>See “Editing disk cartridge properties” on page 329.</p>

Table 15-41 Retrieve Media from Vault report (*continued*)

Item	Description
Vault Name	Displays the name of the vault where the media is located.
Media Set Name	Displays the name of the media set.
Offsite Return Date	Displays the date that the media was returned to the off-site vault.
Recycle Date	Displays the date after which the media can be overwritten.
No Append Date	Displays the date on which Backup Exec can no longer append data to the media.
Rule Name	Displays the name of vault media rule that is applied to the media.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Robotic Library Inventory report

The Robotic Library Inventory report lists the contents of slots in robotic libraries that are attached to Backup Exec servers. Usage statistics are provided for each piece of media.

Table 15-42 Robotic Library Inventory report

Item	Description
Server	Name of the server where the robotic library is located.
Device Name	Name of the robotic library.
Slot	Sequential number of the slot in the robotic library.
Media Label	Media label assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned barcode label.
State	State of operation of the slot: paused, disabled, enabled, offline, or online.
Modified	Date the media in the slot was last accessed.
Write	Number of bytes that have been written to this media.
Full	Space available on a media; "1" indicates that media is full and "0" indicates that there is space available on the media.

Table 15-42 Robotic Library Inventory report (*continued*)

Item	Description
Hours	Total number of hours this media has been in use.
Mounts	Total number of times this media has been mounted.
Append	The time remaining in the media's append period.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Scheduled Server Workload report

The Scheduled Server Workload report displays the estimated scheduled workload for a server during the next 24-hour period or a user-defined time period. The report only displays recurring jobs that have already run at least one time, not jobs scheduled to run once.

Table 15-43 Scheduled Server Workload report

Item	Description
Backup Exec server	Name of the Backup Exec server that will process the scheduled jobs.
Job	Name of the job that is scheduled to run.
Next Due Date	Time and day the next job is scheduled to run.
Backup Size	Estimated amount of data to be processed during the next 24 hours.
Total Size	Total amount of data to be processed on the server during the next 24 hours.
Total Size	Total amount of data to be processed on all Backup Exec servers.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Scratch Media Availability report

The Scratch Media Availability report shows the aging distribution of media, how many media are available for overwrite, and when other media will become available for overwrite.

Table 15-44 Scratch Media Availability report

Item	Description
Cartridge Label	Cartridge label that was assigned by Backup Exec, assigned by the administrator, or contained on a pre-assigned barcode label You can rename the cartridge. See “Editing disk cartridge properties” on page 329.
Media Location Name	Name of the storage device that contains the actual media.
Total Capacity	Total native capacity of the scratch media without using compression.
Append Hours Remaining	Capacity of scratch media available for append.
Remaining Capacity	Total amount of remaining native capacity of the scratch media without compression.
Retention Hours Remaining	The amount of time remaining to retain and protect the media from being overwritten.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Test Run Results report

The Test Run Results report displays the results for the test run jobs that are set for the selected period.

Table 15-45 Test Run Results report

Item	Description
Backup Exec server	Name of the Backup Exec server on which the job ran.
Job Date and Time Run	Date and time the backup job was processed.
Job Name	Name of the test run job.
Backup Sets	Name of the backup set.
Credential Check	Indicates if the Backup Exec logon account was verified as correct for the resources being backed up.

Table 15-45 Test Run Results report (*continued*)

Item	Description
Backup Size	Size in kilobytes, megabytes, or gigabytes of the backup.
Media Type	Type of media used, such as 4mm.
Device Name	Name of the device, such as the name of the robotic library.
Max Needed	Amount of space that is needed on the media to run the job.
Online	Capacity of media available in the device to which data can be appended.
Media Total	Total amount of appendable media available to the system.
Online	Capacity of media available in the device to which data can be written.
Media Total	Total amount of overwritable media available to the system.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Vault Store Usage Details report

The Vault Store Usage Details report displays the archives that are in each store and the size of each archive.

Table 15-46 Vault Store Usage Details report

Item	Description
Vault Store	Displays the name of the vault store where the Backup Exec archives are stored.
Archive Name	Displays the name that the Archiving Option gives to the archive.
Archive Type	Displays the type of data that you are archiving. Archive types include: <ul style="list-style-type: none"> ■ File System Archive ■ Mailbox Archive
Number of Archived Items	Displays the number of archived items that are in the vault store.

Table 15-46 Vault Store Usage Details report (*continued*)

Item	Description
Total Size	Displays the total size of the archived items in the vault store.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Vault Store Usage Summary report

The Vault Store Usage Summary report displays the archived items that are in each vault store and the total size of the vault store.

Table 15-47 Vault Store Usage Summary report

Item	Description
Vault Store	Displays the name of the disk-based vault store where the Backup Exec archives are stored.
Database Name	Displays the name of the vault store database that contains the configuration data and information about each of the archives in the partition.
Vault Store Open Partition	Displays the name of the vault store open partition where the Backup Exec archives are stored.
Vault Store Partition Free Size	Displays the amount of available free space in a vault store open partition.
Number of Archives in Vault Store	Displays the total number of existing Backup Exec archives in the vault store.
Total Size	Displays the total size of the existing Backup Exec vault store archives.

See [“Running a report now”](#) on page 631.

See [“Creating a custom report”](#) on page 633.

Troubleshooting Backup Exec

This chapter includes the following topics:

- [Troubleshooting hardware-related issues in Backup Exec](#)
- [How to get more information about alerts and error messages](#)
- [Troubleshooting backup issues in Backup Exec](#)
- [Troubleshooting failed components in the SAN](#)
- [How to improve Backup Exec's performance](#)
- [Accessing Symantec Online](#)
- [Searching the Symantec Knowledge Base](#)
- [Contacting Backup Exec Technical Support](#)
- [Using Symantec Remote Assistance](#)
- [Managing your Backup Exec support cases](#)
- [About Backup Exec diagnostic tools](#)
- [Running the Symantec Help Tool \(SymHelp\)](#)
- [Generating a diagnostic file for troubleshooting Backup Exec](#)
- [Running the begather utility to troubleshoot Backup Exec components on Linux servers](#)
- [Using the Backup Exec Debug Monitor for troubleshooting](#)
- [About the Backup Exec debug tool](#)

Troubleshooting hardware-related issues in Backup Exec

For common hardware-related issues, review the following frequently asked questions:

Table 16-1 Hardware-related questions

Question	Answer
How do I know if my storage device is supported?	<p>You can find a list of compatible devices at the following URL:</p> <p>http://entsupport.symantec.com/umi/V-269-2</p> <p>If your drive is listed on the hardware compatibility list, run the Configure Storage wizard and install Symantec device drivers.</p> <p>See “Installing Symantec tape device drivers” on page 360.</p> <p>From the Configure Storage wizard, the Symantec Device Driver Installation Wizard finds and installs the most suitable driver for your tape drive.</p>
How can I troubleshoot issues with a robotic library or tape drive?	<p>You can find troubleshooting tips and a video tutorial about configuring and troubleshooting tape drive hardware at the following URL:</p> <p>http://www.symantec.com/business/support/index?page=content&id=TECH24414</p> <p>See “Starting and stopping Backup Exec services” on page 621.</p> <p>See “Deleting a storage device” on page 450.</p>

Table 16-1 Hardware-related questions (*continued*)

Question	Answer
I'm getting an error "Storage device [device] reported an error on a request to read or write data to or from media. Error reported: Data error (cyclic redundancy check)." What should I do?	

Table 16-1 Hardware-related questions (*continued*)

Question	Answer
	<p>Many factors can cause the cyclic redundancy check (CRC) error.</p> <p>The following list contains the most common reasons for this error and potential ways to resolve the problem:</p> <ul style="list-style-type: none"> ■ Contaminated read and write heads of the tape device. Check with the hardware manufacturer for proper cleaning techniques. ■ Bad media. Replace the media. Try a new tape that the hardware manufacturer has certified. ■ Tape driver. Load the appropriate Backup Exec tape driver. You can find a list of compatible devices at the following URL: http://entsupport.symantec.com/urni/V-269-2 ■ Wide negotiation for the SCSI controller is not configured properly. If the device is a wide (68 pin) SCSI device, then wide negotiation may and should be used. If the device is a narrow (50 pin) SCSI device, disable wide negotiation. Use the manufacturer's SCSI installation program to disable wide negotiation on the SCSI controller card. ■ SCSI controller transfer rate is too fast. Use the manufacturer's SCSI installation program to lower the SCSI transfer rate. Check with the manufacturer of the controller and the device for the proper configuration for the SCSI transfer rate. ■ SCSI controller synchronous negotiation enabled. Use the manufacturer's SCSI installation program to disable synchronous negotiation on the SCSI

Table 16-1 Hardware-related questions (continued)

Question	Answer
	<p>controller card. Check with the manufacturer of the controller and the device for the proper configuration for SCSI synchronous negotiation.</p> <ul style="list-style-type: none">■ Incorrect termination or bad cables. Verify that the SCSI cable is good and that it is configured to provide proper SCSI termination. Do not mix passive and active termination.■ Confirm that the tape drive functions properly. Check with the tape drive manufacturer for diagnostic software to test the condition of the tape drive hardware.■ General SCSI problems. Isolate the tape drive on its own controller card or try a different SCSI card.
Why does my DLT tape drive pause when it catalogs some tapes?	<p>The DLT tape drive maintains internal information about the tape on a tape directory track. The directory track is updated before the tape is ejected from the drive. If the drive is turned off without ejecting the tape first, this information is lost.</p> <p>Re-generating the tape directory information takes several hours to complete, which makes it seem like the drive is hung. Allow sufficient time for the operation to complete and then eject the tape. Normal operation resumes after the directory track is updated.</p>

Table 16-1 Hardware-related questions (continued)

Question	Answer
A backup to my DLT tape drive is stuck at 99% complete. What should I do?	<p>The backup most likely fails to complete because the storage option Eject the media after job completes is selected, and the tape drive does not support the operation. Some tape drives require you to manually remove the tape, such as Digital Linear Tape (DLT), Linear Tape-Open (LTO), Travan, and Onstream drives.</p> <p>To remedy this situation, either uncheck the storage option Eject the media after job completes, or configure an automatic response to the active alert.</p> <p>See “Configuring network options for backup jobs” on page 204.</p>

How to get more information about alerts and error messages

Backup Exec generates an error message when a condition occurs that is important enough to warrant your attention, or requires that you submit a response. Most alerts and error messages are self explanatory, but there may be times when you need to get more information to resolve a condition.

You can get more information on Backup Exec alert and error messages in the following ways:

- On the alert message, click the link for more information, or look in the job log and click the UMI link. This code is a hyperlink to the Symantec Technical Support website. You can access the technical notes that are related to the alert.
See [“Linking from the job log to the Symantec Technical Support website”](#) on page 270.
- Search the Symantec Technical Support knowledge base for the error.
See [“Searching the Symantec Knowledge Base”](#) on page 699.

Troubleshooting backup issues in Backup Exec

If you have problems with backing up data, review the following questions.

Table 16-2 Backup questions

Question	Answer
I am unable to back up certain files on my system that are in use by other processes. Why is that?	<p>For non-snapshot backups, when Backup Exec encounters a file that is in use by another process, it either skips the file or waits for the file to become available. These actions depend on the options for no-snapshot backups that you configure when you create the backup.</p> <p>See “Configuring file and folder options for backup jobs” on page 552.</p> <p>If you configure Backup Exec to back up open files with a lock, it attempts to open the files in a different mode. It locks these files during backup to prevent other processes from writing to them. Symantec recommends that you close the applications that leave files open so that the files are backed up in a consistent state.</p> <p>To back up open files on Windows computers, use the Advanced Open File options to configure the backups that use snapshot technology.</p> <p>See “Configuring Advanced Open File options for backup jobs” on page 540.</p>
Why does the Backup Exec Administration Console continue to own a storage device even when it's not running?	<p>Backup Exec is a client/server application that must always be available to process the jobs that are submitted from both local and remote administrative consoles.</p> <p>The Backup Exec services claim all of the storage devices that are attached to the Backup Exec server whenever the services are running. Backup Exec requires constant control of the storage devices to collect statistics on media and storage device usage, and to provide media overwrite protection when necessary.</p>

Table 16-2 Backup questions (*continued*)

Question	Answer
When I run a local backup, the total number of bytes backed up by Backup Exec does not match the number of bytes displayed by Windows. Why?	<p>The type of partition for which the system is formatted may cause this problem.</p> <p>If you have a Windows NTFS compressed partition, Backup Exec displays the uncompressed byte count of the files that are backed up. Meanwhile, Windows Explorer displays the compressed byte count of the files on the hard drive. For example, Windows compresses an NTFS partition that contains 1 GB of data to 500 MB. Backup Exec reports that 1 GB of data was backed up, even though Windows Explorer displays that only 500 MB of compressed data exists on the hard drive.</p> <p>If you have a FAT partition, Backup Exec reports the actual number of bytes of the files being backed up while File Manager reports an inflated amount of disk space. For example, a 2 GB FAT partition has a 32-K cluster size and File Manager displays 1.9 GB of used space. Backup Exec reports that 1.4 GB of data was backed up. Assuming that a 50-MB pagefile.sys is excluded from the backup, there is a 450-MB difference in the number of bytes.</p> <p>Converting to NTFS regains disk space since it is more efficient and the default cluster size (automatically set by Windows) in NTFS is less than FAT. Windows lets you specify a cluster size other than the default; however system performance may decrease. For more information, see the Windows documentation.</p>

Troubleshooting failed components in the SAN

Various problems can occur at any location in a SAN.

For Backup Exec to work properly, a storage device has to be recognized in the following locations:

- The bridge or router must recognize it as a SCSI device
- The operating system must recognize it as a device
- Backup Exec must recognize it as a supported device

In some cases, hardware issues may require you to contact your hardware vendor for technical support.

You may need to replace a component of your SAN, such as a bridge or a switch. For specific steps for replacing your equipment, refer to your hardware vendor's documentation.

See ["Troubleshooting offline storage devices in a SAN "](#) on page 694.

Troubleshooting offline storage devices in a SAN

If a device in your SAN has gone offline, follow these steps to determine the source of the problem.

Before you begin troubleshooting, verify that your storage devices are on the Backup Exec supported device list.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Also, verify that all hardware drivers are up to date and are started. If you find errors with your hardware, contact your hardware vendor for specific instructions.

Table 16-3 Troubleshooting offline storage devices in a SAN

Step	Action
Step 1	<p>Use the Windows Device Manager to verify that the operating system recognizes the device.</p> <p>If the device is not recognized, you may need to troubleshoot the device.</p> <p>See "Finding hardware errors in a SAN" on page 696.</p>
Step 2	<p>For robotic libraries, verify that robotic library support is installed.</p> <p>See "About the Library Expansion Option " on page 358.</p>

Table 16-3 Troubleshooting offline storage devices in a SAN (*continued*)

Step	Action
Step 3	<p>Check the system event log for the following errors, which indicate SAN communication errors: SCSI errors 9, 11, and 15, or timeout errors relating to storage. Check the application event log for multiple events 33152. These events indicate SAN communication errors.</p> <p>See "Finding hardware errors in a SAN" on page 696.</p> <p>You may need to contact your hardware vendor.</p>
Step 4	<p>If the library is online, but some or all of the drives are offline, use Backup Exec to initialize the library.</p> <p>See "Initializing a robotic library" on page 456.</p>
Step 5	<p>If initializing the library does not bring the storage devices online, check the library for an error display on the front panel, mechanical problems, or tapes inappropriately in the drives. Correct any errors that you find.</p>
Step 6	<p>If no errors exist on the library or if you corrected the errors and the storage devices are still offline, stop Backup Exec services and then restart.</p> <p>See "Starting and stopping Backup Exec services" on page 621.</p>
Step 7	<p>If restarting the services does not bring the storage devices online, restart the operating system. Be sure that no Backup Exec jobs are running when you restart.</p>
Step 8	<p>If restarting the operating system does not bring the storage devices online, reset the SAN to help identify problem tape storage. Recycling the SAN may also resolve Fibre Channel problems.</p> <p>See "Resetting the SAN" on page 696.</p>

Finding hardware errors in a SAN

Use the following steps to find the common hardware errors that occur in a SAN. If you find errors with your hardware, contact your hardware vendor for specific instructions.

Table 16-4 Finding hardware errors in a SAN

Step	Action
Step 1	Verify that the proper device drivers are installed.
Step 2	Verify that the fibre cable is securely connected to the HBA and to the fibre switch.
Step 3	Verify that the fibre connection is properly connected to the robotic library from the fibre switch.
Step 4	Check for a failed hardware component between the server and the fibre switch. Sometimes some of the servers in the SAN recognize the storage device, but other servers do not. If none of the servers in the SAN recognize the storage device, check for a failed hardware component between the fibre switch and the storage device.
Step 5	Reset the SAN, which may identify problem hardware components and may resolve fibre problems. See “Resetting the SAN” on page 696.

Resetting the SAN

Resetting the SAN involves turning off the components of the SAN, and then powering them on in a specific order.

Table 16-5 Resetting the SAN

Step	Action
Step 1	<p>Turn off all servers, robotic libraries, and fibre bridges in the SAN.</p> <p>In rare cases, you may need to turn off the fibre switch also. If you need to turn off the switch, you should power it on before any other components. Wait for all checks to complete before you turn on the other components.</p>
Step 2	<p>Turn on the robotic library.</p> <p>See “Initializing a robotic library” on page 456.</p>
Step 3	<p>Verify that the fibre switch recognizes the robotic library.</p>
Step 4	<p>Turn on the central administration server.</p>
Step 5	<p>Verify that the operating system recognizes the robotic library and its drives.</p>
Step 6	<p>Turn on one of the managed Backup Exec servers. Wait for the managed Backup Exec server to start before you turn on the other managed Backup Exec servers.</p>

Bringing storage devices online after an unsafe device removal event in a SAN

If a storage device is in use by Backup Exec at the time of an unsafe device removal event, the device appears as offline in Backup Exec.

Table 16-6 How to bring a device online after an unsafe device removal event

Step	Action
Step 1	<p>Verify that no Backup Exec jobs are running in the SAN.</p>
Step 2	<p>Use Backup Exec to initialize the robotic library if the library is online but the drives are offline.</p> <p>See “Initializing a robotic library” on page 456.</p>

Table 16-6 How to bring a device online after an unsafe device removal event
(continued)

Step	Action
Step 3	<p>Stop all Backup Exec services and then restart them if the library is offline or if the drives are offline after initialization.</p> <p>If the device is not online, you may need to troubleshoot the device.</p> <p>See "Finding hardware errors in a SAN" on page 696.</p>

How to improve Backup Exec's performance

To get the best performance from Backup Exec, you should review several factors:

- Data transfer path
- Backup Exec agent performance
- Network performance
- Backup Exec server performance
- Storage device performance

For more information on how to measure and tune the performance of these items, go to the following URL:

<http://entsupport.symantec.com/umi/V-269-38>

Accessing Symantec Online

You can access Symantec community forums, learn about training courses, and view Symantec websites

Table 16-7 Symantec Online menu items

Item	Description
Share Your Ideas	Connects you to the SymConnect forum where you can post your ideas for improving Backup Exec. You can also access SymConnect at the following URL: http://www.symantec.com/connect/
Education Services	Provides the links to all Symantec Education training and custom learning services.
Backup Exec Tech Center	Provides the links to Backup Exec self-paced training modules.
Backup Exec Page	Provides the links to resources for Backup Exec.
Symantec Home Page	Connects you to the Symantec website.

To access Symantec Online

- ◆ Click the Backup Exec button, select **Symantec Online**, and then select the appropriate menu item.

See “[Searching the Symantec Knowledge Base](#)” on page 699.

Searching the Symantec Knowledge Base

The Symantec Knowledge Base is a centralized location where you can find more information about Symantec products. The knowledge base contains information about how to install, upgrade, configure, and use Symantec products. It also contains information about requirements, best practices, and how to troubleshoot problems. The Symantec Knowledge Base is accessible from within Backup Exec.

Note: You must have an active Internet connection to access the Symantec Knowledge Base.

The knowledge base uses a keyword-based search technology. It focuses on the important keywords in a search and compares them to other search phrases to provide the best possible results. You can use Boolean search features and expression queries to provide search parameters. For best results, focus on a few keywords that best represent your question.

When you search the knowledge base, a new browser window launches and displays the search results.

To search the Symantec Knowledge Base

- 1 Do either of the following:
 - Click the Backup Exec button, select **Technical Support**, and then select **Search Knowledge Base**.
 - On the **Home** tab, in the **Support** group, ensure that **Technical Support** is checked. Then, in the **Technical Support** panel, click **Symantec Technical Support**.
- 2 Enter a keyword or phrase, and then click the search icon.

Contacting Backup Exec Technical Support

If you have tried to solve a problem, but still need a resolution, you can contact Symantec Technical Support for Backup Exec over the Internet or by phone.

To expedite the Technical Support process, do the following:

- Know your Backup Exec version and revision number.
- Use one of the diagnostic utilities that are included with Backup Exec to collect the information that technical support can use to diagnose your issue.

To contact Backup Exec Technical Support

- ◆ Click the Backup Exec button, select **Technical Support**, and then select **Backup Exec Technical Support**.

See [“Displaying the version information for Backup Exec”](#) on page 136.

See [“About Backup Exec diagnostic tools”](#) on page 701.

See [“Accessing Symantec Online”](#) on page 698.

Using Symantec Remote Assistance

Symantec Remote Assistance launches a WebEx session on the Internet, which lets you join a support session or start a support session.

To use Symantec Remote Assistance

- ◆ Click the Backup Exec button, select **Technical Support**, and then select **Symantec Remote Assistance**.

Managing your Backup Exec support cases

From Backup Exec, you can launch the MySymantec website, where you can create, review, and manage cases related to technical product support.

To manage your Backup Exec support cases

- ◆ Click the Backup Exec button, select **Technical Support**, and then select **Manage Support Cases**.

About Backup Exec diagnostic tools

The following diagnostic tools help to troubleshoot issues in Backup Exec:

Table 16-8 Backup Exec diagnostic tools

Item	Description
Symantec Help Tool (SymHelp)	Scans the local computer and generates a report about common issues in your Backup Exec environment. See “Running the Symantec Help Tool (SymHelp)” on page 702.
Backup Exec diagnostic application	Gathers the pertinent information about a Windows computer for troubleshooting. See “Generating a diagnostic file for troubleshooting Backup Exec” on page 702.
Gather utility for Linux servers	Creates and compiles a Packet file. The file contains detailed information about installation, diagnostics, and error reports. See “Running the begather utility to troubleshoot Backup Exec components on Linux servers” on page 706.
Backup Exec Debug Monitor	Captures the debug output from Backup Exec and saves it into debug logs. See “Using the Backup Exec Debug Monitor for troubleshooting” on page 707.

Table 16-8 Backup Exec diagnostic tools (*continued*)

Item	Description
Backup Exec debug tool	<p>Generates the diagnostic information about the Backup Exec processes that shut down unexpectedly.</p> <p>See “About the Backup Exec debug tool” on page 707.</p>

Running the Symantec Help Tool (SymHelp)

The Symantec Help Tool (SymHelp) is a multi-product diagnostic utility designed for identifying common issues, gathering data for support assisted troubleshooting and providing links to other customer self-help and support resources. SymHelp performs checks by gathering information from the local computer and analyzing it. SymHelp does not permanently alter any files on the computer unless you explicitly agrees to this when prompted. SymHelp does not permanently install anything on your computer when it runs.

You can find more information about SymHelp at the following URL:

<http://www.symantec.com/business/support/index?page=content&id=TECH170735>

To run the Symantec Help Tool

- 1 Click the Backup Exec button, select **Technical Support**, and then select **Symantec Help Tool**.
- 2 Follow the on-screen prompts.

Generating a diagnostic file for troubleshooting Backup Exec

Backup Exec includes a diagnostic application (Bediag.exe) that gathers information about a Windows computer for troubleshooting purposes. This application can be run from within Backup Exec or it can be run from a command line. This application can be run for a local server or a remote server. You can run diagnostics on a remote Backup Exec server if Backup Exec is installed on the remote server and the Backup Exec services are running.

The Bediag application collects the following types of information:

- Account groups, account privileges, and environment settings.

- Backup Exec software version and registry information, Backup Exec agent listing, Windows version information, SCSI hardware configuration, SQL Server information, Driver services information, and Windows Services information.
- Server information, supported shared directories, and Windows sockets information.

To generate a diagnostic file from within Backup Exec

- 1 Click the Backup Exec button, select **Technical Support**, and then select **Backup Exec Diagnostics**.
- 2 Select the server from the drop-down list.
 If the server that you want to select is not in the drop-down list, click **Browse** to select it from a list of available servers.
- 3 Enter the logon information for the server that you selected.
- 4 Click **Run Diagnostics**.
- 5 Click **Close**.
- 6 Open the "Bediag.txt" from the directory that contains Bediag.exe (by default `Program Files\Symantec\Backup Exec`).

To generate a diagnostic file from a command line

- 1 Launch the command prompt.
- 2 Do one of the following:

To generate a diagnostic file for a Backup Exec server From the directory `Program Files\Symantec\Backup Exec\`, type *bediag [switches] servername* .

See ["Command line switches for a diagnostic file"](#) on page 704.

To generate a diagnostic file for a remote computer From the directory `Program Files\Symantec\Backup Exec\`, type *bediag [switches] workstationname*.

See ["Command line switches for a diagnostic file"](#) on page 704.

- 3 Open the "Bediag.txt" from the directory that contains Bediag.exe (by default `Program Files\Symantec\Backup Exec`).

Command line switches for a diagnostic file

You can add the following switches to gather additional information when you generate a diagnostic file for troubleshooting.

See [“Generating a diagnostic file for troubleshooting Backup Exec”](#) on page 702.

Table 16-9 Command line switches for a diagnostic file

Switch	Description
/a	Dumps the Agent List.
/ad	Dumps Active Directory information.
/adamm	Appends ADAMM file information.
/agents	Dumps agent information for the favorite resources.
/all	Dumps everything.
/app	Dumps the Application Event log.
/b2d	Dumps Backup Exec backup-to-disk information.
/basicscsi	Dumps basic SCSI hardware subkey from the registry.
/beallfiles	Dumps all Backup Exec files and directories.
/bedb	Dumps Backup Exec Database information.
/befiles	Dumps Backup Exec file information.
/bereg	Dumps Backup Exec software configuration from the registry.
/beupdate	Dumps Backup Exec update information.
/bex	Dumps only the Backup Exec entries that are in the Application Event log.
/c	Dumps the Backup Exec software configuration from the registry.
/caso	Dumps information about the central administration server and managed Backup Exec servers.
/cluster	Dumps cluster information
/cps	Dumps CPS registry information.
/detailnic	Dumps the server's detailed network card information.
/detailscsi	Dumps the server's detailed SCSI adapter information.

Table 16-9 Command line switches for a diagnostic file (*continued*)

Switch	Description
/dirsvc	Dumps the directory service event log.
/dlo	Dumps DLO ini file.
/dns	Dumps the DNS event log.
/err	Dumps only the error events from any event log.
/evdb	Dumps Symantec Enterprise Vault database information.
/evevents	Dumps the Symantec Enterprise Vault event logs.
/evreg	Dumps Symantec Enterprise Vault registry information.
/frs	Dumps the file replication service event log.
/o:[file]	Specifies the output job log for append. Omitting [file] sends output to the screen.
/h	Dumps the SCSI hardware subkey from registry.
/instapp	Dumps information about installed applications.
/lic	Dumps Backup Exec server license information.
/liveupdate	Dumps LiveUpdate information.
/lotus	Dumps the Lotus Notes information.
/n	Dump Windows Socket Network Protocols.
/networkinformation	Dumps TCP/IP settings, Winsock information, and Windows firewall information.
/oracle	Dumps Oracle information.
/p	Dumps the user privileges.
/power	Dumps the PowerShell event log.
/recs:n	Dumps only the newest records from given event logs. The bex, err, and recs switches must be used with the app switch and the sys switches.
/s	Dumps the information on Services.
/sec	Dumps the security event log.

Table 16-9 Command line switches for a diagnostic file (*continued*)

Switch	Description
/server	Dumps server information, such as CPU, memory, disk information, and more.
/services	Dumps information about services.
/sql	Dumps the Microsoft SQL Server information.
/symreginfo	Dumps the Symantec registry information.
/sys	Dumps the system event log.
/sys32info	Dumps driver information from the system32\drivers folder.
/svcacct	Dumps user names used for Backup Exec services.
/u	Dumps Microsoft update information.
/userinfo	Dumps user privileges and group information.
/v	Dumps Server Information.
/w	Dumps Windows version information.
/winpower	Dumps the Windows PowerShell event log.
/winupdate	Dumps Microsoft update information.
/x	Dumps Microsoft Exchange Server Information.
/?	Displays usage information.

Running the begather utility to troubleshoot Backup Exec components on Linux servers

The begather utility brings together the files that help you diagnose issues with Backup Exec components on Linux servers. After you run it, the begather utility displays the name of the Packet file that it creates. The files that are gathered contain detailed information regarding installation, diagnostics, and error reporting. Reviewing these files before contacting technical support can reveal the source of the issue. If the solution is not evident based on the gathered files, have the Packet file available when contacting support. The support technician may request an email that contains the Packet file.

Run the begather utility to troubleshoot Backup Exec components on Linux servers

- 1 Log on as root to the Linux server on which the Backup Exec components are installed.
- 2 Navigate to the following directory:
`/opt/VRTSralus/bin`
For example:

```
cd /opt/VRTSralus/bin
```
- 3 Start the begather utility.
For example:

```
./begather
```
- 4 Note the location of the Packet file that is displayed on the screen.

Using the Backup Exec Debug Monitor for troubleshooting

The Backup Exec Debug Monitor, or SGMon, is a diagnostic tool that captures debug output from Backup Exec and saves it in debug logs. SGMon debug logs can help you troubleshoot backup issues. Furthermore, debug logs can help Symantec Technical Support diagnose and repair problems.

When you open SGMon, it automatically captures debug data from Backup Exec's services. To collect debug information while SGMon is closed, enable debug log creation outside of SGMon and specify a directory in which to save the logs.

For more information about how to configure the Debug Monitor and read log files, refer to the help within the Debug Monitor.

To use the Backup Exec Debug Monitor for troubleshooting

- ◆ Click the Backup Exec button, select **Technical Support**, and then select **Collect debug output**.

About the Backup Exec debug tool

Backup Exec includes a debug tool (BEDBG) that generates diagnostic information about the Backup Exec processes that shut down unexpectedly. The diagnostic information helps Symantec Technical Support diagnose and repair the problem. The Backup Exec debug tool runs by default in Backup Exec. The data that the tool gathers is copied into the BEDBG folder which is located in \Program Files\Symantec\Backup Exec.

You can find more information about the Backup Exec debug tool at the following URL:

<http://seer.entsupport.symantec.com/docs/327408.htm>

Using Backup Exec in cluster environments

This chapter includes the following topics:

- [About Backup Exec and clusters](#)
- [Requirements for clustering Backup Exec in a Microsoft Cluster Server environment](#)
- [How Backup Exec works in a Microsoft Cluster Server environment](#)
- [Requirements for installing Backup Exec on a Microsoft Cluster Server](#)
- [Installing Backup Exec on a Microsoft Cluster Server](#)
- [Upgrading Backup Exec on a Microsoft cluster](#)
- [Installing additional Backup Exec options on a Microsoft cluster](#)
- [Uninstalling Backup Exec from a Microsoft cluster](#)
- [Creating storage device pools for Microsoft Cluster Servers](#)
- [Specifying a different failover node](#)
- [Designating a new central administration server in a Microsoft Cluster Server](#)
- [Configurations for Backup Exec and Microsoft Cluster Servers](#)
- [Using the Central Admin Server Option with Microsoft clusters and a storage area network](#)
- [About backing up Microsoft Cluster Servers](#)
- [About restoring data to a Microsoft cluster](#)

- [Disaster recovery of a cluster](#)
- [Restoring the Microsoft Cluster Server data files](#)
- [Recovering all shared disks in a Microsoft cluster](#)
- [Recovering Backup Exec in a Microsoft cluster](#)
- [Changing the Quorum disk signature for clusters on Windows Server 2003](#)
- [Manually joining two cluster disk groups and resynchronizing volumes](#)
- [Troubleshooting clusters](#)

About Backup Exec and clusters

In a server cluster, Backup Exec can protect data on local disks and shared disks, as well as protect Microsoft SQL and Exchange databases that are configured as virtual server applications; that is, they contain an IP address resource, a Network Name resource, and are displayed on the network with a unique server name (the virtual server name). Clustered servers provide high availability of applications and data to users. In a clustered server, several servers (called nodes) are linked in a network, and run cluster software that allows each node access to the shared disks. If a node becomes unavailable, cluster resources migrate to an available node (called failover). The shared disks and the virtual server are kept available. During failover, users experience only a short interruption in service.

See [“Installing Backup Exec on a Microsoft Cluster Server”](#) on page 713.

See [“Configurations for Backup Exec and Microsoft Cluster Servers”](#) on page 719.

See [“About backing up Microsoft Cluster Servers”](#) on page 727.

See [“About restoring data to a Microsoft cluster”](#) on page 727.

See [“Requirements for clustering Backup Exec in a Microsoft Cluster Server environment”](#) on page 710.

See [“Disaster recovery of a cluster”](#) on page 728.

Requirements for clustering Backup Exec in a Microsoft Cluster Server environment

The following are requirements for clustering Backup Exec:

- Symantec strongly recommends that if you plan to cluster Backup Exec, use the default database instance (MSDE) that Backup Exec installs.

- Symantec supports using a remote SQL Server instance to host the Backup Exec Database.
 Only one installed instance of Backup Exec can be installed into the remote SQL Server instance on a clustered node. All other installed instances of Backup Exec in the cluster must use the default Backup Exec MSDE database instance.

Note: You must run the Backup Exec cluster wizard on the cluster node that uses the remote SQL Server instance.

- For off-host backups that use the hardware provider in a Microsoft cluster (MSCS) environment, the Backup Exec server and the remote computer must be in different cluster groups. The cluster applications cannot support device logical unit numbers (LUNs) that have duplicate signatures and partition layouts. Therefore, the snapshots containing the LUNs must be transported to a host or remote computer that is outside the cluster.
- For Windows Server 2008 and later, you must use the same level of operating system on the Backup Exec server and on the computer that hosts the remote SQL Server instance in the following scenarios:
 - To use a remote clustered SQL Server instance to host the Backup Exec Database
 - To use the Backup Exec Utility to reconfigure the clustered Backup Exec installation or the remote clustered SQL Server instance

For Windows Server 2012, Backup Exec supports parallel backups across Cluster Shared Volumes (CSV) and cluster nodes.

Backup Exec does not support virtual machines that store data on both CSVFS and NTFS volumes.

See [“Installing Backup Exec on a Microsoft Cluster Server”](#) on page 713.

See [“About backing up Microsoft Cluster Servers”](#) on page 727.

How Backup Exec works in a Microsoft Cluster Server environment

When you install Backup Exec into a Microsoft cluster, you install it as a virtual server application. You assign an IP address resource, a Network Name resource (the virtual server name), and a disk resource to Backup Exec.

When a failover occurs, backup jobs that were running are rescheduled. The Backup Exec services are restarted on a designated failover node, and the backup jobs are

restarted by default. Backup Exec provides an additional rule for cluster failover restart called checkpoint restart. A checkpoint restart option allows backup jobs to continue from the point at which the jobs were interrupted rather than starting the backup again. If you enable the rule to retry jobs on a cluster failover, then you can specify an additional option to do a checkpoint restart when retrying the job. Checkpoint restart is the only property available for the Cluster Failover Rule. You can change the default so that jobs are not restarted.

When the failed server comes back online, the Microsoft cluster can automatically rebalance the workload in a cluster, called failback, by moving cluster groups back to the server that has rejoined the cluster. However, by design, Backup Exec does not failback. The backup jobs will continue to run on the designated failover node. By continuing to run backup jobs on the designated failover node, any further risk of having to restart the jobs again when the failed server rejoins the cluster is avoided. Then, when it is convenient, you can move the Backup Exec cluster group back to the controlling node.

Specific details of how Backup Exec runs in a cluster vary depending on the configuration you use in the cluster.

See [“Configuring checkpoint restart”](#) on page 543.

See [“Configurations for Backup Exec and Microsoft Cluster Servers”](#) on page 719.

See [“Installing Backup Exec on a Microsoft Cluster Server”](#) on page 713.

Requirements for installing Backup Exec on a Microsoft Cluster Server

The following items are required to install Backup Exec on a Microsoft cluster:

- Two-node clusters are supported with Backup Exec on Microsoft Windows Server 2003 Enterprise/DataCenter, and Windows Server 2008 R2 Enterprise/DataCenter.
- Four-node clusters are supported with Backup Exec on Microsoft Windows Server 2008 R2 Enterprise/DataCenter.
- Up to eight-node clusters are supported with Backup Exec on Microsoft Windows Server 2003 DataCenter.
- Backup Exec clusters can be installed on Windows Server 2003/2008 R2 majority node configurations. However, there must be a shared disk in the configuration in order for Backup Exec to share its database files between nodes. In this type of configuration, if the majority of the cluster nodes fail, then the entire cluster will fail. This configuration normally uses more than two nodes in the cluster configuration.

- The controlling node and designated failover nodes must be online during installation of Backup Exec into the cluster.
- A unique IP address and a unique network name for the Backup Exec virtual server is required during installation.
- During installation of a Backup Exec cluster, the node that runs the installation should own the shared disk. If you use a physical disk resource that belongs to another application, the Backup Exec Cluster Wizard will move all the resources that belong to the other application into the Backup Exec group. It is recommended that Backup Exec not be installed on the cluster quorum.
- An individually licensed copy of Backup Exec, as well as any applicable agents and options, is required for each active node in the cluster as defined in the End User License Agreement. When installing a trial version of Backup Exec, a cluster environment is automatically detected and licenses are not required.
- When you install Backup Exec clusters in a Central Admin Server Option (CASO) configuration, all Backup Exec installations must have the same server configuration. Either all nodes should be database servers or all nodes should be managed Backup Exec servers that connect to the central administration server.
- All Backup Exec installations into a cluster must either be part of a single cluster group, or be locally installed on each node. If cluster-aware Backup Exec is installed in a cluster as well as a locally installed version of Backup Exec (not cluster-aware), then you cannot log on to the locally installed Backup Exec server. You can only log on using the Backup Exec virtual server name. To be able to log on to the locally installed Backup Exec server, you must first use the Cluster Configuration Wizard to uninstall cluster-aware Backup Exec from all the nodes in the cluster.
- Use the same account for Backup Exec services on all nodes in the cluster. If nodes in a cluster use Backup Exec and have different accounts, change the services to use the same account.

See [“Installing Backup Exec on a Microsoft Cluster Server”](#) on page 713.

See [“Configurations for Backup Exec and Microsoft Cluster Servers”](#) on page 719.

Installing Backup Exec on a Microsoft Cluster Server

Symantec does not recommend installing Backup Exec on the same disk that the cluster quorum is installed on. If you have to specify a new drive letter for the quorum disk during a recovery process, Backup Exec will not recognize the new drive and will not run.

Note: By default, failover from the controlling node to a designated node occurs in alphabetical order according to the machine name of each node. To change the order in which failover will occur on the designated nodes, rename the machines.

The Backup Exec Agent for Windows is automatically installed on all the nodes in the cluster. If this installation of Backup Exec is to be used to back up remote servers outside the cluster, install the Agent for Windows on those remote servers as well.

To install Backup Exec on a Microsoft Cluster Server

- 1 Install Backup Exec on all the nodes that you want in the cluster. Use the same installation path for each node.
- 2 From the node that you want to be the active node, start Backup Exec.
- 3 Click the Backup Exec button, select **Configuration and Settings**, and then select **Cluster Configuration Wizard**.
- 4 Follow the instructions on the screen.

On the **Virtual Server Information** screen, Backup Exec automatically displays a default name called BKUPEXECVRS for the virtual server. Type a new default name if you do not want to use the default.

- 5 When the Cluster Configuration Wizard completes, create a storage device pool that contains all the locally attached storage devices on each node to be used when failover occurs. This ensures that jobs can be run on the storage devices that are attached to the failover nodes.

See [“Creating storage device pools for Microsoft Cluster Servers”](#) on page 716.

- 6 Repeat step 5 for all nodes.

See [“Configurations for Backup Exec and Microsoft Cluster Servers”](#) on page 719.

See [“Specifying a different failover node”](#) on page 717.

Upgrading Backup Exec on a Microsoft cluster

You can upgrade Backup Exec on the nodes in a cluster without taking the nodes out of the cluster.

Table 17-1 Upgrading Backup Exec on a Microsoft cluster

Step	Action
Step 1	Select a node to upgrade and make that node the active Backup Exec cluster node.

Table 17-1 Upgrading Backup Exec on a Microsoft cluster (*continued*)

Step	Action
Step 2	Run the Backup Exec installation program on the active node.
Step 3	Move the cluster group to the next node you want to upgrade, and then run the Backup Exec installation program on that node. All of the resources except for the disk should be offline when moved over to each node for upgrade.
Step 4	Repeat step 3 for each node in the cluster.

See [“Installing additional Backup Exec options to the local Backup Exec server”](#) on page 76.

Installing additional Backup Exec options on a Microsoft cluster

Install additional Backup Exec options on each node of the cluster. For details on installing each option, see the appropriate section in this guide, or in online Help.

To install additional Backup Exec options

- 1 On the controlling node, make sure the Backup Exec group is online before you start installing additional options.
- 2 Install the additional options.

See [“Installing additional Backup Exec options to the local Backup Exec server”](#) on page 76.
- 3 After the installation is complete on the controlling node, use the cluster administrator to move the Backup Exec group to the next appropriate node, and repeat step 2.

Be sure to install the same options with the same settings for each node in the cluster.
- 4 To install the Agent for Oracle on Windows or Linux Servers on other nodes, map a drive to the shared disks where Backup Exec is installed on the cluster, and run SETUP.

See [“Uninstalling Backup Exec from a Microsoft cluster”](#) on page 716.

Uninstalling Backup Exec from a Microsoft cluster

You use the Cluster Configuration Wizard to remove Backup Exec.

To uninstall Backup Exec from a cluster

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Cluster Configuration Wizard**.
- 2 Use the wizard to remove cluster-aware Backup Exec from all selected servers.
When unclustering the active node, you can either leave the Backup Exec data on the shared drive or delete it. If you delete the data, you can make the data available on the active node.
- 3 Uninstall Backup Exec from all the nodes.
- 4 After Backup Exec has been uninstalled, move any resource disks from the Backup Exec cluster group to another group, and then delete the Backup Exec cluster group.
- 5 On any node, click **Start**, point to **Settings**, and then click **Control Panel** to uninstall Backup Exec.
- 6 Double-click **Add/Remove Programs**, and then in the list of currently installed programs, select **Symantec Backup Exec** and click **Change/Remove**.
- 7 Repeat step 5 for all nodes.

See [“Creating storage device pools for Microsoft Cluster Servers”](#) on page 716.

Creating storage device pools for Microsoft Cluster Servers

When Backup Exec is installed on a cluster, it creates a set of default storage pools for each node in the cluster. If a node has storage devices, those storage devices are automatically assigned to an appropriate device pool, which is also the default destination device on that node when you create backup or restore jobs. However, to allow jobs to run on the storage devices attached to a failover node after a failover occurs, you must create a storage pool that includes the storage devices from all of the nodes. If the cluster is also configured with tape devices on a shared SCSI bus, then add the tape device name used by each node to the storage pool. You must also select this storage pool as the destination device for all jobs that you want to be restarted.

You can create either a single storage pool, or you can create storage pools for device or media types so that when jobs fail over they can be restarted on "like" devices and media.

To create a storage pool for a cluster

- 1 From the controlling node, open Backup Exec.
- 2 Create a new storage device pool and add storage devices.
See [“Creating storage device pools”](#) on page 416.
- 3 Exit Backup Exec. If there are tape devices on a shared SCSI bus, then add the tape device name used by each node.
- 4 Using the cluster administrator, move the Backup Exec resource group to the next appropriate node.
- 5 Open Backup Exec, add storage devices for this node to the previous storage pool and then exit Backup Exec. If there are tape devices on a shared SCSI bus, then add the tape device name used by each node.
- 6 Repeat step 4 and step 5 for each node in the cluster.

Specifying a different failover node

You can do the following:

- Change the order in which the nodes fail over.
- Add a failover node to the cluster.
- Remove a failover node from the cluster.

Note: By default, in a MSCS cluster, failover from the controlling node to a designated node occurs in alphabetical order according to the machine name of each node. To change the order in which failover will occur on the designated nodes, rename the machines in the order in which they should fail over.

Before you add a node to the Backup Exec cluster configuration, you must install Backup Exec on it. Cluster services for a node should be online before you add or remove it from the cluster.

If you are removing a node, do not run the cluster configuration wizard from the node you want to remove.

To add or remove a failover node

- 1 On the controlling node, click the Backup Exec button, select **Configuration and Settings**, and then select **Cluster Configuration Wizard**.
- 2 Follow the instructions on the screen to add or remove a node.
- 3 If you have added a failover node, also add any locally attached storage devices that are to be used when failover occurs to the cluster storage pool. This ensures that jobs can be run on the storage devices that are attached to the failover nodes.

If you remove some but not all nodes in a cluster, an uninstall of Backup Exec results in a password being requested for the virtual server and the services continuing to run. You must remove Backup Exec from all nodes on the cluster.

See [“Uninstalling Backup Exec from a Microsoft cluster”](#) on page 716.

See [“Specifying a different failover node”](#) on page 717.

See [“Configurations for Backup Exec and Microsoft Cluster Servers”](#) on page 719.

Designating a new central administration server in a Microsoft Cluster Server

To designate a new central administration server for a cluster environment, use BEUtility.exe. BEUtility enables you to do various types of configuration and maintenance operations on your Backup Exec servers.

Note: In a cluster environment, do not use **Change Service Account** in BEUtility.exe.

To change a Backup Exec cluster from a Database Server to a Member Server

- 1 Install the new server as a managed Backup Exec server with the Library Expansion Option installed.

Make sure connections to the Backup Exec cluster and other member servers are working properly.
- 2 Using the cluster administrator, shut down the Backup Exec cluster services.

Be sure to keep the Disk resource online.
- 3 Move the catalog files from the Backup Exec cluster installation path to the respective installation paths on the new database server.

- 4 Use BEUtility.exe to connect all Backup Exec servers to the new database server and to start all Backup Exec services.
- 5 Stop and restart the Backup Exec Services on the new database server.
- 6 Using the Cluster Administrator, move the Backup Exec resource group to the failover node and make sure services start on that node.
- 7 Use BEUtility.exe to stop and restart the Backup Exec Services on all the member servers of the SAN in order for them to connect to the new database server.

To change a Backup Exec cluster from a central administration server to a managed Backup Exec server

- 1 Install the new server as a managed Backup Exec server.
 Make sure connections to the Backup Exec cluster and other managed Backup Exec servers are working properly.
- 2 Using the cluster administrator, shut down the Backup Exec cluster services.
 Be sure to keep the Disk resource online.
- 3 Move the catalog files from the Backup Exec cluster installation path to the respective installation paths on the new central administration server.
- 4 Use BEUtility.exe to connect all Backup Exec servers to the new central administration server and to start all Backup Exec services.
- 5 Stop and restart the Backup Exec Services on the central administration server.
- 6 Using the Cluster Administrator, move the Backup Exec resource group to the failover node and make sure services start on that node.
- 7 Use BEUtility.exe to stop and restart the Backup Exec Services on all the managed Backup Exec servers in order for them to connect to the new central administration server.

See [“Multi-node clusters on a fibre channel SAN with the Central Admin Server Option”](#) on page 724.

Configurations for Backup Exec and Microsoft Cluster Servers

Backup Exec supports various cluster configurations of between two and eight nodes on a fibre channel SAN, with locally attached storage devices, or with storage devices on a shared SCSI bus. You can use any combination of these configurations.

Note: If you install the cluster on a private network, use the Cluster Administrator to enable public communication if necessary.

If you are using a cluster on a fibre channel SAN or with storage devices on a shared SCSI bus and failover occurs, depending on the capability of your various SAN components, media might be orphaned in the tape drive until the failed node becomes active again.

If end-of-job markers were not written to the media before the failover occurred, the media may be marked as unappendable by the Backup Exec engine when the next append backup job is run. The media remains unappendable until it is overwritten (or erased, or the retention period expires, etc.).

If the storage device is a robotic library, you can review the Robotic Library Inventory report to discover if the media was marked unappendable by the Backup Exec engine. If the Full column reports a 3, the Backup Exec engine has marked the media as unappendable.

To add or remove hot-swappable devices in a cluster, run the Hot-swap Device Wizard on all Backup Exec Cluster nodes. If a server is not updated to recognize a new device, any job that is targeted to that device may fail.

See [“Adding or replacing devices with the Hot-swappable Device Wizard”](#) on page 358.

The following examples illustrate various cluster configurations:

- See [“Two-node cluster with locally attached storage devices”](#) on page 720.
- See [“Two-node cluster with tape devices on a shared SCSI bus”](#) on page 721.
- See [“Configuring a shared SCSI bus for tape devices”](#) on page 722.
- See [“Multi-node clusters on a fibre channel SAN with the Central Admin Server Option”](#) on page 724.

Two-node cluster with locally attached storage devices


In this configuration, cluster-aware Backup Exec is installed on the controlling node, failover occurs to designated nodes in the cluster, and storage devices are locally attached to each node.

Depending on the type of storage device, each node's locally attached storage devices are automatically assigned to an appropriate storage device pool that Backup Exec creates. These storage device pools also act as default destination devices on that node when you create backup or restore jobs. You must create a storage pool that includes storage devices on the controlling node and on each failover node in order for jobs to run when failover occurs.

See [“Creating storage device pools for Microsoft Cluster Servers”](#) on page 716.

To restore data in this configuration, move media to the failover node's locally attached storage device and reinventory the device before starting a restore operation.

Figure 17-1 Two-node Cluster with Locally Attached Storage Devices



See [“Multi-node clusters on a fibre channel SAN with the Central Admin Server Option”](#) on page 724.

Two-node cluster with tape devices on a shared SCSI bus

In this configuration, cluster-aware Backup Exec is installed on the controlling node, failover occurs to designated nodes in the cluster, and tape devices are attached to a shared SCSI bus that is separate from any shared SCSI bus for disks.

Because each node creates a unique tape device name for the same device, if the drive is not serialized, this configuration requires you to create a storage pool that includes the tape device name used by each node in order for jobs to run when failover occurs.


See [“Creating storage device pools for Microsoft Cluster Servers”](#) on page 716.

When failover occurs, a SCSI bus reset is issued. Therefore, tape devices and shared drives should not be connected to the same SCSI bus; each should be connected to separate SCSI buses.

See [“Configuring a shared SCSI bus for tape devices”](#) on page 722.

Note: If you are using a serialized tape device in a shared SCSI cluster configuration, media that is orphaned in a device because of a failover will be ejected from the tape device. If you are using a tape device that is not serialized, you need to manually eject the media from the device or reboot the device.

Figure 17-2 Two-node cluster with tape devices on a shared SCSI bus



See [“Multi-node clusters on a fibre channel SAN with the Central Admin Server Option”](#) on page 724.

Configuring a shared SCSI bus for tape devices

To configure tape devices on a shared SCSI bus, you must have SCSI cables, SCSI terminators, a SCSI adapter in each cluster node to provide a shared external bus between the nodes, and at least one tape device on the shared bus.

The tape devices must be connected to a bus that uses the same method of transmission that the device does (single-ended or differential). Only one transmission method can be used on a single SCSI bus, however, if the devices use different transmission methods, you can install a signal converter between the devices. A signal converter converts single-ended SCSI signals to differential SCSI signals.


Note: You must use a signal converter to connect single-ended and differential devices in order to avoid hardware damage.

You must terminate the SCSI bus at both ends so that commands and data can be transmitted to and from all devices on the bus. Each SCSI bus must have two terminators and they must be at each end of the segment.

If a tape device is in the middle of the bus, remove any internal termination in that device.

If the tape device is at the end of the bus, and the tape device has internal termination, you can use the device's internal termination to terminate the bus.

Figure 17-3 Example of a shared bus with tape devices at the end of the bus



Following are methods you can use to terminate a bus:

- **SCSI adapters.** This method is not recommended because if the server is disconnected from the shared bus, or if there is a power supply failure, the bus may not be properly terminated and may be inoperable.
- **Pass-through (or feed-through) SCSI terminators.** These can be used with SCSI adapters and with some tape devices. If the tape device is at the end of the bus, you can attach a pass-through SCSI terminator to terminate the bus. The internal terminators in the tape device must be disabled. This is a recommended method.

Note: To ensure termination if a power supply failure occurs, turn off the on-board terminators on the SCSI controller (using the host adapter manufacturer's recommended method) and physically terminate the controller with a terminator.

- **Y cables.** These can be used with some tape devices. If the tape device is at the end of the bus, you can attach a terminator to one branch of a Y cable to terminate the bus. The internal terminators in the tape device must be disabled. This is a recommended method.
- **Trilink connectors.** These can be used with some tape devices. If the tape device is at the end of the bus, you can attach a terminator to one of the trilink

connectors to terminate the bus. The internal terminators in the tape device must be disabled. This is a recommended method.

Besides terminating the bus, Y-cables and trilink connectors also allow you to isolate the devices from the shared bus without affecting the bus termination. You can maintain or remove that device without affecting the other devices on the shared SCSI bus.

To configure a shared SCSI bus for tape devices

- 1 Install the SCSI controllers for the shared SCSI bus.

Make sure that the SCSI controllers for the shared SCSI bus are using different SCSI IDs. For example, on the controlling node, set the SCSI controller ID to 6 and on the failover node, set the SCSI controller ID to 7.

- 2 Prepare the SCSI controllers for the shared SCSI bus. Refer to your SCSI host adapter manufacturer's documentation for details.

Do not have power on to both nodes while configuring the computers, or if both nodes have power on, do not connect the shared SCSI buses to both nodes.

- 3 Connect the shared SCSI tape devices to the cable, connect the cable to both nodes, and then terminate the bus segment using one of the methods discussed in the previous section.

See [“Two-node cluster with tape devices on a shared SCSI bus”](#) on page 721.

Multi-node clusters on a fibre channel SAN with the Central Admin Server Option

In this configuration, one or more clusters are attached to a fibre channel storage area network (SAN), with cluster-aware Backup Exec and the Central Admin Server Option (CASO) installed on the controlling node in each cluster. Shared secondary storage devices are attached to the fibre channel, although a single storage device can be shared between one or more clusters. Failover occurs (in alphabetical order of the machine name) to other designated nodes in the cluster.

Note: When using multiple clusters in a CASO environment, it is strongly recommended that the cluster nodes be connected to the storage devices using a fibre switch. If you use a hub rather than a fibre switch, the hub will receive a reset command during a failover event that causes all other components attached to the hub to be disconnected. You can designate any server on the fibre channel SAN as the central administration server.

You should create a failover storage pool for the cluster.

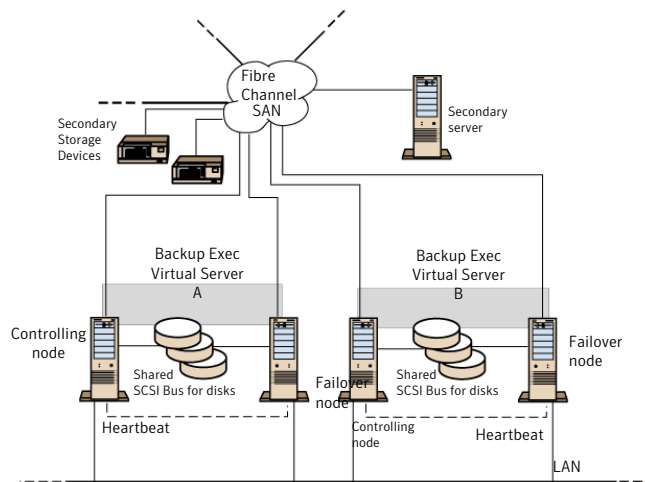
See [“Creating storage device pools for Microsoft Cluster Servers”](#) on page 716.

This configuration offers increased performance since backups are performed locally instead of over a network. Additionally, centralized media catalogs are available. Because CASO uses a shared catalog database, a tape that has already been cataloged can be physically moved from one device to another and not have to be recataloged.

Note: The CASO option must be installed on each failover node, with the same settings that were used on the primary node. Either all nodes should be database servers or all nodes should be managed Backup Exec servers.

The following are examples of multi-node clusters:

Figure 17-4 Two 2-node Clusters on a Fibre Channel SAN with the Central Admin Server Option



You can have a four-node cluster.

See [“Using the Central Admin Server Option with Microsoft clusters and a storage area network”](#) on page 726.

Using the Central Admin Server Option with Microsoft clusters and a storage area network

Managed Backup Exec servers can be clustered. However, this configuration is not recommended because the central administration server recovers all failed jobs in a distributed job environment.

The following configurations can be used when installing Backup Exec clusters with the Central Admin Server Option (CASO).

- Backup Exec cluster with CASO
- Backup Exec cluster with the managed Backup Exec server configuration

To install Backup Exec in a Microsoft cluster with CASO

- 1 Install Backup Exec with CASO and any additional options onto your Microsoft cluster nodes.
- 2 From the node that you want to designate as the active node, start Backup Exec.
- 3 Click the Backup Exec button, select **Configuration and Settings**, and then select **Cluster Configuration Wizard**.
- 4 Follow the instructions on the screen.
- 5 When the Cluster Configuration Wizard completes, install the managed Backup Exec server. Use the virtual Backup Exec cluster name when prompted for the central administration server.

To install Backup Exec in a Microsoft cluster with the managed Backup Exec server configuration

- 1 Install Backup Exec with the managed Backup Exec server option and any additional options onto your Microsoft cluster nodes.

All nodes that run Backup Exec in the Microsoft cluster configuration must access the same central administration server. If the nodes do not access the same central administration server, failovers do not occur properly.
- 2 From the node that you want to designate as the active node, start Backup Exec.
- 3 Click the Backup Exec button, select **Configuration and Settings**, and then select **Cluster Configuration Wizard**.
- 4 Follow the instructions on the screen.

See [“About backing up Microsoft Cluster Servers”](#) on page 727.

About backing up Microsoft Cluster Servers

To protect all data in the cluster, which includes file shares, databases, and the cluster quorum, back up the following:

- Local disks, Shadow Copy Components, and System State on each node. The cluster quorum, which contains recovery information for the cluster and information about changes to the cluster configuration, is included in the System State backup.
- All shared disks, including the data in the Microsoft Cluster Server folder on the Quorum disk.
- Virtual servers, which may contain data or contain applications such as Microsoft SQL Server or Exchange Server. Use Backup Exec database agents to back up databases.

Note: For offhost backup jobs that use the hardware provider, the Backup Exec server and the remote computer must be in different cluster groups. The cluster applications cannot support device logical unit numbers (LUNs) that have duplicate signatures and partition layouts, therefore, the snapshots containing the LUNs must be transported to a host, or remote computer, that is outside the cluster.

The Backup Exec Management Command Line Interface (BEMCLI) can be used with Backup Exec when Backup Exec is installed in a cluster. The only limitation is that you cannot use the BEMCLI to specify a device for backup. You can use the BEMCLI to target a storage pool, but not a specific device in that pool.

See [“Backing up data”](#) on page 158.

See [“Editing backup definitions”](#) on page 206.

See [“Adding a stage to a backup definition”](#) on page 219.

About restoring data to a Microsoft cluster

For all file restore operations, including redirecting restores, use the normal restore procedures.

When restoring files to shared drives, direct those files to the virtual server or the controlling node of the resource. When restoring individual database files, such as Microsoft SQL Server or Exchange Server, direct those files to the virtual server name of a specific installation of the SQL or Exchange database.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

Disaster recovery of a cluster

Prepare for recovery by creating a disaster preparation plan.

Prepare to restore SQL, Exchange, and Lotus Domino databases in a cluster after a disaster by reading the sections on preparing for disaster recovery in the appropriate chapters.

In addition to the initial preparation instructions, further action is required to completely protect the Microsoft clusters.

If a disaster occurs, the following information is required to successfully recover the cluster:

- General Cluster Information
 - Cluster name
 - Cluster IP address and subnet mask
 - Cluster node names
 - Node IP addresses
 - Local and shared drive letters and partition scheme
 - Disk signatures
- Cluster Groups
 - Group name
 - Preferred nodes
 - Failover/failback policies
- Cluster Resources
 - Resource name
 - Resource type
 - Group membership
 - Possible owners
 - Resource dependencies
 - Restart and Looks Alive/Is Alive properties
 - Resource-related parameters
 - Application-specific configuration (SQL Database Character Set)

- If you are recovering a Microsoft Cluster Server, run the Clusterrecovery.exe from the Microsoft 2003 Resource Kit to retrieve the disk signatures from the shared disk.

See [“Creating a disaster preparation plan \(DPP\)”](#) on page 144.

See [“How to prepare for disaster recovery on a Lotus Domino server”](#) on page 1079.

See [“Disaster recovery of an Exchange Server”](#) on page 968.

Using Simplified Disaster Recovery to prepare for disaster recovery of a cluster

Backup Exec provides a fully-automated disaster recovery solution called Simplified Disaster Recovery, which allows you to quickly and efficiently recover the nodes that comprise the server cluster after a disaster. Oracle servers cannot be restored using Simplified Disaster Recovery. For more information about disaster recovery for these options, see the appropriate chapters.

See [“Preparing computers for use with Simplified Disaster Recovery”](#) on page 739.

Note: To change the setup, use hardware, or a hardware configuration that is different from the original configuration, you must perform a manual recovery.

Recovering nodes on the cluster using Simplified Disaster Recovery

If you used Backup Exec's Simplified Disaster Recovery to prepare for a disaster, you can use Simplified Disaster Recovery to recover the nodes to their pre-disaster state.

Note: You must create disaster recovery media for each Windows 2003 cluster node. Disaster recovery media is customized for a single computer. You will not be able to use the media interchangeably between the nodes in a cluster.

When recovering both nodes in a cluster, make sure that the drive letters match the original cluster configuration. The scaled-back version of Windows that runs the recovery wizard may detect the hard drives in a different order than what was originally configured under the original version of Windows.

If the original configuration does not match, then to a certain extent, you can control the hard drive numbering scheme that Windows devises.

If you cannot get the SDR **Recover This Computer Wizard** to properly detect the hard drive order, you can still manually set up hard drive partitions using the **Advanced Disk Configuration** option within the **Recover This Computer Wizard**.

After this is done, you can continue with automated restore of your backup media afterward.

Note: After Windows has been installed, you cannot change the system drive's letter. You must restore the system to the same drive letter from which it was backed up.

To recover nodes on the cluster using Simplified Disaster Recovery

- 1 If you are recovering more than one node, disconnect the shared disks. If you are recovering only one node, the shared disks do not need to be disconnected.

If all nodes in the cluster are unavailable and must be recovered, the cluster cannot fail over. Disconnect the shared disks before recovery begins.
- 2 Restore the nodes.

See [“Methods for restoring data in Backup Exec”](#) on page 232.
- 3 Reconnect the shared drives and bring the nodes online.
- 4 To restore a database to the shared drives, use the appropriate Backup Exec agent.

See [“About Simplified Disaster Recovery”](#) on page 736.

Recovering Backup Exec on a Microsoft cluster using Simplified Disaster Recovery

To fully restore a cluster on which Backup Exec is installed, you can use Simplified Disaster Recovery to restore the cluster node and all shared disks or you can rebuild the cluster. To remotely restore the cluster catalog the media that contains the backup sets of the cluster nodes and the shared disk.

To recover Backup Exec on a Microsoft Cluster with Simplified Disaster Recovery:

- 1 Replace all shared disks, if necessary.
- 2 Run the **SDR Recover This Computer Wizard** on one of the nodes. During this process, use **Advanced Disk Configuration** to repartition all shared disks to their original configuration. Restore the local disk, system state, and the data files to the shared disk.
- 3 Restart the server.

The cluster service and all other cluster applications should come online.
- 4 Run the **Recover This Computer** wizard on all other nodes. Restore only the local disk and system state.

See [“About Simplified Disaster Recovery”](#) on page 736.

Recovering the entire cluster using a manual disaster recovery procedure

As part of the manual recovery process, you must reinstall Windows, including the last service pack applied before the failure.

See [“How to prepare for disaster recovery on a Lotus Domino server”](#) on page 1079.

See [“About restoring Oracle resources”](#) on page 1012.

To recover the entire cluster manually

- 1 On the first node you want to recover, reinstall Windows, including the last service pack applied before the failure.
- 2 On the other nodes you want to recover, reinstall Windows, including the last service pack applied before the failure.
- 3 Reinstall the cluster services and bring the cluster online.

Do the following:

- If you are recovering a Microsoft Cluster Server, after booting the nodes in a cluster, make sure that the drive letters match the original cluster configuration. If the original configuration does not match, then to a certain extent, you can control the hard drive numbering scheme that Windows devises by using the Disk Administrator.
- If you are recovering a Microsoft Cluster Server, use the Cluster Wizard to reinstall Backup Exec on the cluster. You must use the same settings used during the initial installation.

See [“Installing Backup Exec on a Microsoft Cluster Server”](#) on page 713.

- 4 Catalog the media in the cluster.
- 5 On the **Backup and Restore** tab on the active node, click **Restore**.
- 6 Restore last full backup sets made of the active node, and then restore the System State.
- 7 If you are recovering a Microsoft Cluster Server, select the **Restore cluster quorum** option in the **Restore Wizard**.
- 8 Start the restore operation.
- 9 When the restore has completed, reboot the active node.
- 10 For each node that you need to recover, repeat step 5 through step 9.

- 11 After all nodes are recovered, restore the Backup Exec data files, and all other data files, to the shared disks.
- 12 To restore a database to the shared disks, use the appropriate Backup Exec agent.

Restoring the Microsoft Cluster Server data files

To fully recover the cluster, the cluster files in the MSCS folder may need to be restored. If the Quorum disk is still available and has not changed, then you do not have to restore the data files.

For versions of Windows prior to Windows Server 2008, if the Quorum disk is new, you must restore the data files to the new Quorum disk. You should disable the cluster disk driver before restoring the data files.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

Recovering all shared disks in a Microsoft cluster

Recover shared disks using Cluster recovery from the Microsoft 2003 Resource Kit, which helps automate the recovery process or by performing a manual recovery.

To recover all shared disks using Dumpcfg

- 1 Disable the cluster disk driver on all nodes in order to gain access to the new disk.
- 2 On the **Computer Management** menu, select **System Tools**, then select **Device Manager**.
- 3 Right-click the cluster disk driver, and then select **Disable**.
- 4 Replace and then repartition the shared disk. Use Disk Manager to verify that all nodes have access to the same shared disk.
- 5 Run Dumpcfg or Clusterrecovery to replace the disk signature for the Quorum disk.
- 6 Using a remote Backup Exec server, restore the cluster files to the Quorum disk via the node that has access to the disk.
- 7 Enable the cluster disk driver on all nodes.
- 8 On the **Computer Management** menu, select **System Tools**. Then select **Device Manager**.
- 9 Right-click the cluster disk driver, and then select **Enable**.
- 10 Reboot all cluster nodes.

To recover all shared disks without using Dumpcfg

- 1 Uninstall all cluster applications and the cluster software from both nodes.
- 2 Replace and then use Disk Manager to repartition the shared disk to the previously saved configuration.
- 3 Reinstall the cluster software.
- 4 Reinstall the cluster-aware version of Backup Exec on the cluster.
- 5 Reinstall additional cluster-aware software applications on the shared disk.
- 6 Use Backup Exec to restore any data from the catalogs.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

Recovering Backup Exec in a Microsoft cluster

If you used Simplified Disaster Recovery's **Create Simplified Disaster Recovery Disk Wizard** to prepare disaster recovery media for the shared disks, you must use a manual process to recover Backup Exec on a shared disk.

To use a manual process to recover Backup Exec on a shared disk

- 1 Replace the shared disk if necessary, and add that disk to the cluster as a disk resource.
- 2 Reinstall the cluster-aware version of Backup Exec on the cluster using the same information used in the original installation.
- 3 Use Backup Exec to restore any data from the catalogs.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

Changing the Quorum disk signature for clusters on Windows Server 2003

For clusters on Windows Server 2003, the cluster service may not start because the disk signature on the Quorum disk is different from the original signature. You can change the disk signature.

To change the Quorum disk signature for clusters on Windows Server 2003

- 1 Start the cluster service on one node with the -Fixquorum option in the startup parameters.
- 2 Open the Cluster Administrator and right-click the cluster, and then select **Properties**.
- 3 Select the **Quorum** tab.

- 4 In the **Quorum resource** field, select a different disk.
- 5 Click **OK**.
- 6 Stop the cluster services and then restart them without the -Fixquorum option.
You may run the -Fixquorum option as many times as needed to redesignate a Quorum disk signature.
- 7 Bring all other nodes online.

See [“Troubleshooting clusters”](#) on page 734.

Manually joining two cluster disk groups and resynchronizing volumes

If an Advanced Disk-based backup failed due to the application virtual server failover, you may need to re-join the cluster disk groups.

To manually re-join the two cluster disk groups and resynchronize the volumes

- 1 Import the cluster disk group into the node, if the original cluster disk group is not already imported into the node where the production virtual server is currently online.
- 2 Rejoin the new cluster disk group with the original cluster disk group.
- 3 Snap back the snapped volumes with their original volumes. Ensure that the option to synchronize using the original volume is selected.

If you are not able to import the new cluster disk group into the node where the original cluster disk group is currently located, failover the application virtual server back to its original node before rejoining the two cluster disk groups.

See [“Troubleshooting clusters”](#) on page 734.

Troubleshooting clusters

If you experience problems with using Backup Exec in a cluster environment, review the questions and answers in this section.

Table 17-2 Cluster troubleshooting questions and answers

Question	Answer
After I recovered my cluster and all shared disks, the cluster service will not start. Why won't it start and how can I get it started?	The cluster service may not start because the disk signature on the Quorum disk is different from the original signature. See "Changing the Quorum disk signature for clusters on Windows Server 2003" on page 733.
I used the checkpoint restart option for my backups. During one of my backups, a Microsoft cluster failover occurred. Multiple backup sets were created. When I try to verify or restore using these backup sets, an "Unexpected End of Data" error occurs on the set that contains the data that was backed up prior to the failover. Why does this occur? Is my data safe?	You received this error because failover occurred in the middle of backing up the resource, therefore the backup set was not closed on the media. However, the objects that were partially backed up in the first backup set were completely backed up again during restart, ensuring data integrity. Therefore, all of the objects on the media for the given backup set should still be restored and verified.
I clustered a central administration server with a managed Backup Exec server. Now the device and media service on the managed Backup Exec server fails. Why?	This occurs when the managed Backup Exec server becomes the active node and attempts to connect to the Backup Exec database on the central administration server, which is no longer available. To correct this, you must use the Backup Exec Utility (BEUTILITY.EXE) or reinstall the managed Backup Exec server to become a central administration server.

See ["About Backup Exec and clusters"](#) on page 710.

Simplified Disaster Recovery

This chapter includes the following topics:

- [About Simplified Disaster Recovery](#)
- [Requirements for using Simplified Disaster Recovery](#)
- [Preparing computers for use with Simplified Disaster Recovery](#)
- [How to ensure that backups are enabled for Simplified Disaster Recovery](#)
- [How Simplified Disaster Recovery uses disaster recovery information files](#)
- [Setting or changing the alternate location for the disaster recovery information file](#)
- [Creating a Simplified Disaster Recovery disk image](#)
- [Preparing to recover from a disaster by using Simplified Disaster Recovery](#)
- [Recovering a computer with Simplified Disaster Recovery](#)
- [Advanced Disk Configuration on the Recover This Computer Wizard](#)

About Simplified Disaster Recovery

Simplified Disaster Recovery (SDR) is automatically installed with Backup Exec so that you can perform disaster recovery on Windows computers on which the Agent for Windows is installed. By default, Backup Exec is configured to back up all of the critical system components that you need to perform a disaster recovery.

After a computer's critical system components are backed up, use the **Create Simplified Disaster Recovery Disk Wizard** to create a Simplified Disaster Recovery

disk image. Use the recovery disk to perform disaster recovery of the computers that are backed up.

When you use SDR to perform a recovery, Backup Exec uses the system-level information from the SDR backup to rebuild the server and restore it to a functional state. Recovery includes a bare metal or dissimilar hardware restore operation.

Through integration with the Microsoft Volume Shadow Copy Service (VSS), the SDR backups include all selected elements, even if they are components of the active operating system or are in an open state. VSS integration ensures that the backups are in a consistent state and have been properly placed into a quiescent state when the backup occurs.

Simplified Disaster Recovery is available only for servers on which the Agent for Windows is installed and that are backed up through the Agent for Windows. You must purchase the Agent for Windows separately, and then install it on the remote computers that you want to protect.

The Simplified Disaster Recovery feature now requires the Microsoft Assessment and Deployment Kit (ADK) to create a Simplified Disaster Recovery disk image (.iso).

See [“Preparing computers for use with Simplified Disaster Recovery”](#) on page 739.

See [“How to ensure that backups are enabled for Simplified Disaster Recovery”](#) on page 743.

See [“Requirements for using Simplified Disaster Recovery”](#) on page 737.

See [“Creating a Simplified Disaster Recovery disk image”](#) on page 750.

See [“Recovering a computer with Simplified Disaster Recovery”](#) on page 764.

Requirements for using Simplified Disaster Recovery

The following items are required for Simplified Disaster Recovery (SDR):

- Backup Exec or the Backup Exec Agent for Windows must be installed on any computers that you want to protect with SDR.
- Encryption key files must exist for all volumes that you encrypt with Windows BitLocker Drive Encryption.
- A third-party ISO 9660-compliant CD or DVD burning application must be available to burn the SDR-created bootable image to a CD or DVD.
- A writable or rewritable CD or DVD device must be available.
- An Internet connection so that you can download the Microsoft Assessment and Deployment Kit.

- The option **Use storage-based catalogs** must be enabled. If you disable this option, the backup sets that you create for use with SDR cannot be restored during an SDR recovery operation. As a result, SDR cannot recover the failed computer. To ensure that this option is enabled, click the Backup Exec button, click **Configuration and Settings > Backup Settings > Catalog**.

Note: If you use deduplication disk storage devices, be aware that there are limitations in their use with SDR.

See [“Recovery notes for using Simplified Disaster Recovery with Exchange, SQL, SharePoint, CASO, Hyper-V hosts, and the Deduplication Option”](#) on page 771.

Additional requirements exist when you create a Simplified Disaster Recovery disk image and when you run the **Recover This Computer Wizard**, as described in the following sections.

Requirements for creating a Simplified Disaster Recovery disk image

The following items are required to create the Simplified Disaster Recovery disk image:

- The Simplified Disaster Recovery disk image must be the same version of Backup Exec as the Backup Exec server. You cannot use SDR to restore the backups that were created with previous versions of Backup Exec.
- The Microsoft Assessment and Deployment Kit (ADK). ADK 8.0 is supported only on Windows Server 2008 and later. If the Backup Exec server does not run on Windows Server 2008 or later, or if the server does not have an Internet connection, methods are provided to let you create the Simplified Recovery Disk.
- 5 GB of disk space to download and install ADK.
- 1 GB of disk space to store the required files and folders to create the ISO image.

Requirements for running the Recover This Computer Wizard

The following items are required to run the **Recover This Computer Wizard**:

- The Simplified Disaster Recovery disk image must be the same version of Backup Exec as the Backup Exec server.
- If the computer that you want to recover was backed up to a tape device, deduplication storage, or to a virtual disk, then SDR cannot store the disaster recovery information file with the backup sets. Instead, you must provide the path to the default location or to the alternate location when the **Recover This Computer Wizard** prompts you. The default location is C:\Program

Files\Symantec\Backup Exec\sdr\Data. If the file is unavailable, you cannot recover the computer with SDR.

- The backup set that contains all of the critical system components for the computer that you want to restore.
- The boot drive on the computer that you want to recover must have from 3-GB to 5-GB of free space depending on the operating system and configuration. If a blank screen appears and the computer does not restart after you use the SDR disk, ensure that the boot drive has the necessary amount of free space. Then, restart the computer again.
- The target computer that you want to recover to must have an amount of RAM that is equal to or greater than the minimum that the restored operating system requires.
- The target computer that you want to recover to must have disks with enough free space to contain the data that you restore from the backup of critical volumes. A volume is considered critical if it is required for the computer to start successfully.
- If you recover a Windows computer that has BitLocker encryption enabled, you must turn on BitLocker encryption after the restore. See your Microsoft documentation for more information on BitLocker drive encryption.
- If the computer you recover contains a RAID setup, you may be required to configure the RAID before you start it with the SDR disk. Use the computer manufacturer's RAID software to configure the RAID system.
- If you restore Windows 8 or Windows Server 2012 storage spaces and storage pools, you should be aware of the possible restore scenarios.
See ["Recovery notes for using Simplified Disaster Recovery with storage pools and storage spaces"](#) on page 769.

See ["Creating a Simplified Disaster Recovery disk image"](#) on page 750.

See ["Recovering a computer with Simplified Disaster Recovery"](#) on page 764.

Preparing computers for use with Simplified Disaster Recovery

Prepare computers for use with Simplified Disaster Recovery (SDR) by performing the following steps:

Table 18-1 Preparing computers for use with Simplified Disaster Recovery

Step	Description	More information
Step 1	Specify an alternate location where copies of the disaster recovery information files are stored.	<p>Backup Exec creates the disaster recovery information file after a backup job that includes all critical system components completes successfully. Backup Exec then stores the disaster recovery information file in the default storage location with the backup set on a disk storage or disk cartridge device, and in the alternate storage locations. Catalog entries from subsequent backups are automatically added to the disaster recovery information file.</p> <p>Warning: If you back up a computer to a tape device, deduplication storage, or a virtual disk, Backup Exec cannot store a disaster recovery information file with the backup set. You must have a disaster recovery information file in an alternate location to ensure that you can use SDR to recover the computer.</p> <p>See “Setting or changing the alternate location for the disaster recovery information file” on page 747.</p>

Table 18-1 Preparing computers for use with Simplified Disaster Recovery
(continued)

Step	Description	More information
Step 2	<p>Run the backup jobs that include all critical system components (SDR-enabled backups) for the computers that you want to protect. These are the backups for which the Simplified Disaster Recovery indicator is ON.</p> <p>Note: For environments running the Central Admin Server Option, run a database maintenance job before you run SDR-enabled backups. Otherwise, the central administration server denies communication attempts from the managed Backup Exec servers.</p> <p>See “Configuring database maintenance” on page 567.</p>	<p>By default, Backup Exec selects all critical system components when you select a computer for backup. When all critical system components are included in the backup job selections, the Simplified Disaster Recovery indicator on the backup selections appears as ON.</p> <p>If you deselect one or more critical system component files, the indicator appears as OFF. Symantec recommends that you select the entire computer for backup; otherwise, Backup Exec cannot create the system-specific disaster recovery information file.</p> <p>Critical system components include the following:</p> <ul style="list-style-type: none"> ■ System volume (including EFI and utility partitions) ■ Boot volume (executing operating system) ■ Services application volumes (boot, system, and automatic startup) ■ System State devices and volumes (including Active Directory, system files, and so on) <p>See “How to ensure that backups are enabled for Simplified Disaster Recovery” on page 743.</p>

Table 18-1 Preparing computers for use with Simplified Disaster Recovery
(continued)

Step	Description	More information
Step 3	Create additional copies of the disaster recovery information files and store them in a safe place.	<p>Backup Exec stores the important disaster recovery information files in the default path, the alternate location, and with the backup set if disk storage or disk cartridge devices are used as the destination storage. Symantec recommends that you also make additional copies of the files and store them in a safe place. Without the disaster recovery information files, you cannot recover Backup Exec servers by using Simplified Disaster Recovery (SDR). Having multiple copies of the disaster recovery information files ensures that you can successfully recover Backup Exec servers with SDR.</p> <p>By default, the disaster recovery information files reside in the \Program Files\Symantec\Backup Exec\SDR\Data directory on the Backup Exec server. Use Windows Explorer or another Copy utility to copy the disaster recovery information files from the default location to another storage location of your choice. Backup Exec does not automatically update these copies, but they do let you restore a computer to an older point-in-time if the other disaster recovery files are not available</p>
Step 4	Use the Create Simplified Disaster Recovery Disk Wizard to create a disk image, and then a bootable CD or DVD recovery disk.	<p>Backup Exec generates alerts after each backup job until you create the Simplified Disaster Recovery disk image. You can disable these alerts, but Symantec recommends that you create the Simplified Disaster Recovery disk image.</p> <p>See “Creating a Simplified Disaster Recovery disk image” on page 750.</p>

How to ensure that backups are enabled for Simplified Disaster Recovery

By default, Backup Exec is configured to back up all of the critical system components that you need to use Simplified Disaster Recovery to recover a computer.

When the Simplified Disaster Recovery indicator for the backup selections is green, or **ON**, it indicates that the critical system components are selected for backup. The backup is SDR-enabled. If the indicator is gray, or **OFF**, then the backup is not enabled for SDR. You can click the icon to select all of the necessary components for disaster recovery, or to disable disaster recovery for backups created by this job.

To ensure that a backup is enabled for Simplified Disaster Recovery, you can view the backup selections that appear when you create or edit a backup job.

To ensure that backups are enabled for Simplified Disaster Recovery

- 1 On the **Backup and Restore** tab, create a new backup job or edit an existing one.
- 2 On the **Backup Definition Properties** dialog box, on the selections pane, do one of the following:
 - Ensure that the icon to the left of the computer name is green.
 - Click **Edit**, and ensure that the icon on the right is green, and that the text indicates that Simplified Disaster Recovery is on.

Note: If you hover the mouse over the icon, the text indicates whether the Simplified Disaster Recovery is on or off.

- 3 If the icon is not green, or if the text indicates that Simplified Disaster Recovery is off, click the icon and then select the option **Select all necessary components for disaster recovery**.

Figure 18-1 Simplified Disaster Recovery indicator is ON in the **Backup Definition Properties** dialog box

Simplified Disaster Recovery indicator


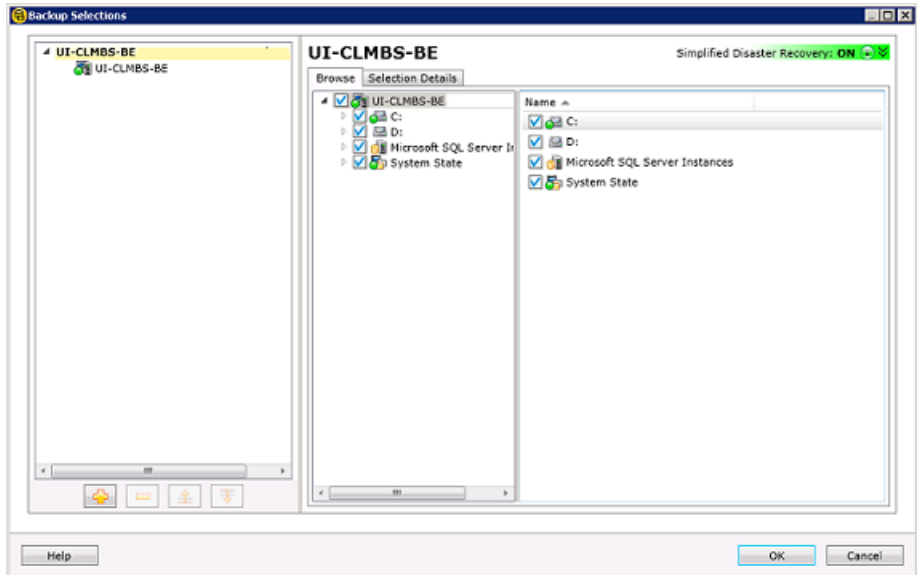


Figure 18-2 Simplified Disaster Recovery indicator is ON in the **Backup Selections** dialog box



See [“Preparing computers for use with Simplified Disaster Recovery”](#) on page 739.

How Simplified Disaster Recovery uses disaster recovery information files

For each computer that you back up and for which the Simplified Disaster Recovery indicator is **ON**, Backup Exec creates a disaster recovery information file. A disaster recovery information file contains computer-specific information for the computer that is backed up. Each time that a backup of all critical system components runs, the disaster recovery information files are automatically updated. Each disaster information recovery file uses the file name <computer_name>.DR. SDR uses the computer-specific information that is contained in the file when you run the **Recover This Computer** wizard. Without a disaster recovery information file, a recovery of the computer is not possible with SDR.

A disaster recovery information file contains the following information for the computer that is backed up:

- Hardware-specific information, such as hard disk partition information, mass storage controller information, and network interface card information.

- A list of catalog entries that identify the backup sets and storage media that are needed to recover the computer.
- The Windows Automated System Recovery configuration information file (asr.xml) for Windows Vista/Windows Server 2008/Windows Server 2008 R2/Windows 7. The ASR file is necessary to recreate partitions on Windows Vista/Windows Server 2008/Windows Server 2008 R2/Windows 7 computers during the recovery process.
- The Windows Automated System Recovery (ASR) configuration information files (asr.sif and asrnpn.sif) for Windows XP/Windows Server 2003 computers. The ASR files are necessary to recreate partitions on Windows XP/Windows Server 2003 computers during the recovery process.

Backup Exec stores the disaster recovery information file in the following locations:

- With the backup sets if the backup storage is disk storage or a disk cartridge device.
- On the Backup Exec server's hard drive in the following path:

`C:\Program Files\Symantec\Backup Exec\sdr\Data\`
- In an alternate location that you specify, on a different computer than the Backup Exec server.

Note: Symantec recommends that you specify an alternate storage location. If the Backup Exec server crashes, you cannot retrieve the disaster recovery information file from the default location. However, you can retrieve it from the alternate location. You should also create additional copies of the disaster recovery information files and store them in a safe place. Use Windows Explorer or another Copy utility to copy the disaster recovery information files from the default location to another storage location of your choice. Backup Exec does not automatically update these copies, but they do let you restore a computer to an older point-in-time if the other disaster recovery files are not available.

If the disaster recovery information file is stored with the backup sets, then SDR automatically uses that file to perform the recovery. If the backup sets are stored on a tape storage device, deduplication storage, or on a virtual disk, then SDR cannot store the file with the backup sets. Instead, you must provide the path to the default location or the alternate location of the disaster recovery file when the **Recover This Computer** wizard prompts you.

See [“Setting or changing the alternate location for the disaster recovery information file”](#) on page 747.

See [“Changing the default path for the disaster recovery information files”](#) on page 748.

Setting or changing the alternate location for the disaster recovery information file

Symantec recommends that you specify an alternate location where Backup Exec can store the disaster recovery information files. The disaster recovery information files contain computer-specific information for each computer that you back up with Simplified Disaster Recovery (SDR). When you recover a Backup Exec server from a disaster, you must have the Backup Exec server's disaster recovery information file available. Without it, you cannot use SDR to recover the Backup Exec server.

Symantec recommends that the alternate location be on another computer or on a different physical drive than the default location. If the Backup Exec server's hard drive is damaged, you can access a copy of the disaster recovery information from the alternate location. You can specify a drive letter that is mapped to a network share as the alternate location.

To use a remote computer's hard drive as the alternate path, establish a valid connection to the remote computer. Specify a UNC path as the alternate path, and then check the directory to ensure that the disaster recovery information files are copied.

For a clustered Backup Exec server or remote Backup Exec servers, specify an alternate location on a shared drive or on a drive that is outside the cluster.

You should also create additional copies of the disaster recovery information files and store them in a safe place.

See [“How Simplified Disaster Recovery uses disaster recovery information files”](#) on page 745.

To set or change the alternate location for the disaster recovery information file

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, select **Simplified Disaster Recovery**.
- 3 In the **Alternate path** field, enter the path where you want to store a copy of the disaster recovery information file, or click **Browse** to navigate to a path.
- 4 Specify the logon account to use to access the alternate path.
- 5 Click **OK**.

See [“Changing the default path for the disaster recovery information files”](#) on page 748.

Changing the default path for the disaster recovery information files

You can change the default path for the disaster recovery information files. However, Symantec recommends that you do not change the default path.

Copies of the disaster recovery information files are necessary to automate the recovery of a Backup Exec server.

Backup Exec automatically creates the disaster recovery information file during a backup and stores a copy of it in the following path:

```
C:\Program Files\Symantec\Backup Exec\sdr\Data\
```

If you change the disaster recovery data path, ensure that you copy the existing disaster recovery information files to the new path. The copies let you recover a computer with the previous SDR backups.

See [“How Simplified Disaster Recovery uses disaster recovery information files”](#) on page 745.

To change the default path for the disaster recovery information files

- 1 Click the Symantec Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, select **Simplified Disaster Recovery**.
- 3 In the **Path** field, change the path to the location where you want to store the disaster recovery information file, or click **Browse** to navigate to a location.
- 4 Click **OK**.

See [“Setting or changing the alternate location for the disaster recovery information file”](#) on page 747.

Disaster recovery information file data paths

Simplified Disaster Recovery (SDR) creates the disaster recovery information files that are required for recovery.

See [“How Simplified Disaster Recovery uses disaster recovery information files”](#) on page 745.

Table 18-2 Storage locations for the disaster recovery information file

Item	Description
Path	<p>Indicates a directory path where you want to store the disaster recovery information files for the computers that you back up. Backup Exec automatically creates the files after a backup job runs for which all critical system components are selected. Backup Exec then stores the disaster information recovery files in the following default location:</p> <p>C:\Program Files\Symantec\Backup Exec\sdr\Data\<computer name>.dr.</p> <p>Note: Symantec recommends that you do not change the default path.</p>
Alternate path	<p>Indicates an alternate path where you can store copies of the disaster recovery information files for the computers that you back up. Backup Exec automatically creates or updates the files after a backup job runs for which all critical system components are selected. Backup Exec then stores copies of the disaster recovery information files in this location.</p> <p>Symantec recommends that you specify an alternate path that is not on the Backup Exec server, or that is on a different physical drive than the default location. You can specify a drive letter mapped to a network share as the alternate location. If the Backup Exec server's hard drive is unavailable during a recovery, you can copy the disaster recovery information file from the alternate path to any location to recover the failed computer.</p> <p>To use a remote computer's hard drive as the alternate path, establish a valid connection to the remote computer. Specify a UNC path as the alternate path, and then check the directory to ensure that the disaster recovery information files are copied.</p>

Creating a Simplified Disaster Recovery disk image

The **Create Simplified Disaster Recovery Disk Wizard** guides you through the process of creating a startup recovery disk image that you can use to recover backed up computers.

The following items are required to create the Simplified Disaster Recovery disk image:

- The Microsoft Assessment and Deployment Kit (ADK). ADK is only supported on Windows Server 2008 and later.
 - If the Backup Exec server does not run on Windows Server 2008 or later, the **Create Simplified Disaster Recovery Disk Wizard** guides you through the download and installation of the ADK. You can then create the recovery disk by connecting to a remote Backup Exec server, or by using the Remote Administration Console.
 - If the Backup Exec server on which you want to create the Simplified Disaster Recovery disk image does not have an internet connection, and does not have ADK installed, instructions are provided on how to install the ADK. You can then use the **Create Simplified Disaster Recovery Disk Wizard** to create the recovery disk.
- 5 GB of disk space to download and install the ADK.
- 1 GB of disk space to store the required files and folders to create the ISO image.

Warning: The ADK download may take several hours.

After you run Simplified Disaster Recovery-enabled backups of the computers that you want to protect, run the **Create Simplified Disaster Recovery Disk Wizard** to create the Simplified Disaster Recovery disk image. Simplified Disaster Recovery-enabled backups are those backups that display the green indicator on the backup properties pane that indicates the critical system components on the computer are selected.

See [“How to ensure that backups are enabled for Simplified Disaster Recovery”](#) on page 743.

The wizard installs the ADK, and then uses the required files and folders from the ADK to create a startup recovery disk image in ISO format. You can then burn the image to a CD or DVD by using a third-party image burning application. For disaster recovery of a computer, you use the Simplified Disaster Recovery disk to start the computer, and then to recover it.

When you create the recovery disk, you can add language and time zone selections. You can also select the computers for which you want to add network and storage drivers. You can further customize the recovery disk by adding new OEM drivers.

Note: If you create the Simplified Disaster Recovery disk image before you run backups, the drivers for the backed-up computers are not included on the recovery disk. You must create a new recovery disk if you want the drivers of the backed-up computers to be included on the recovery disk. You can then specify the existing recovery disk as the source, and then select the computers for which you want to add the drivers to the recovery disk.

Depending on your environment, different methods are available for you to create the Simplified Disaster Recovery disk image.

Table 18-3 Methods for creating the Simplified Disaster Recovery disk image

Environment	Method
If the Backup Exec server runs on Windows Server 2008 or later	See the section called “Creating a Simplified Disaster Recovery disk image if the Backup Exec server runs on Windows Server 2008 or later” on page 751.
If the Backup Exec server runs on Windows Server 2003 or earlier If the Backup Exec server is not available	See the section called “Creating a Simplified Disaster Recovery disk image if the Backup Exec server runs on Windows Server 2003, or is not available” on page 754.
If the Backup Exec server on which you want to create the Simplified Disaster Recovery disk image does not have an Internet connection and does not have ADK installed	See the section called “Creating a Simplified Disaster Recovery disk image on a Backup Exec server that does not have an internet connection and does not have ADK installed” on page 757.

Creating a Simplified Disaster Recovery disk image if the Backup Exec server runs on Windows Server 2008 or later

If the Microsoft Assessment and Deployment Kit (ADK) is not already installed on this server, then the **Create Simplified Disaster Recovery Disk Wizard** guides you through the download and installation of the ADK. The ADK installation defaults are already selected. Symantec recommends that you use these defaults.

To create a Simplified Disaster Recovery disk image if the Backup Exec server runs on Windows Server 2008 or later

- 1 Do any of the following:

- On the **Backup and Restore** tab, in the **Restores** group, click **Create Disaster Recovery Disk**.
 - Click the Backup Exec button, click **Configuration and Settings**, and then click **Create Disaster Recovery Disk**.
 - On the **Home** tab, in the **Simplified Disaster Recovery** panel, click **Create Disaster Recovery Disk**.
 - Click the **Start** button, and then click **Create Disaster Recovery Disk**.
- 2 Review the **Welcome** panel of the **Simplified Disaster Recovery Disk Creation** wizard, and then click **Next**.
- 3 Do one of the following:

To create a Simplified Disaster Recovery disk image for the first time

Click **Create a new Simplified Disaster Recovery Disk image (.iso)**.

To select network and storage drivers from an existing Simplified Disaster Recovery disk CD or DVD, or from a folder

Do the following in the order listed:

- Click **From a previous CD or DVD, or from a folder**.
- Enter the location of the previous recovery disk, or type the path to a folder that contains the files that are extracted from the recovery disk.

To select network and storage drivers from an existing image (.iso) file

Do the following in the order listed:

- Click **From an image (.iso) file**.
- Enter the location of the recovery image file.

- 4 Click **Next**.
- 5 Do one of the following:

If the ADK is already installed on the computer

Continue with step [6](#)

If the ADK is not installed on the computer Do the following in the order listed:

- Click **Next** to continue to the ADK download and installation wizard.
- On the ADK wizard panels, click **Next** to accept the recommended defaults.

Note: The ADK download may take several hours, depending on the bandwidth.

- After the ADK is installed, click **Close** to continue with the **Create Simplified Disaster Recovery Disk Wizard**.

6 Specify the startup options that you want when you recover a computer, and then click **Next**.

7 Do one of the following:

To add network drivers and storage drivers for the computers for which you ran SDR backups Select the computers for which you want to add the drivers to the recovery disk, and then click **Next**.

To add network drivers and storage drivers that were found on the computers for which you ran SDR backups that are not already included in the source image Do the following in the order listed:

- On the **Select computers to use the drivers from** panel, click **Next**.
- On the **Drivers to Include** panel, select the drivers that you want to include in the custom recovery disk, and then click **Next**.

To add network drivers and storage drivers from OEM media Do the following in the order listed:

- On the **Select computers to use the drivers from** panel, click **Next**.
- On the **Drivers to Include** panel, click **Add Driver**, navigate to the location of the media, select it, and then click **Open**.
- When the drivers display on the **Select computers to use the drivers from** panel, click **Next**.

- 8 Type the volume label and the path for the image (.iso) file to specify the location to store the startup image, and then click **Next**.

Note: If you use any non-English characters in the volume label, then those characters are not properly displayed when you mount the .iso image.

Note: Symantec recommends that you store the image that you want to burn on the computer where the CD image or DVD image burning application is installed. By default, the location is as follows:

```
C:\Program Files\Symantec\Backup Exec\sdr\CustomSDRImage.iso
```

- 9 Review the summary panel, and then click **Create Image**.
- 10 After the Simplified Disaster Recovery disk CD or DVD image is created, click **Next**, and then click **Finish**.
- 11 Burn the startup ISO 9660 image to a CD or DVD by using a third-party image burning application.

Creating a Simplified Disaster Recovery disk image if the Backup Exec server runs on Windows Server 2003, or is not available

The Microsoft Assessment and Deployment Kit (ADK) is only supported on Windows Server 2008 and later. If the Backup Exec server does not run on Windows Server 2008 or later, or if the server is not available, you can use the Backup Exec installation media to install the **Create Simplified Disaster Recovery Disk Wizard** and the Remote Administration Console on a standalone server that runs Windows Server 2008 or later.

The **Create Simplified Disaster Recovery Disk Wizard** guides you through the download and installation of the ADK. The ADK installation defaults are already selected. Symantec recommends that you use these defaults. You can then create the recovery disk by connecting to a remote Backup Exec server, as guided by the wizard.

To create a Simplified Disaster Recovery disk image if the Backup Exec server runs on Windows Server 2003, or is not available

- 1 Insert the Backup Exec installation media into a server that runs Windows Server 2008 or later.
- 2 From the installation media browser, click **Installation**.

- 3 Click **Simplified Disaster Recovery Disk Creation Wizard**, and then click **Install**.

Note: The **Simplified Disaster Recovery Disk Creation Wizard** and the Backup Exec Remote Administration Console are installed.

- 4 Accept the terms of the license agreement, and then click **Next**.
- 5 Accept the default installation location, or specify a new location, and then click **Next**.
- 6 By default, after the computer is restarted, the **Create Simplified Disaster Recovery Disk Wizard** starts; if you unchecked the default, to start the wizard, click **Start > All Programs > Symantec Backup Exec > Create Disaster Recovery Disk**.
- 7 On the **Welcome** panel of the **Simplified Disaster Recovery Disk Creation** wizard, click **Connect to a Backup Exec server**.
- 8 Specify a Backup Exec server that runs on Windows Server 2003, enter your credentials and the domain, and then click **Next**.
- 9 Do one of the following:

To create a Simplified Disaster Recovery disk image for the first time

Click **Create a new Simplified Disaster Recovery Disk Image (.iso)**.

To select network and storage drivers from an existing Simplified Disaster Recovery disk CD or DVD, or from a folder

Do the following in the order listed:

- Click **From a previous CD or DVD, or from a folder**.
- Enter the location of the previous recovery disk, or type the path to a folder that contains the files that are extracted from the recovery disk.

To select network and storage drivers from an existing image (.iso) file

Do the following in the order listed:

- Click **From an image (.iso) file**.
- Enter the location of the recovery image file.

- 10 Click **Next**.
- 11 Do one of the following:

If the ADK is already installed on this computer

Continue with step [12](#)

If the ADK is not installed on this computer Do the following in the order listed:

- Click **Next** to continue to the ADK download and installation wizard.
- On the ADK wizard panels, click **Next** to accept the recommended defaults.
Note: The ADK download may take several hours.
- After the ADK is installed, click **Close** to continue with the **Create Simplified Disaster Recovery Disk Wizard**.

12 Specify the startup options that you want when you recover a computer, and then click **Next**.

13 Do one of the following:.

To add network drivers and storage drivers from the computers for which you ran SDR backups

Select the computers for which you want to add the drivers to the recovery disk, and then click **Next**

To add network drivers and storage drivers that were found on the computers for which you ran SDR backups that are not already included in the source image

Do the following in the order listed:

- On the **Select computers to use the drivers from** panel, click **Next**.
- On the **Drivers to Include** panel, select the drivers that you want to include on the custom recovery disk, and then click **Next**.

To add network drivers and storage drivers from OEM media

Do the following in the order listed:

- On the **Select computers to use the drivers from** panel, click **Next**.
- On the **Drivers to Include** panel, click **Add Driver**, navigate to the location of the media, select it, and then click **Open**.
- When the drivers display on the **Select computers to use the drivers from** panel, click **Next**.

To add a computer to the list

Do the following in the order listed:

- Click **Add DR file**.
- Navigate to the location of the DR file, select it, and then click **OK**.
- On the **Select computers to use the drivers from** screen, select the computer that you added, and then click **Next**.

- 14 Type the volume label and the path for the image (.iso) file to specify the location to store the startup image, and then click **Next**.

Note: Symantec recommends that you store the image that you want to burn in a location on the computer where the CD image or DVD image burning application is installed. By default, the location is as follows:

```
C:\Program Files\Symantec\Backup Exec\sdr\CustomSDRImage.iso
```

- 15 Review the summary panel, and then click **Create Image**.
- 16 After the Simplified Disaster Recovery disk CD or DVD image is created, click **Next**, and then click **Finish**.
- 17 Burn the startup ISO 9660 image to a CD or DVD by using a third-party image burning application.

Creating a Simplified Disaster Recovery disk image on a Backup Exec server that does not have an internet connection and does not have ADK installed

You can download the ADK executable file to a computer that has Internet access, and then copy the downloaded files to the computer on which you want to install ADK. Then, you can run the **Simplified Disaster Recovery Disk Creation** wizard to create the recovery disk.

To create a Simplified Disaster Recovery disk image on a Backup Exec server that does not have an Internet connection and does not have ADK installed

- 1 On a computer that has an Internet connection, click the following link to download the adksetup.exe:
<http://go.microsoft.com/fwlink/?LinkID=219682>
- 2 Run the adksetup.exe.

- 3 On the ADK **Specify Location** panel, click **Download the Assessment and Deployment Kit for installation on a separate computer**, enter the path to the Backup Exec server to which you want to install the ADK components, and then click **Next**.
- 4 Click **Yes** or **No** to join the Microsoft Customer Experience Improvement Program, and then click **Next**.
- 5 Click **Accept** to accept the terms of the license agreement, and to start the ADK download.
- 6 After the download is complete, verify that the adksetup.exe file and the Installers folder have been downloaded to the computer.
- 7 Copy the downloaded files and folder to the computer on which you want to install ADK.

Note: ADK is only supported on Windows Server 2008 and later.

- 8 Start adksetup.exe from the copied folder.
- 9 On the ADK **Specify Location** panel, click **Install the Assessment and Deployment Kit to this computer**, enter the path to which you want to install the ADK, and then click **Next**.
- 10 Click **Accept** to accept the terms of the license agreement.
- 11 On the ADK **Select the features you want to install** panel, click **Deployment Tools** and **Windows Preinstallation Environment (Windows PE)**, and then click **Install**.
- 12 After the installation is complete, on the Backup Exec server to which you installed the ADK, do any of the following:
 - On the **Backup and Restore** tab, in the **Restores** group, click **Create Disaster Recovery Disk**.
 - Click the Backup Exec button, click **Configuration and Settings**, and then click **Create Disaster Recovery Disk**.
 - On the **Home** tab, in the **Simplified Disaster Recovery** panel, click **Create Disaster Recovery Disk**.
 - Click the **Start** button, and then click **Create Disaster Recovery Disk**.
- 13 Review the **Welcome** panel of the **Simplified Disaster Recovery Disk Creation** wizard, and then click **Next**.
- 14 Do any of the following:

To create a Simplified Disaster Recovery disk image for the first time

Click **Create a new Simplified Disaster Recovery Disk image (.iso)**.

To select network and storage drivers from an existing Simplified Disaster Recovery disk CD or DVD, or from a folder

Do the following in the order listed:

- Click **From a previous CD or DVD, or from a folder**.
- Enter the location of the previous recovery disk, or type the path to a folder that contains the files that are extracted from the recovery disk.

To select network and storage drivers from an existing image (.iso) file

Do the following in the order listed:

- Click **From an image (.iso) file**.
- Enter the location of the recovery image file.

15 Click **Next**.

16 Specify the startup options that you want when you recover a computer, and then click **Next**.

17 Do one of the following:

To add network drivers and storage drivers for the computers for which you ran SDR backups

Select the computers for which you want to add the drivers to the recovery disk, and then click **Next**.

To add network drivers and storage drivers that were found on the computers for which you ran SDR backups that are not already included in the source image

Do the following in the order listed:

- On the **Select computers to use the drivers from** panel, click **Next**.
- On the **Drivers to Include** panel, select the drivers that you want to include in the custom recovery disk, and then click **Next**.

To add network drivers and storage drivers from OEM media

Do the following in the order listed:

- On the **Select computers to use the drivers from** panel, click **Next**.
- On the **Drivers to Include** panel, click **Add Driver**, navigate to the location of the media, select it, and then click **Open**.
- When the drivers display on the panel, click **Next**.

- 18 Type the volume label and the path for the image (.iso) file to specify the location to store the startup image, and then click **Next**.

Note: Symantec recommends that you store the image that you want to burn on the computer where the CD image or DVD image burning application is installed. By default, the location is as follows: `C:\Program Files\Symantec\Backup Exec\sdr\CustomSDRImage.iso`

- 19 Review the summary panel, and then click **Create Image**.
- 20 After the Simplified Disaster Recovery disk CD or DVD image is created, click **Next**, and then click **Finish**.
- 21 Burn the startup ISO 9660 image to a CD or DVD by using a third-party image burning application.

Contents of the Simplified Disaster Recovery disk image

The Simplified Disaster Recovery disk image provides the **Recover This Computer** wizard that guides you through the steps to recover a computer.

The Simplified Disaster Recovery disk image also includes the tools that you can use to manage the server's network configuration while in the recovery environment. Other utilities are included on the disk that you can use for additional recovery environment operations, such as dynamic loading of drivers, and accessing a command prompt.

Network tools include the following menu items:

- **Start My Networking Services**
- **Map a Network Drive**
- **Configure Network Connection Settings**
- **Run IP Config Utility**
- **Ping a Remote Computer**

Utility tools include the following menu items:

- **Gather Log Files for Symantec Technical Support**
- **View Log File**
- **Start PowerShell**
- **Start Command Prompt**
- **Edit the Windows boot.ini File**

- **Load a Driver**
- **Select Keyboard**

Simplified Disaster Recovery also provides you with a hard disk configuration tool called Advanced Disk Configuration. Advanced Disk Configuration lets you run advanced hard disk operations on the computer that you want to recover.

For example, you can do tasks such as:

- Create or delete volumes.
- Resize a volume.
- Convert a basic hard disk to a dynamic disk.
- Change or assign drive letters.
- Create spanned, striped, and mirrored volumes.

See [“Advanced Disk Configuration on the Recover This Computer Wizard”](#) on page 773.

See [“Requirements for using Simplified Disaster Recovery ”](#) on page 737.

Preparing to recover from a disaster by using Simplified Disaster Recovery

When a disaster occurs, you can use Simplified Disaster Recovery (SDR) to return the computer to its pre-disaster state.

To prepare to recover a computer, use the following steps:

Table 18-4 Preparing to recover from a disaster by using Simplified Disaster Recovery

Step	Description
Step 1	Disconnect any storage area network or cluster that is attached to the computer being recovered. Otherwise, the hard drives on those computers may also be repartitioned and reformatted.
Step 2	Plan any hardware changes to the computer to be recovered. See “Hardware replacement during disaster recovery” on page 762.

Table 18-4 Preparing to recover from a disaster by using Simplified Disaster Recovery *(continued)*

Step	Description
Step 3	<p>Review additional requirements for IBM computers if the computer to be recovered is an IBM computer.</p> <p>See “Prepare to recover IBM computers with Simplified Disaster Recovery” on page 763.</p>
Step 4	<p>Review recovery notes if you are recovering a SQL server, Exchange server, SharePoint server, or an environment in which the Central Admin Server Option is installed.</p> <p>See “Recovery notes for using Simplified Disaster Recovery with Exchange, SQL, SharePoint, CASO, Hyper-V hosts, and the Deduplication Option” on page 771.</p> <p>Review recovery notes if you are recovering Windows 8 or Windows Server 2012 storage pools and storage spaces.</p> <p>See “Recovery notes for using Simplified Disaster Recovery with storage pools and storage spaces” on page 769.</p>

When you complete these steps, you can start the recovery process.

See [“Recovering a computer with Simplified Disaster Recovery”](#) on page 764.

Hardware replacement during disaster recovery

You can use Simplified Disaster Recovery (SDR) to recover a computer that is no longer functioning. For example, if the computer’s main system board fails, you can restore the computer’s data after you replace the system board. You can also restore the data even if the new board is a different model or if it contains multiple processors.

If you plan to change the hardware in the computer being recovered, consider the following information before you use SDR to recover the computer:

Table 18-5 Hardware considerations when recovering failed computers

Item	Description
Hard drives	Any hard drives that you replace should be the same size or larger than the original drives.
System boards	After you replace a faulty system board and after you use SDR to recover the computer, you must use the system board manufacturer's driver CD to re-install additional functionality such as onboard sound and video.
Network interface cards	<p>If you change the network interface card in the computer that you want to recover, you must install the necessary network drivers. Without the network drivers, you cannot access the network if you want to use a remote Backup Exec server or remote legacy backup-to-disk folders to recover the computer. After you complete the recovery, you must install new network interface card drivers that match the network card that is currently in the computer.</p> <p>In most instances, it is not necessary to install the drivers manually. Most drivers are available on the Simplified Disaster Recovery disk image. When you create the Simplified Disaster Recovery disk image, you can customize it to include any drivers that you may need.</p>

See [“Recovering a computer with Simplified Disaster Recovery”](#) on page 764.

Prepare to recover IBM computers with Simplified Disaster Recovery

To prepare to recover an IBM computer that has an IBM ServeRAID card by using Simplified Disaster Recovery, use the following steps:

Table 18-6 Prepare to recover an IBM computer

Step	Description
Step 1	Install and configure the IBM ServeRAID controller card and ServeRAID software so that a boot volume is recognizable by the Windows operating system.
Step 2	Before you use the SDR startup media, start the server by using the IBM server's ServeRAID Configuration and Management CD in the CD-ROM drive. This starts the IBM ServeRAID utilities configuration and installation process to view and update the current BIOS and firmware levels.

Refer to the IBM ServeRAID documentation for complete installation instructions for installing Windows on an IBM Server with the ServeRAID controller. Create and initialize the ServeRAID disks so that volumes are recognizable by the Windows operating system.

See [“Recovering a computer with Simplified Disaster Recovery”](#) on page 764.

Recovering a computer with Simplified Disaster Recovery

You can use the Simplified Disaster Recovery **Recover This Computer Wizard** to run a local recovery or a remote recovery of a Backup Exec server or a Windows computer.

You can perform a local recovery if all of the following conditions are met:

- You want to restore a Backup Exec server or a Windows computer.
- The backup data for the computer is located on devices that you can locally attach to the computer on which you are run the **Recover This Computer Wizard**.
- You can provide the disaster recovery information file for the Backup Exec server or the Windows computer that you want to recover.

Note: If the computer that you want to recover was backed up to a tape device, to deduplication storage, or to a virtual disk, then SDR was unable to store the disaster recovery information file with the backup sets. You must provide the path to the default location or to the alternate location for the disaster recovery information file when the **Recover This Computer Wizard** prompts you. If you cannot provide a disaster recovery information file, then you must perform a manual disaster recovery. Click the following URL for information on how to perform a manual disaster recovery:

<http://entsupport.symantec.com/umi/V-290-706>

You can perform a remote recovery if all of the following conditions are met:

- You want to restore a Backup Exec server or a Windows computer.
- The backup data for the computer is located on devices that are attached to a remote Backup Exec server
- You have network access to the remote Backup Exec server.

Before you start, review the steps for preparing to recover.

See “[Preparing to recover from a disaster by using Simplified Disaster Recovery](#)” on page 761.

During recovery operations, you can recover the computer from the most recent backup. Or, you can recover to a previous point-in-time backup.

The **Recover This Computer Wizard** supports the recovery of computers with encrypted backup sets. If the Simplified Disaster Recovery (SDR) backups are encrypted during backup, the wizard prompts you for the pass phrase of each encrypted backup set that is required to complete the recovery.

See “[Encryption key management](#)” on page 586.

To restore data by using the **Recover This Computer Wizard**, the following items are required:

- The Simplified Disaster Recovery disk image, which must be the same version of Backup Exec as the Backup Exec server.
- If the computer that you want to recover was backed up to a tape device, deduplication storage, or to a virtual disk, then the disaster recovery information file is not stored with the backup sets. Instead, you must provide the path to the default location or to the alternate location when the **Recover This Computer Wizard** prompts you. The default location is: C:\Program Files\Symantec\Backup Exec\sdr\Data. If the file is unavailable, you cannot recover the computer with SDR.

- The backup set that contains all of the critical system components for the computer that you want to restore.
- The boot drive on the computer that you want to recover must have from 3 GB to 5 GB of free space, depending on the operating system and configuration.

Note: If a blank screen appears and the computer does not restart after you run SDR, ensure that the boot drive has the necessary amount of free space. Then, restart the computer again.

- The computer that you want to recover must have an amount of RAM that is equal to or greater than the minimum that the restored operating system requires.
- The computer that you want to recover must have disks with enough free space to contain the data that you restore from the backup of critical volumes. A volume is considered critical if it is required for the computer to start successfully.
- If you recover a Windows computer that has BitLocker encryption enabled, you must turn on BitLocker encryption after the restore. See your Microsoft documentation for more information on BitLocker drive encryption.
- If the computer that you want to recover contains a RAID setup, you may be required to configure the RAID before you start it with the SDR disk. Use the computer manufacturer's RAID software to configure the RAID system.
- If appropriate, review the following recovery notes:
 - See [“Recovery notes for using Simplified Disaster Recovery with storage pools and storage spaces”](#) on page 769.
 - See [“Recovery notes for using Simplified Disaster Recovery with Exchange, SQL, SharePoint, CASO, Hyper-V hosts, and the Deduplication Option”](#) on page 771.

Note: Boot managers, such as System Commander or the OS/2 Boot Manager, cannot be restored with SDR. Boot managers are usually installed at a very low level that Backup Exec cannot protect. For example, the OS/2 Boot Manager resides in its own hard drive volume that Backup Exec cannot access. Because of the many different boot managers available, you may not be able to restart the computer after an SDR recovery, even though the operating system was restored. You can reinstall the boot manager to resolve this issue.

To recover a computer by using the Recover This Computer Wizard

- 1 Place the startup Simplified Disaster Recovery Disk in the CD drive or DVD drive of the computer that you want to recover, and then start the computer.
- 2 On the **End User License Agreement** screen, click **Accept**.
- 3 On the Simplified Disaster Recovery **Welcome** screen, click **Recover This Computer**.
- 4 On the left, click **Network**, and then click **Configure Network Connection Settings**.
- 5 Select the appropriate network adapter configuration, and then click **OK**.
- 6 Do one of the following:

If the backup data for this computer is located on devices that are attached to a remote Backup Exec server (remote recovery)

Do the following in the order listed:

- Click **The data is located on devices attached to a remote Backup Exec server**
- Enter the name and domain of the remote Backup Exec server where the backup data is located, and the appropriate administrator or administrator-equivalent credentials.
- (Optional) To configure network adapter settings, click **Configure network adapter settings**, and then do one of the following:
 - To assign a static IP address for each detected network adapter, select the appropriate options, and then click **OK**.
 - To configure an IPv6 network controller, click **Configure IPv6**, select the appropriate options, and then click **OK**.
- (Optional) To load network adapter drivers, click **Load network adapter drivers**. Then, click **Install Driver** for any inactive network controller that the wizard detects. Navigate to the device that contains the network controller driver, and then click **Open**. Select the driver, and then click **Open** again.
- Click **Next**.

If the computer that you want to recover was backed up to locally attached devices such as tape drives, robotic libraries, disk storage devices, or disk cartridge devices (local recovery)

Do the following in the order listed:

- Click **The data is located on devices locally attached to this computer.**
- Enter the appropriate administrator or administrator-equivalent credentials when prompted, and then click **Next**.

- 7 Select the disaster recovery information file that you want to use, and then click **Next**.
- 8 Select the backup sets that you want to use to recover the computer, and then click **Next**.
- 9 Select the storage device that contains the backup data for the computer, and then click **Next**.
- 10 Do any of the following:

To use the volume layout as shown

Click **Next**.

To view the disk geometry as it presently exists., or to view a graphical representation of proposed changes

Click **Preview**.

To install required SCSI or RAID controller drivers

Do the following in the order listed:

- Click **Load Storage controller drivers**.
- Click **Install Driver** for any inactive network controller that the wizard detects.
- Navigate to the device that contains the network controller driver, and then click **Open**.
- Select the driver, and then click **Open** again.

To automatically create a volume layout on the available hard disks if mismatched volumes appear in the simplified volume layout view

Check **Erase hard disks and recreate the volume layout shown above**. Existing data on these disks will be lost.

To change volume sizes or other disk-related operations

Click **Advanced Disk Configuration** .

See [“Advanced Disk Configuration on the Recover This Computer Wizard”](#) on page 773.

- 11

Review the **Recovery Summary** and click **Back** to make any changes, or click **Recovery** to start the recovery process.
- 12

When the recovery is complete, do any or all of the following:
- To troubleshoot any issues that occurred during the hardware discovery phase

Click **View Hardware Discovery log**.

To troubleshoot any issues that occurred during the recovery process

Click **View Recover This Computer log**.

To restart this computer after you click **Finish**

Select **Restart this computer**.

To complete the recovery process and close the wizard

Click **Finish**.

Note: If the recovered computer contains multiple hard disks, ensure that the computer's BIOS is configured to start the computer from the hard disk that contains the Windows operating system. On many computers, the BIOS menu can be accessed by pressing F2 or DEL on the computer keyboard.

Recovery notes for using Simplified Disaster Recovery with storage pools and storage spaces

You can perform Simplified Disaster Recovery (SDR) for Windows 8 storage spaces and Windows Server 2012 storage pools and storage spaces using either of the scenarios that are described in the following table

Table 18-7 Scenarios for SDR for storage pools and storage spaces

Scenario	Description
Restore Windows 8 storage spaces and Windows Server 2012 storage pools and storage spaces to the same computer	<p>In this scenario, the Recover This Computer Wizard prompts you to reconfigure storage pools and storage spaces, or remove them, or to restore them as they are.</p> <p>To reconfigure or remove the storage pools and storage spaces, click the PowerShell button on the displayed message dialog box to open a PowerShell window.</p> <p>For more information about storage commands in PowerShell, see the following URL:</p> <p>http://technet.microsoft.com/library/hh848705.aspx</p>
Restore a version of Windows that is before Windows 8/Windows Server 2012 to a computer on which storage pools and storage spaces or virtual disks are already configured.	<p>In this scenario, you restore to a computer that contains configured storage pools and storage spaces. However, the operating system that you restore does not recognize storage pools and storage spaces.</p> <p>The Recover This Computer Wizard prompts you to do either of the following:</p> <ul style="list-style-type: none">■ To remove the storage pools and storage spaces, click the PowerShell button on the displayed message dialog box to open a PowerShell window. <p>For more information about storage commands in PowerShell, see the following URL:</p> <p>http://technet.microsoft.com/library/hh848705.aspx</p> <ul style="list-style-type: none">■ To continue to restore the storage pools and storage spaces, click Next on the displayed message dialog box. The Recover This Computer Wizard blocks any volume operations on the virtual disk, such as creating a volume, deleting a volume, assigning a drive letter, and so on.

Limitations for restoring storage pools and storage spaces

The **Recover This Computer Wizard** automatically maps volumes to their original virtual disks or storage spaces if you restore to the same computer. If you do not

restore to the same computer, then you must use PowerShell and the SDR Advanced Disk Configuration utility to manually map the volumes to the virtual disks or storage spaces.

The **Recover This Computer Wizard** does not allow Boot, System, and Recovery volumes to be mapped to virtual disks.

Simplified Disaster Recovery does not support Windows Server 2012 R2 storage pools.

See [“Advanced Disk Configuration on the Recover This Computer Wizard”](#) on page 773.

Recovery notes for using Simplified Disaster Recovery with Exchange, SQL, SharePoint, CASO, Hyper-V hosts, and the Deduplication Option

Review the following recovery notes if you recover a SQL Server, Exchange Server, SharePoint server, or an environment in which the Central Admin Server Option is installed.

Microsoft SQL Server recovery notes

After you use Simplified Disaster Recovery (SDR) to recover the Windows server, SDR automatically replaces the damaged master and model databases with copies of the master database and the model database. After SQL is restarted and the latest master database backup and all other system databases are restored, you must still restore all user databases. Use the Backup Exec Restore wizard to restore the latest backups.

Microsoft Exchange recovery notes

After you use SDR to recover the Windows server, use the Backup Exec Restore wizard to restore the Exchange Server databases from the most recent Exchange Server database backups.

SharePoint Portal Server recovery notes

You can use SDR to recover a Windows server that has SharePoint Portal Server installed. After you restore the Windows computer, you must restart it. After the computer restarts, however, the SharePoint Portal Server software is installed but is not functional. You must remove SharePoint Portal Server and reinstall it before the SharePoint data can be restored.

Central Admin Server Option (CASO) recovery notes

When you use SDR to recover a computer in a CASO environment, you can submit the remote restore job to either of the following:

- The central administration server.
- The managed Backup Exec server that performed the original backup job.

Microsoft Hyper-V hosts

After you use SDR to recover the Windows server, use the Backup Exec Restore wizard to restore the Microsoft Hyper-V host from the most recent Microsoft Hyper-V backups.

Backup Exec Deduplication Option recovery notes

If the Backup Exec Deduplication Option is installed and you use deduplication disk storage devices, consider the following:

- If you back up the local Backup Exec server, do not select a deduplication disk storage device on the server as the destination storage device. The **Recover This Computer** wizard cannot restore data from a local deduplication disk storage device.
- Simplified Disaster Recovery (SDR) cannot recover a deduplication disk storage device.
- Before you can use SDR to restore a remote computer that was backed up with client-side deduplication, you must first delete the direct access device.
See [“Selecting storage devices for direct access sharing”](#) on page 817.

If you use SDR to recover a Backup Exec server that contains a deduplication disk storage device, consider the following:

- Any existing backup sets that were sent to the deduplication disk storage device after it was backed up cannot be restored.
- The deduplication disk storage device may not be in an operational state after the recovery.

You can also use manual disaster recovery to recover deduplication disk storage devices and OpenStorage devices.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

See [“Disaster recovery of deduplication disk storage devices”](#) on page 827.

See [“Disaster recovery information file data paths”](#) on page 748.

See [“Setting or changing the alternate location for the disaster recovery information file”](#) on page 747.

See [“About the Central Admin Server Option”](#) on page 1096.

Advanced Disk Configuration on the Recover This Computer Wizard

The **Recover This Computer Wizard** restores the hard drive volumes to the same sizes they were before the disaster. If the hard drive in the failed computer is larger than the hard drive that was in place before the disaster, it may result in unused and unallocated space. You can run Advanced Disk Configuration to alter the volume sizes to reflect the larger hard drive size.

The following is an example of why the hard drive volumes should be resized:

The pre-disaster computer hardware contains a 40 GB hard drive with two 20-GB volumes. You replace it with a 90-GB model. SDR then uses the disaster recovery information file to rebuild the hard disk partition table by using the partition information that is found on the original 40-GB hard drive. As a result, only 40 GB of space is allocated on the new 90 GB hard drive, with a partition map that consists of two 20-GB partitions.

You can access Advanced Disk Configuration from within the **Recover This Computer Wizard**.

Note: You should be familiar with Microsoft Disk Management concepts before you run Advanced Disk Configuration.

The following table provides details about the additional disk-related operations that you can do with Advanced Disk Configuration.

Table 18-8 Advanced Disk Configuration tasks

Task	Description
Create a simple volume	A simple volume is a partition on a disk that contains a file system.
Format a volume	Disk volumes must be formatted before data can be stored on them.
Extend the size of a volume	If a disk contains some unallocated disk space that is adjacent to a functional volume, you can extend the volume to include the free space. To extend the volume, it must be either raw or formatted with the Windows NTFS file system.

Table 18-8 Advanced Disk Configuration tasks (*continued*)

Task	Description
Shrink the size of a volume	<p>You can decrease the size of a volume by shrinking the volume into the contiguous, unallocated disk space that is on the same disk.</p> <p>When you shrink a volume, there is no need to reformat the volume. Ordinary files are automatically relocated on the disk to create the new, unallocated disk space.</p>
Create a spanned volume	<p>A spanned volume spans more than one physical disk. You can create a spanned volume by spanning it across multiple physical disks, or by spanning the volume into unallocated disk space.</p> <p>To create a spanned volume, you must have a startup volume and at least two dynamic volumes.</p> <p>Note: Spanned volumes are not fault-tolerant.</p>
Create a striped volume	<p>Striped volumes store data in stripes across two or more physical disks. Although striped volumes do not provide fault-tolerance protection, they do offer the best performance of all the volumes in Windows.</p>
Create a mirrored volume	<p>A mirrored volume provides a copy of the data that is written to a selected volume. Because all data is written to both the mirrored volume and the selected volume, mirroring reduces the capacity of both volumes by 50%.</p>
View volume properties	<p>You can view properties for each volume in the Current Disk Layout view or in the Original Disk Layout view.</p>
Change an assigned drive letter	<p>You can change assigned drive letters for all volumes if you want to organize your drive letters in a certain way.</p>

Table 18-8 Advanced Disk Configuration tasks (*continued*)

Task	Description
Delete a volume	Deleting a volume erases all data from the volume so Symantec recommends caution when considering the use of this option.
Convert a basic disk to a dynamic disk	Converting basic disks to dynamic disks lets you create the volumes that span multiple disks. Dynamic disks also let you create fault-tolerant volumes, such as mirrored volumes and RAID-5 volumes. All volumes on dynamic disks are referred to as dynamic volumes.
Convert a Master Boot Record (MBR) disk to a Guid Partition Table (GPT) disk	<p>MBR disks use the standard BIOS interface. GPT disks use extensible firmware interface (EFI).</p> <p>You can convert MBR disks to GPT disks if the disk does not contain partitions or volumes.</p>
Convert a Guid Partition Table (GPT) disk to a Master Boot Record (MBR) disk	<p>GPT disks use extensible firmware interface (EFI). Master boot record (MBR) disks use the standard BIOS interface.</p> <p>GPT disks can be converted to MBR disks if the disk does not contain partitions or volumes.</p>

Table 18-8 Advanced Disk Configuration tasks (continued)

Task	Description
View the original disk layout geometry	<p>The original disk layout shows the actual hard disk layout that existed during the backup job. Within the simplified layout view, you can accept the disk geometry as it originally existed before the disaster, or you can alter the geometry by changing the volume sizes. Depending on the size of the existing disks, you can alter volume sizes in megabytes, gigabytes, or terabytes.</p> <p>The simplified layout view has a Preview tab that lets you view the disk geometry as it presently exists. If you change the disk geometry and change volume sizes, click the Preview tab to see a graphical representation of your proposed changes.</p> <p>If mismatched volumes appear in the simplified volume layout view, you can use the option Erase hard disks and recreate the volume layout shown above to automatically create a volume layout on the available hard disks. You can also manually create a volume layout by using the Advanced Disk Configuration option.</p>

Symantec Backup Exec Agent for Windows

This appendix includes the following topics:

- [About the Agent for Windows](#)
- [Requirements for the Agent for Windows](#)
- [Stopping and starting the Agent for Windows](#)
- [Establishing a trust between the Backup Exec server and a remote computer](#)
- [About the Backup Exec Agent Utility for Windows](#)
- [Using the Backup Exec Agent Utility Command Line Applet](#)
- [Backup Exec Agent Utility Command Line Applet switches](#)

About the Agent for Windows

The Symantec Backup Exec Agent for Windows (Agent for Windows) is installed as a separate add-on component. The Agent for Windows enables Windows Servers network administrators to perform backup and restore operations on Backup Exec agents that are connected to the network.

The Agent for Windows is a system service that runs on remote Windows servers and workstations. The Agent for Windows provides faster backup processing by locally performing the tasks that require extensive network interaction in typical backup technologies. The Agent for Windows processes backup data into a continuous stream that the Backup Exec server then processes as a single task. This method provides better data transfer rates over traditional technologies, which require multiple requests and acknowledgments between the Backup Exec server and the remote server.

The Agent for Windows enables you to do the following:

- Back up and restore in firewall environments.
- Back up and restore using a specified local network if the Backup Exec server and the remote computer are on the same subnet.
- Attain significant performance increases when running modified backups (for example, differential and incremental). This occurs because file selection is performed locally by the Agent for Windows instead of across the network as performed by traditional network backup applications.

Note: Network hardware has a major effect on performance. Performance is directly related to the capabilities of the networking hardware in the Backup Exec server and the remote device. Higher network bandwidth ratings also contribute to faster operation processing.

See [“Requirements for the Agent for Windows”](#) on page 778.

See [“Methods for installing the Agent for Windows”](#) on page 86.

See [“Changing network and security options for Backup Exec”](#) on page 576.

See [“Using Backup Exec with firewalls”](#) on page 578.

See [“Backup Exec Shadow Copy Components file system”](#) on page 187.

See [“About the Backup Exec Agent Utility for Windows”](#) on page 780.

Requirements for the Agent for Windows

Because an Agent for Windows is also a Client Access License (CAL), you must install the Agent for Windows on any remote Windows computer that you want to back up. You cannot fully protect resources on a remote server until an Agent for Windows has been installed.

At the Backup Exec server, you must enter Agent for Windows licenses for each remote Windows computer that you want to protect. To back up a remote Windows computer from more than one Backup Exec server, you must enter the same Agent for Windows licenses on each Backup Exec server.

Backup Exec Agent for Applications and Databases also includes an Agent for Windows that lets you protect one remote Windows computer. The Agent for Windows license is enabled when you install the database agents on the Backup Exec server.

To protect the Workstation versions of the supported Windows platforms, you must install the Agent for Windows on each platform.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Note: If a previous version of the Agent for Windows is installed, it is automatically upgraded when you initiate a new Agent for Windows installation. Previous versions of the Agent for Windows are automatically detected on the remote computers and replaced with the new version during installation of the new Agent for Windows. The name of the system service may have changed when the upgrade is complete.

You can install the Agent for Windows using many methods, depending on your environment.

See “[Methods for installing the Agent for Windows](#)” on page 86.

Stopping and starting the Agent for Windows

The Agent for Windows is automatically started as a service when Windows is started on the remote computer.

To stop or start the Agent for Windows

1 Do one of the following:

On Windows 7/Vista/Server 2008 R2/Server 2008	Right-click Computer .
On a Windows Server 2003	Right-click My Computer .
On Windows 8	In Windows Explorer, right-click This PC .
On Windows Server 2012/Server 2012 R2	In Windows Explorer, right-click Computer .

2 Click **Manage**.

3 Do one of the following:

On Windows Server 2008 R2/Server 2008	On the Server Manager dialog box, expand Configuration .
On Windows 7/Vista/Server 2003/8/Server 2012/Server 2012 R2	On the Computer Management dialog box, double-click Services and Applications .

4 Click **Services**.

5 In the **Results** pane, right-click **Backup Exec Remote Agent for Windows**.

6 Do one of the following:

To stop the Agent for Windows Click **Stop**.

To start the Agent for Windows Click **Start**.

See [“About the Agent for Windows”](#) on page 777.

Establishing a trust between the Backup Exec server and a remote computer

When you connect to a remote computer from the Backup Exec server, you must establish a trust between the Backup Exec server and the remote computer to ensure secure communication. You must also establish this trust if you want to configure a remote computer to perform client-side deduplication. You should manually verify the identity of the remote computer to ensure that the remote computer is a trusted source before you establish the trust. After you verify that the remote computer is a trusted resource, you can establish the trust with the Backup Exec server.

Backup Exec issues a security certificate for both the Backup Exec server and the remote computer. The security certificate is valid for approximately one year and is automatically renewed during normal operations. However, if the certificate expires you must re-establish the trust.

You can establish a trust between the Backup Exec server and the remote computer by adding the remote computer to the list of servers on the **Backup and Restore** tab.

To establish a trust for a remote computer

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Click Microsoft Windows computers and servers.
- 3 Follow the on-screen prompts.

About the Backup Exec Agent Utility for Windows

The Backup Exec Agent Utility is installed when the Agent for Windows is installed on a remote Windows computer.

You can perform the following tasks with the Backup Exec Agent Utility:

- Start the Backup Exec Agent Utility each time you log on.
See [“Starting the Backup Exec Agent Utility”](#) on page 781.
- View current activity on the remote Windows computer.
See [“Viewing the activity status of the remote computer from the system tray”](#) on page 782.
- Configure the Agent for Windows to send information such as the version and the IP address to a Backup Exec server.
See [“About publishing the Agent for Windows to Backup Exec servers”](#) on page 784.
- Configure the Backup Exec Agent Utility for backup and restore operations of Oracle instances.
See [“Configuring an Oracle instance on Linux servers”](#) on page 1001.
- Configure the Backup Exec Agent Utility for Backup Exec server database access for Oracle operations.
See [“Configuring database access for Oracle operations”](#) on page 787.
- Remove the security certificate for a Backup Exec server.
See [“Removing Backup Exec servers that the Agent for Windows can publish to”](#) on page 786.

Starting the Backup Exec Agent Utility

You access the Backup Exec Agent Utility from the Windows taskbar.

See [“Viewing the activity status of the remote computer in the Backup Exec Agent Utility”](#) on page 782.

See [“About publishing the Agent for Windows to Backup Exec servers”](#) on page 784.

To start the Backup Exec Agent Utility

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec Agent Utility**.

When the Backup Exec Agent Utility is running, an icon appears in the system tray. You can double-click the icon to view the utility.

- 2 To open the registry editor, the Services window, and the Event Viewer on the remote Windows computer, right-click the Backup Exec Agent Utility icon in the system tray, and then click **Tools**.

Viewing the activity status of the remote computer in the Backup Exec Agent Utility

You can use the Backup Exec Agent Utility to view the activity status of the remote Windows computer.

To view the activity status of the remote computer in the Backup Exec Agent Utility

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec Agent Utility**.

If the Backup Exec Agent Utility is already running, you can double-click its icon in the system tray.

- 2 Click the **Status** tab.

You can view any of the following information about the remote Windows computer:

Backup Exec server	Displays the name of the Backup Exec server that is processing the current operation.
Source	Displays the media or share that is being processed.
Current folder	Displays the name of the current directory, folder, or database (depending on the specific agent) that is being processed.
Current file	Displays the name of the current file that is being processed.

- 3 Click **OK**.

See [“About the Backup Exec Agent Utility for Windows”](#) on page 780.

Viewing the activity status of the remote computer from the system tray

You can view the activity status for a remote computer.

Possible statuses are as follows:

- A backup job is running.
- A restore job is running.
- A backup and a restore job are running.
- Snapshot in progress.

- The Backup Exec client service named Beremote.exe is not running on the computer.
- Idle.

To view the activity status of a remote computer

- ◆ Position the cursor over the Agent for Windows icon in the system tray.
- See [“About the Backup Exec Agent Utility for Windows”](#) on page 780.

Starting the Backup Exec Agent Utility automatically on the remote computer

You can start the Backup Exec Agent Utility automatically each time you log on to the remote computer.

To start the Backup Exec Agent Utility automatically on the remote computer

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec Agent Utility**.

If the Backup Exec Agent Utility is already running, you can double-click its icon in the system tray.

- 2 Click the **Status** tab.
- 3 Check the **Start the Backup Exec Agent Utility every time you log on** check box.
- 4 Click **OK**.

See [“About the Backup Exec Agent Utility for Windows”](#) on page 780.

Setting the refresh interval on the remote computer

You can display the number of seconds for the Backup Exec Agent Utility to wait before refreshing the status of the computer.

To set the refresh interval on the remote computer

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec Agent Utility**.

If the Backup Exec Agent Utility is already running, you can double-click its icon in the system tray.

- 2 Click the **Status** tab.

3 In the **Refresh interval** box, type the number of seconds to refresh the status.

4 Click **OK**.

See [“About the Backup Exec Agent Utility for Windows”](#) on page 780.

About publishing the Agent for Windows to Backup Exec servers

Use the Backup Exec Agent Utility to add, change, or delete the Backup Exec server names or IP addresses that this remote Windows computer publishes to. Each Backup Exec server that you add to the list on the **Publishing** tab displays the remote computer in the list of servers on the **Backup and Restore** tab. After the remote computer is added to the list of servers, you can right-click the remote computer and click **Establish Trust** to verify that the remote computer is a trusted resource.

This information that the Agent for Windows publishes includes the version of the Agent for Windows and the remote computer's IP addresses. Because the remote computer's IP address is published to the Backup Exec server, the Backup Exec server can connect to and display the remote computer even if it is in an unknown domain.

For each Backup Exec server that is published to, you can specify a local backup network for operations between the Backup Exec server and the remote computer. Directing jobs to a specified local network rather than to a corporate network isolates the backup data traffic so that other connected networks are not affected when operations are performed between the Backup Exec server and the remote computer.

See [“About the Backup Exec Agent Utility for Windows”](#) on page 780.

See [“Backup networks”](#) on page 574.

See [“About the list of servers on the Backup and Restore tab”](#) on page 152.

See [“Adding Backup Exec servers that the Agent for Windows can publish to”](#) on page 784.

See [“Editing Backup Exec server information that the Agent for Windows publishes to”](#) on page 785.

See [“Removing Backup Exec servers that the Agent for Windows can publish to”](#) on page 786.

Adding Backup Exec servers that the Agent for Windows can publish to

You can use the Backup Exec Agent Utility to add a Backup Exec server that the Agent for Windows can publish information.

See [“About publishing the Agent for Windows to Backup Exec servers”](#) on page 784.

See [“Viewing the activity status of the remote computer from the system tray”](#) on page 782.

To add Backup Exec servers that the Agent for Windows can publish to

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec Agent Utility**.

When the Backup Exec Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 Click the **Publishing** tab.
- 3 (Optional) The first time that you start the Backup Exec Agent Utility, click **Change Settings** to enable the options.
- 4 Click **Add**.
- 5 Enter the following information:

Backup Exec server name or IP address	Type the Backup Exec server name or the IP address of the Backup Exec server to which you want to publish information.
User Name	Type the user name for an account that has administrative rights on the Backup Exec server.
Password	Type the password for an account that has administrative rights on the Backup Exec server.

- 6 Click **OK**.

Editing Backup Exec server information that the Agent for Windows publishes to

You can use the Backup Exec Agent Utility to edit a Backup Exec server name or IP address to which the Agent for Windows can publish information.

See [“About publishing the Agent for Windows to Backup Exec servers”](#) on page 784.

To edit Backup Exec server information

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec Agent Utility**.

When the Backup Exec Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 Click the **Publishing** tab.
- 3 (Optional) The first time that you start the Backup Exec Agent Utility, click **Change Settings** to enable the options.
- 4 Select the Backup Exec server that you want to edit from the list.
- 5 Click **Edit**.
- 6 Edit the Backup Exec server name or IP address.
- 7 Click **OK**.

Removing Backup Exec servers that the Agent for Windows can publish to

You can use the Backup Exec Agent Utility to remove a Backup Exec server so that the Agent for Windows no longer publishes information to it.

See [“About publishing the Agent for Windows to Backup Exec servers”](#) on page 784.

To remove Backup Exec servers that the Agent for Windows can publish to

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec Agent Utility**.

When the Backup Exec Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 Click the **Publishing** tab.
- 3 (Optional) The first time that you start the Backup Exec Agent Utility, click **Change Settings** to enable the options.
- 4 Select the Backup Exec server that you want to remove from the list.
- 5 Click **Remove**.
- 6 Click **OK**.

Configuring database access for Oracle operations

You can configure database access to enable the Backup Exec server to authenticate Oracle operations.

See [“Setting authentication credentials on the Backup Exec server for Oracle operations”](#) on page 1006.

To configure database access for Oracle operations

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec Agent Utility**.
- 2 Click the **Database Access** tab.
- 3 (Optional) The first time that you start the Backup Exec Agent Utility, click **Change Settings** to enable the options.

4 Complete the appropriate options to configure database access:

Enable the Backup Exec server to authenticate Oracle operations	Select this option to enable Oracle operations between the Backup Exec server and this computer.
User name	<p>Specify a user name that has administrative rights to this computer. This logon account is what the Backup Exec server uses when it connects to this computer.</p> <p>If you specify an IP address or a fully qualified computer name as part of the user name, the Backup Exec Agent Utility may not be able to verify the user account. If the credentials entered are incorrect, the error “cannot attach to a resource” may be displayed when you run a backup or restore job.</p> <p>You must add this computer name and logon account to the Backup Exec server's list of authentication credentials for Oracle servers. If the authentication fails when the Oracle resources are backed up, the backup job fails. If the authentication fails when you are browsing the backup sets for a restore job, then the backup sets become unavailable, and you must run a DBA-initiated restore job to restore data.</p>
Password	<p>Specify the password for this logon account.</p> <p>Note: For security reasons, the logon credentials are not stored on the remote computer.</p>
Confirm Password	Type the password again to confirm it.
Use a custom port to connect to the Backup Exec server during Oracle operations	<p>Select this option to change the port that is used for communications between this computer and the Backup Exec server during Oracle operations. By default, port 5633 is used.</p> <p>If you change the port number on this computer, you must also change it on the Backup Exec server, and then restart the Backup Exec Job Engine Service on the Backup Exec server.</p>
Port number	Type the port number that you want to use for communications between this computer and the Backup Exec server.

- 5 Click **OK**.
 - 6 On the Backup Exec server, add the name of the Oracle server and the user name that you entered on the **Database Access** tab to the Backup Exec server's list of authentication credentials.
- See [“About the Backup Exec Agent Utility for Windows”](#) on page 780.

Removing a security certificate for a Backup Exec server that has a trust with the Agent for Windows

You can remove the security certificate for a Backup Exec server that has established a trust with the Agent for Windows.

To remove a security certificate for a Backup Exec server

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec Agent Utility**.
- 2 Click the **Security** tab.
- 3 (Optional) The first time that you start the Backup Exec Agent Utility, click **Change Settings** to enable the options.
- 4 Select the Backup Exec server that you want to remove the security certificate for, and then click **Remove**.
- 5 Click **OK**.

See [“Establishing a trust between the Backup Exec server and a remote computer”](#) on page 780.

Using the Backup Exec Agent Utility Command Line Applet

You can use the Backup Exec Agent Utility Command Line Applet from any Windows operating system command prompt to access the Backup Exec Agent Utility. The Backup Exec Agent Utility Command Line Applet is installed when you install the Agent for Windows. If you run the command line utility on a Windows Vista/Server 2008/Server 2008/7/8/Server 2012/Server 2012 R2 computer, you must run it in elevated command prompt.

Note: To run the Backup Exec Agent Utility Command Line Applet on a Microsoft Windows Server 2008/Server 2012/Server 2012 R2 computer, you must use Server Core.

You can run the following Backup Exec Agent Utility functions with the Backup Exec Agent Utility Command Line Applet:

- Set the publishing interval (in minutes).
- List the published name for the agent.
- List the Backup Exec server names to which the agent is publishing.
- Add a Backup Exec server to the publishing list.
- Remove a Backup Exec server from the publishing list.
- View the following status information:
 - Activity status
 - Current source
 - Current folder
 - Current file
 - Currently attached Backup Exec server

To use the Backup Exec Agent Utility Command Line Applet

- 1 Open a command prompt.
- 2 From the Backup Exec installation directory, type `ramcmd.exe` followed by a series of command switches.

The default installation location is `c:\Program Files\Symantec\Backup Exec\RAWS`

See [“Backup Exec Agent Utility Command Line Applet switches”](#) on page 790.

Backup Exec Agent Utility Command Line Applet switches

The following table describes the switches that you can use with the Backup Exec Agent Utility Command Line Applet.

See [“Using the Backup Exec Agent Utility Command Line Applet”](#) on page 789.

Table A-1 Backup Exec Agent Utility Command Line Applet switches

Switch	Description
status:[n]	<p>Status output is repeated every <n> seconds, with a range of 1 - 86400. Press Q to stop the output from running.</p> <p><i>ramcmd /status:[n]</i></p> <p>When you use the /status switch without a time value, the Agent for Windows status appears in the command window and then the applet exits.</p>
/publish:[on off add remove interval][/ms:<Backup Exec server>] [/t:<x>]	<p>Use the following parameters with the /publish switch:</p> <ul style="list-style-type: none">■ No parameter specified- Displays the publishing status and then exits.■ [on] - Turns publishing on. Lets the Agent for Windows send information about itself, such as its version and IP address.■ [off] - Turns publishing off.■ [add], [remove] - Used with /ms. You can use this parameter to add or remove Backup Exec servers from the Agent for Windows publish list.■ [interval] - Used with /t. Specifies the time interval that the Agent for Windows sends information about itself to the Backup Exec server. <p>You can set the time interval in minutes using the /t:<x> parameter.</p> <p>Note: The [interval] switch must be used with the /t: switch. Using [interval] alone on the command line is not supported.</p> <p><i>ramcmd /publish:[on off add remove interval] [/ms<Backup Exec server>] [/t:<x>]</i></p>

Table A-1 Backup Exec Agent Utility Command Line Applet switches
(continued)

Switch	Description
/oracle: [new edit delete] /in:[<instance name>] /ms:[<Backup Exec server address>] /jt:[<job template>] /user:[<username>] /password:[<password> *] /rc: [yes no] /tns:[<TNS name>]	<p>Use the following parameters with the /oracle switch:</p> <ul style="list-style-type: none"> ■ No parameter specified- Displays the existing Oracle instances and then exits. ■ [new], [edit], [delete] - Used with switch /in. ■ /in:[<instance name>] - Used to add, edit, and delete Oracle instance names from the Oracle instance list. ■ /ms:[<Backup Exec server name address>] - Sets the Backup Exec server name or its IP address. ■ /jt:[<job template>] - Sets a Backup Exec job template. ■ /user:[<username>] - Sets a username. ■ /password:[<password> *] - Sets a password to be used with /user:[<username>]. If you omit the password, or you use an asterisk [*], you do not need to enter the password on the command line. After the command runs, a prompt appears asking you for a password. ■ /rc:[yes no] - Turns the Use recover catalog setting on or off. If /rc appears without a parameter, then the current status for that instance is displayed. ■ /tns:[TNS name] - Sets the TNS name alias of an available Oracle database and the server it resides on in the Oracle TNSNAMES file. <pre>ramcmd.exe /oracle:edit /in:<instance name> /rc:[yes no] [/tns:<TNS name>] [/user:<username>] [/password:password *]</pre>
/auth:[on off] [/user:<username>] [/password:<password> *]	<p>Enables or disables Backup Exec server authentication for Oracle operations.</p> <ul style="list-style-type: none"> ■ /auth:on - Turns the state on. Requires /user parameter. ■ /auth:off - Turns the state off. Requires /user parameter. ■ /user:<username> - Sets a username. ■ /password:<password> - Sets a password to be used with /user:<username>. If you enter an asterisk for the password or omit the password, you are prompted for the password.

Table A-1 Backup Exec Agent Utility Command Line Applet switches
(continued)

Switch	Description
/port:[<port>]	<p>Displays or sets a custom port that is used to connect to the Backup Exec server during Oracle operations.</p> <ul style="list-style-type: none">■ /port - Displays the current port number. If the port is the default port, displays "(default)".■ /port:<port> - Sets the port number to <port>. To change the port to the default port number, type [/port:0].
/log_path:[<log path>]	<p>Displays or sets a custom path for debug logs.</p> <ul style="list-style-type: none">■ /log_path - Displays the log directory path and then exits.■ /log_path:<"logs path"> - Creates the directory <"logs path">. If the path has a space in the name, enclose the path in quotes. For example, "C:\Program files\LogsFolder".

Symantec Backup Exec Deduplication Option

This appendix includes the following topics:

- [About the Deduplication Option](#)
- [Deduplication methods for Backup Exec agents](#)
- [Requirements for the Deduplication Option](#)
- [Installing the Deduplication Option](#)
- [Configuring an OpenStorage device](#)
- [Editing the properties of an OpenStorage device](#)
- [Specifying a Backup Exec server that has proximity to a shared OpenStorage device](#)
- [Creating or importing deduplication disk storage](#)
- [Selecting storage devices for direct access sharing](#)
- [Changing the location of a deduplication disk storage device](#)
- [Sharing a deduplication device between multiple Backup Exec servers](#)
- [How to use client-side deduplication](#)
- [How to set up backup jobs for deduplication](#)
- [Using optimized duplication to copy deduplicated data between OpenStorage devices or deduplication disk storage devices](#)
- [Copying deduplicated data to tapes](#)

- [Using deduplication with encryption](#)
- [Restoring a deduplication disk storage device or deduplicated data](#)
- [Disaster recovery of deduplication disk storage devices](#)
- [Disaster recovery of OpenStorage devices](#)

About the Deduplication Option

The Backup Exec Deduplication Option supports a data-reduction strategy by optimizing storage and network bandwidth. The Deduplication Option supports integrated deduplication at the Backup Exec server and on remote computers that have the Agent for Windows or the Agent for Linux installed. It also allows data to be deduplicated and stored on intelligent disk devices from Symantec and other vendors.

Table B-1 Types of deduplication

Type of deduplication	Where deduplication occurs	Benefits
Backup Exec server-side deduplication	On the Backup Exec server.	Reduces the size of backups, which reduces storage requirements.
Client-side deduplication	On the remote computer where the data is located. Note: The Agent for Windows is required on the remote Windows computer to perform Windows client-side deduplication. The Agent for Linux is required on the Linux computer to perform Linux client-side deduplication.	Reduces network traffic because only unique data is sent across the network. It also reduces the backup window.
Appliance deduplication	On an intelligent disk device, such as Symantec PureDisk, or a device from a third-party vendor.	Reduces the size of backups, which reduces storage requirements. It also reduces the backup window.

With a single Deduplication Option license key, you can use two types of deduplication devices.

Table B-2 Types of deduplication devices that work with the Deduplication Option

Type of device	Description
OpenStorage device	<p>Backup Exec uses Symantec's OpenStorage technology, which allows intelligent disk devices to integrate with Backup Exec. You can back up data to the Symantec PureDisk device and to storage devices from other vendors.</p> <p>You can find a list of compatible types of storage at the following URL:</p> <p>http://entsupport.symantec.com/umi/V-269-2</p> <p>See "Configuring an OpenStorage device" on page 804.</p>
Deduplication disk storage	<p>Deduplication disk storage provides integrated deduplication on the Backup Exec server. Deduplication disk storage is a disk-based backup folder that is located on the Backup Exec server.</p> <p>See "Creating or importing deduplication disk storage" on page 809.</p>

In addition to reducing storage requirements and network traffic, the Deduplication Option lets you do the following:

- Copy deduplicated data from an OpenStorage device or deduplication disk storage to tape for long-term or off-site storage.
- Use optimized duplication, which lets you copy deduplicated data between OpenStorage devices from the same vendor and between deduplication disk storage devices.
- Use Symantec's Granular Recovery Technology (GRT) with jobs that use deduplication devices.
- Share OpenStorage devices and deduplication storage devices among multiple Backup Exec servers when you use the Central Admin Server Option.

See "Installing the Deduplication Option" on page 803.

See "Requirements for the Deduplication Option" on page 798.

See "Sharing a deduplication device between multiple Backup Exec servers" on page 821.

See [“Using optimized duplication to copy deduplicated data between OpenStorage devices or deduplication disk storage devices”](#) on page 823.

See [“Copying deduplicated data to tapes”](#) on page 826.

Deduplication methods for Backup Exec agents

Backup Exec supports the following deduplication methods:

- Client-side deduplication, either on an intelligent disk device or to a deduplication disk storage device.
- Backup Exec server-side deduplication with a deduplication disk storage device.
- Appliance deduplication on an OpenStorage device.

The following table lists the deduplication methods that are available for the Backup Exec agents.

Table B-3 Deduplication methods for Backup Exec agents

Agent	Client-side deduplication (file system backups or VSS-snapshot enabled backups, whichever is supported)	Client-side deduplication (with Granular Recovery Technology enabled)	Backup Exec server-side deduplication (file system backups or VSS-snapshot enabled backups, whichever is supported)	Backup Exec server-side deduplication (with Granular Recovery Technology enabled)	Appliance deduplication on an OpenStorage device
Agent for Windows	Yes	Not applicable	Yes	Not applicable	Yes
Agent for VMware and Hyper-V	Yes (for Hyper-V only) Note: The Agent for Windows must be installed on the Hyper-V host.	Yes (for Hyper-V only) Note: The Agent for Windows must be installed on the Hyper-V host.	Yes	Yes	Yes
Agent for Linux	Yes	No	Yes	Not applicable	Yes

Table B-3 Deduplication methods for Backup Exec agents (*continued*)

Agent	Client-side deduplication (file system backups or VSS-snapshot enabled backups, whichever is supported)	Client-side deduplication (with Granular Recovery Technology enabled)	Backup Exec server-side deduplication (file system backups or VSS-snapshot enabled backups, whichever is supported)	Backup Exec server-side deduplication (with Granular Recovery Technology enabled)	Appliance deduplication on an OpenStorage device
Agent for Enterprise Vault	No	No	Yes	No	No
Exchange Agent	Yes	Yes	Yes	Yes	Yes
SQL Agent	Yes	Not applicable	Yes	Not applicable	Yes
SharePoint Agent	Yes	Yes	Yes	Yes	Yes
Active Directory Agent	Yes	Yes	Yes	Yes	Yes
Agent for Oracle	Linux: Yes Windows: Yes	No	Yes	No	Yes
Agent for Lotus Domino	Yes	No	Yes	No	Yes
Agent for Mac	No	Not applicable	Yes	Not applicable	Yes

See [“About the Deduplication Option”](#) on page 795.

Requirements for the Deduplication Option

The requirements for the Deduplication Option vary depending on the type of storage devices you want to use and the type of deduplication you want to use. Before you install the Deduplication Option, you should determine what type of storage devices

you want to use with it and what type of deduplication you want to use. Then, verify that your system meets the requirements for the storage devices you want to use.

Warning: Symantec strongly recommends that you exclude the deduplication disk storage device from all antivirus scans. If an antivirus scanner deletes or quarantines the files from the deduplication disk storage device, access to the device may be disabled.

Table B-4 Requirements and recommendations for the Deduplication Option

Item	Requirements and Recommendations
Deduplication disk storage devices	

Table B-4 Requirements and recommendations for the Deduplication Option
(continued)

Item	Requirements and Recommendations
	<p>The following items are required:</p> <ul style="list-style-type: none">■ A 64-bit Backup Exec server.■ A Backup Exec server with a minimum of 4 cores. Symantec recommends 8 cores.■ For 64 TB of stored deduplicated data, 8 cores are required.■ A dedicated volume to use as the location to store the deduplication disk storage. The dedicated volume must have at least 5 gigabytes (GB) of free space.■ The deduplication disk storage device must be excluded from antivirus scans. If an antivirus scanner deletes or quarantines the files from the deduplication disk storage device, access to the deduplication disk storage device may be disabled.■ 8 GB RAM for 5 terabytes (TB) of stored deduplicated data. For more than 5 TB of data up to 64 TB, use the following calculation to find the required amount of RAM: $1.5 \text{ GB} \times n$ where n is the number of TBs of deduplicated data to be stored. For example, to calculate the required amount of RAM for 64 TB of data: $1.5 \times 64 = 96 \text{ GB RAM}$■ Symantec recommends the following as minimum disk speeds per individual read, write, or verify operation:<ul style="list-style-type: none">■ Up to 32 TBs of storage:<ul style="list-style-type: none">■ 130 MB per second■ 200 MB per second for enterprise-level performance■ 32 to 48 TBs of storage: 200 MB per second■ 48 to 64 TBs of storage: 250 MB per second <p>The above recommendations are for the performance of a single operation. You may need more capability depending on your objectives for writing to and reading from disk.</p> <p>Note: You should be aware of the effects that computer disk speeds have on deduplication performance.</p> <p>Computer disk speeds can have the following effects on deduplication performance:</p> <ul style="list-style-type: none">■ Computers with disk speeds greater than 200 MB per second have optimal read and write performance for deduplication.■ Computers with disk speeds between 150-200 MB per second have sufficient read and write speed for deduplication.

Table B-4 Requirements and recommendations for the Deduplication Option
(continued)

Item	Requirements and Recommendations
	<ul style="list-style-type: none">■ Computers with disk speeds between 100-150 MB per second have some operations with degraded performance.■ Computers with disk speeds less than 100 MB per second experience poor performance. You should improve disk reads and writes before you install and run deduplication.
Deduplication disk storage connection	<ul style="list-style-type: none">■ Storage area network (Fibre Channel or iSCSI), direct-attached storage (DAS), or internal disks are supported. Removable disks including USB, eSATA, and FireWire devices are not supported.■ The Backup Exec server should have redundant connectivity to the storage.■ The storage network must be a dedicated, low latency network with a maximum of 1-millisecond latency per round trip.■ The storage network must have enough bandwidth to meet your throughput objectives. Symantec supports the following storage network bandwidths:<ul style="list-style-type: none">■ iSCSI SANs with a bandwidth of at least 10 Gb per second.■ Fibre Channel SANs with a bandwidth of at least 4 Gb per second.■ Symantec requires a minimum bandwidth of 130 MB per second for read and write performance. Bandwidth that is less than 130 MB per second may be used in smaller, less resource-intensive environments. However, as usage increases, deduplication requires more bandwidth to ensure adequate throughput for deduplication processes and backups. Otherwise, performance and stability are negatively affected.
Deduplication disk storage credentials	<p>The following requirements are for the password credentials for a deduplication disk storage device:</p> <ul style="list-style-type: none">■ The password cannot be blank■ The password cannot contain the following characters: &, ", <, >, %, ^■ The password cannot start with the hyphen character: -■ The password cannot end with the backslash character: \

Table B-4 Requirements and recommendations for the Deduplication Option
(continued)

Item	Requirements and Recommendations
OpenStorage devices	<p>To use a Symantec PureDisk device or a storage device from another vendor as an OpenStorage device, you must purchase the device and the appropriate OpenStorage connector from the device's vendor.</p> <p>You can use the Deduplication Option with OpenStorage devices on either a 32-bit Backup Exec server or a 64-bit Backup Exec server. The standard system requirements for Backup Exec apply to the Deduplication Option when you use OpenStorage devices.</p>
Client-side deduplication for Windows	<p>On the server where the Agent for Windows is installed, 1.5 GB of memory is required.</p> <p>Both 32-bit and 64-bit Windows operating systems are supported.</p>
Client-side deduplication for Linux	<p>You can find a list of compatible types of storage for Linux client-side deduplication in Backup Exec's Software Compatibility List: http://entsupport.symantec.com/umi/V-269-1</p> <p>A 64-bit Linux operating system is required.</p> <p>The following deduplication devices can be used:</p> <ul style="list-style-type: none">■ Deduplication disk storage device■ Symantec PureDisk OpenStorage device, which is the only type of OpenStorage device that supports client-side deduplication for Linux. <p>The following Backup Exec options are required:</p> <ul style="list-style-type: none">■ Agent for Linux■ Deduplication Option

See ["Installing the Deduplication Option"](#) on page 803.

Installing the Deduplication Option

The Deduplication Option is installed using the Backup Exec installation media. You install it locally as a separate, add-on component of Backup Exec. Before you attempt to install the Deduplication Option, verify that your system meets the requirements.

See ["Requirements for the Deduplication Option"](#) on page 798.

See ["Installing additional Backup Exec options to the local Backup Exec server"](#) on page 76.

Configuring an OpenStorage device

OpenStorage is a Symantec technology that allows intelligent disk devices to integrate with Backup Exec.

You can find a list of compatible types of storage at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

Some intelligent disk devices can include multiple logical storage units. However, each logical storage unit is added as a single OpenStorage device. When you add an OpenStorage device, Backup Exec can automatically locate the logical storage units on that device.

Note: When you delete or erase the media from an OpenStorage device, it may take up to 48 hours for more space to become available. Backup Exec cannot always calculate the amount of space that will be made available.

If you use Backup Exec Central Admin Server Option, you can share an OpenStorage device between multiple Backup Exec servers. You can enable sharing when you add an OpenStorage device. You can select new Backup Exec servers to share an OpenStorage device or remove the sharing ability for Backup Exec servers at any time. You can specify a different Backup Exec server on which to run backup operations when the OpenStorage device is the source of a job, such as a duplicate job, or a verify job.

See “[Sharing a deduplication device between multiple Backup Exec servers](#)” on page 821.

Before you configure an OpenStorage device, have the following information available:

- What you want to name the OpenStorage device.
- The name of the provider for the OpenStorage device.
- The name of the server and the logon account to use to access the OpenStorage device. You cannot use the System Logon Account. Symantec recommends that you select or create a logon account that you use exclusively for the deduplication disk storage device. You should not use this account for any other purpose. This account should not contain credentials that are subject to password update policies.
- Whether you want to enable encryption while data is transmitted to the OpenStorage device and while the data is stored on it.
- The number of concurrent operations to run on the device. This setting determines the number of jobs can run at the same time on this device. The

number of jobs varies depending on your hardware and environment, so you may need to adjust this setting more than once. Symantec recommends that you set it low enough to avoid overloading your system, but high enough to process your jobs in a timely manner.

To prevent a deduplication job from being sent to a non-deduplication device in a storage pool, you cannot add an OpenStorage device to any storage pools.

To configure an OpenStorage device

1 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.

2 Do one of the following:

If the Central Admin Server Option is not installed Click **Network storage**, and then click **Next**.

If the Central Admin Server Option is installed Do the following in the order listed:

- Select a Backup Exec server, and then click **Next**.
- Click **Network storage**, and then click **Next**.

3 Click **OpenStorage**, and then click **Next**.

4 Enter a name and description for the OpenStorage device, and then click **Next**.

5 Do one of the following:

If the provider for the OpenStorage device is listed Select the provider, and then click **Next**.

If the provider for the OpenStorage device is not listed Do the following in the order listed:

- Select **My provider is not listed here**, and then click **Next**.
- Enter the provider name, and then click **Next**.

6 Enter the name of the server and the logon account to use to access the OpenStorage device, and then click **Next**.

7 Specify the number of concurrent operations that can run on the OpenStorage device, and then click **Next**.

8 Review the summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To create or import the OpenStorage device

Click **Finish**.

See [“Specifying a Backup Exec server that has proximity to a shared OpenStorage device”](#) on page 808.

Editing the properties of an OpenStorage device

You can view all of the properties of an OpenStorage device and you can change some of the properties.

To edit the properties for an OpenStorage device

- 1 On the **Storage** tab, double-click the name of the OpenStorage device.
- 2 In the left pane, select **Properties**.

3 Change the following properties as needed:

Name	Indicates the user-defined name for this OpenStorage device.
Description	Indicates the user-defined description of this OpenStorage device.
Backup Exec server with device proximity	<p>Indicates the server that you want to run backup operations when the OpenStorage device is shared between servers in a Central Admin Server Option environment. You can specify a Backup Exec server that has physical or logical proximity to the OpenStorage device. Proximity of the Backup Exec server to the device is an advantage when the device is the source of a job, such as a duplicate job or a verify job, and you want to avoid running the job over a WAN.</p> <p>By default, the Backup Exec server on which you create the OpenStorage device is the server that runs the backup operations for the device.</p>
Logon account	Indicates the name of the logon account that is required to access the device.
Concurrent operations	Indicates the maximum number of jobs that you want to run at the same time on this device.
Split data stream every	Indicates the size at which you want Backup Exec to span to a new image. The default size is 50 GB.
Data stream size	Indicates the size of a single write operation that Backup Exec issues. The default size varies based on the type of device that is being used.
Stream handler	Indicates whether stream handler is used. Backup Exec sets this option automatically when you select a server type. For some types of devices, this option does not appear at all. If Backup Exec does not set this option, contact the device's vendor for the recommended setting.

Client-side deduplication

Indicates whether client-side deduplication is enabled for this OpenStorage device.

Client-side deduplication enables a remote computer to send data directly to an OpenStorage device. By using client-side deduplication, the Backup Exec server is bypassed, which leaves the Backup Exec server free to perform other operations.

Disk space to reserve for non-Backup Exec operations

Displays the amount of disk space to set aside for applications other than Backup Exec. The default amount is 5%.

- 4 Click **Apply** to save the changes.

Specifying a Backup Exec server that has proximity to a shared OpenStorage device

You can specify a Backup Exec server that has physical or logical proximity to an OpenStorage device to run operations for that device. Proximity of the Backup Exec server to the device is an advantage when the device is the source of a job, such as a duplicate job or a verify job, and you want to avoid running the job over a WAN.

By default, the Backup Exec server on which you create the device is the server that runs the backup operations for the device. You can specify a Backup Exec server with proximity only if the device is shared between servers in a Central Admin Server Option environment.

If you remove the specified Backup Exec server with device proximity, you are prompted to specify another server.

To specify a Backup Exec server that has proximity to a shared OpenStorage device

- 1 On the **Storage** tab, double-click the name of the device.
- 2 In the left pane, select **Properties**.
- 3 In the **Backup Exec server with device proximity** field, on the drop-down menu, select the server that you want to run operations for the device.
- 4 Click **Apply**.

Creating or importing deduplication disk storage

Deduplication disk storage provides a disk-based backup folder that you can use as a destination for backup jobs. When you use deduplication disk storage, only unique data is stored.

Before you create a deduplication disk storage device, review the requirements. Symantec recommends a dedicated volume and a large amount of RAM for deduplication disk storage.

See [“Requirements for the Deduplication Option”](#) on page 798.

You can create only one deduplication disk storage device on a Backup Exec server. You can create deduplication disk storage on a storage array. However, if a deduplication disk storage device already exists on a Backup Exec server, then you cannot add another device to a storage array that is connected to that Backup Exec server.

When you use Backup Exec's **Delete** option on a deduplication disk storage device, the folder is removed from the Backup Exec database. However, the folder and the files in it remain on the disk. When you delete backup sets from a deduplication disk storage device, it may take up to 48 hours for more space to become available. Backup Exec cannot always calculate the amount of space that will be made available.

If you use the Backup Exec Central Admin Server Option, you can share a deduplication disk storage device between multiple Backup Exec servers. You can enable sharing when you add a deduplication disk storage device. You can select new Backup Exec servers to share deduplication disk storage or remove the sharing ability for Backup Exec servers at any time.

You can create new deduplication disk storage, or you can import an existing deduplication disk storage device from another Backup Exec server.

See [“Sharing a deduplication device between multiple Backup Exec servers”](#) on page 821.

Before you create or import a deduplication disk storage device, have the following information available:

- What you want to name the deduplication disk storage device.
- The volume on which you want to create the deduplication disk storage device.
 - The path of the existing deduplication disk storage device, if you import a deduplication disk storage device.
 - The information for the user account that was used when the existing deduplication disk storage device was originally created.

- The logon account to use to access the deduplication disk storage device. You cannot use the System Logon Account. Symantec recommends that you select or create a logon account that you use exclusively for the deduplication disk storage device. You should not use this account for any other purpose. This account should not contain credentials that are subject to password update policies.

The following requirements are for the password credentials for a deduplication disk storage device:

- The password cannot be blank
- The password cannot contain the following characters: &, ", <, >, %, ^
- The password cannot start with the hyphen character: -
- The password cannot end with the backslash character: \
- Whether you want to enable encryption while the data is transmitted to the deduplication disk storage device and while the data is stored on it. You should not use the Backup Exec encryption options for backup jobs that deduplicate data.
- The number of concurrent operations to run on the device. This setting determines the number of jobs that can run at the same time on this device. The number of jobs varies depending on your hardware and environment, so you may need to adjust this setting more than once. Symantec recommends that you set it low enough to avoid overloading your system, but high enough to process your jobs in a timely manner.

To create or import deduplication disk storage

- 1 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.
- 2 Do one of the following:

If the Central Admin Server Option is not installed

Select **Disk-based storage**, and then click **Next**.

If the Central Admin Server Option is installed

Do the following in the order listed:

- Select the Backup Exec server that you want to configure deduplication disk storage for, and then click **Next**.
- Select **Disk-based storage**, and then click **Next**.

- 3 Click **Deduplication disk storage**, and then click **Next**.

- 4 Enter a name and description for the deduplication disk storage device, and then click **Next**.
- 5 Do one of the following:

To create a new deduplication disk storage device	Click Next to accept the default location that appears in the field.
To import an existing deduplication disk storage device	Enter the path of the existing deduplication disk storage device, and then click Next .
- 6 Click the drop-down arrow and select the logon account to use to access the deduplication disk storage device, or click **Add/Edit** to create a logon account, and then click **Next**.
- 7 Specify if you want to enable encryption during transmission of data to the deduplication disk storage device and while the data is stored on it, and then click **Next**.
- 8 Specify the number of concurrent operations that can run on the deduplication disk storage device, and then click **Next**.
- 9 Review the storage configuration summary, and do either of the following:

To change any of the selections	Do the following in the order listed: <ul style="list-style-type: none">■ Click a review heading and make the appropriate changes.■ Click Next through the wizard to return to the summary screen.■ Click Finish.
To start the configuration	Click Finish .

Editing the properties of a deduplication disk storage device

You can edit some properties of a deduplication disk storage device.

To edit the properties of a deduplication disk storage device

- 1 On the **Storage** tab, double-click the name of the deduplication disk storage device.
- 2 In the left pane, select **Properties**.

3 Change the following properties as needed:

Name	Indicates the name that was entered when the deduplication disk storage was configured. You can change the name at any time.
Description	Indicates the description that was entered when the deduplication disk storage was configured. You can change the description at any time.
Logon account	<p>Indicates the logon account that is being used to access the device.</p> <p>If you change the logon account after you configure the device, additional steps are required to enable access to the data that is already stored on the device.</p> <p>See "Changing the password for the logon account for deduplication disk storage" on page 816.</p>
Encryption	<p>Enables or disables encryption while the data is transmitted to the device, and while the data is stored on the device.</p> <p>You should not use the Backup Exec encryption options for backup jobs that deduplicate data.</p>
Concurrent operations	<p>Indicates the maximum number of jobs that you want to run at the same time on this device.</p> <p>The number of jobs varies depending on your hardware and environment, so you may need to adjust this setting more than once. Symantec recommends that you set it low enough to avoid overloading your system, but high enough to process your jobs in a timely manner.</p>
Data stream size	Indicates the size of a single write operation that Backup Exec issues. The default size varies based on the type of device being used.

Client-side deduplication	<p>Indicates whether client-side deduplication is enabled for this device.</p> <p>Client-side deduplication enables a remote computer that is configured to send data directly to the deduplication disk storage. After the data is deduplicated, then only unique data is sent directly to the deduplication disk storage. By using this option, the Backup Exec server is bypassed, which leaves the Backup Exec server free to perform other operations.</p>
Percentage of disk space to reserve for non-Backup Exec operations	<p>Displays the amount of disk space to set aside for applications other than Backup Exec. The default amount is 5%</p>
Log level	<p>Indicates the type of information you want to include in the diagnostic logs for this device. The choices range from critical errors only to all types of messages.</p>
Log retention period	<p>Indicates the number of days to keep the diagnostic logs for this device.</p>

Low disk space - Critical

Displays the critically low disk space threshold at which you want Backup Exec to send an alert. Backup Exec sends alerts when the amount of free disk space drops below the low disk space threshold, and again if it drops below the warning threshold. The amount of free disk space does not include the disk space that is reserved for non-Backup Exec operations.

You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes. This threshold must be less than the warning low disk space threshold.

You may want to set the threshold slightly higher than the minimum amount that you need to run jobs. By doing so, you allow time to address the disk space issue before the jobs fail.

The default is 5%.

This property appears only if the deduplication disk storage is on a storage array.

Low disk space - Warning

Displays the low disk space threshold at which you want Backup Exec to send an alert. If free disk space drops below the warning threshold to the critical threshold, another alert is sent. The amount of free disk space does not include the disk space that is reserved for non-Backup Exec operations.

You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes. This threshold must be less than the low disk space threshold.

You may want to set the threshold slightly higher than the minimum amount that you need to run jobs. By doing so, you allow time to address the disk space issue before the jobs fail.

The default is 15%.

This property appears only if the deduplication disk storage is on a storage array.

Low disk space

Displays the low disk space threshold at which you want Backup Exec to send an alert. If free disk space drops below the warning threshold to the critical threshold, another alert is sent. The amount of free disk space does not include the disk space that is reserved for non-Backup Exec operations.

You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes.

You may want to set the threshold slightly higher than the minimum amount that you need to run jobs. By doing so, you allow time to address the disk space issue before the jobs fail.

The default is 25%.

This property appears only if the deduplication disk storage is on a storage array.

- 4 Click **Apply** to save the changes

Changing the password for the logon account for deduplication disk storage

When you specify a Backup Exec logon account for a deduplication disk storage device, an additional user account is created for the deduplication components with the same user name and password. However, if you change the credentials for the Backup Exec logon account, the credentials for the additional user account are not changed automatically. You must use the `spausr.exe` utility to update the password for the additional user account. This account is known as the "User 1" account when you use the `spausr.exe` utility to view a list of user names that are associated with the deduplication disk storage.

To change the password for the logon account for deduplication disk storage

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Logon Accounts**, and then select **Manage Logon Accounts**.
- 3 Select the Backup Exec logon account that you want to change, and then click **Edit**.
- 4 Type the current password for the logon account, and then click **OK**.

- 5 Click **Change Password**.
- 6 Type the new password in the **Password** field and in the **Confirm** field.
- 7 Click **OK**.
- 8 At a command prompt, switch to the Backup Exec program file directory, and then type the following command:

```
spausers.exe -c -u <UserName>
```

The default Backup Exec program file directory is C:\Program Files\Symantec\Backup Exec

The user name is case-sensitive. If you do not know the user name, type the following command to find the user name that is associated with "User 1":

```
spausers.exe -l
```

You are prompted for the old password and a new password. Ensure that the new password is the same as the password that you used in step 6.

Selecting storage devices for direct access sharing

Direct access enables a remote computer to send data directly to storage devices that are hosted by a Backup Exec server. When direct access sharing is enabled, the Backup Exec server is bypassed.

If you use a deduplication disk storage device or an OpenStorage device that supports client-side deduplication, then enabling direct access sharing enables Backup Exec to perform client-side deduplication. Note that client-side deduplication is CPU-intensive.

Direct access sharing becomes available after you create a backup job in which a deduplication device is selected and the following option is selected: **Enable the remote computer to directly access the storage device and to perform client-side deduplication, if it is supported**.

After the correctly configured backup job is created, then the option **Direct access sharing** appears in the following locations.

- On the details screen for a server on the **Backup and Restore** tab.
- On the details screen for a storage device on the **Storage** tab.

In addition, the option **Direct access properties** appears on the details screen for a server on the **Backup and Restore** tab.

To select storage devices for direct access sharing

- 1 Do one of the following:

- On the **Backup and Restore** tab, double-click the server that you want to set up to share devices.
 - On the **Storage** tab, double-click the storage device that you want to share.
- 2 In the left pane, select **Direct access sharing**.
 - 3 Select the check box for the items that you want to share.

Note: After you enable direct access sharing for a client, you must restart the Backup Exec services on the Backup Exec server. Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Services**, and then click **Restart all services**.

See [“Editing server properties for direct access”](#) on page 818.

Editing server properties for direct access

For servers that are enabled for direct access, you can do the following:

- Add or change a description of the server.
- Enable or disable ICMP ping operations to detect the server.
- Add or edit a logon account that is used to access the remote computer.

To edit server properties for direct access

- 1 On the **Backup and Restore** tab, double-click the server that is enabled for direct access.
- 2 In the left pane, select **Direct access properties**.

- 3 Edit the following options as needed:

Server name	Indicates the name of the remote computer or managed Backup Exec server.
Description	Lets you enter a description of the server.
Port	Indicates the port that is used for communications between the Backup Exec server and the remote computer.
Use ICMP ping operations to detect the server	Lets the Backup Exec server use ICMP ping to locate the remote computer.
Logon account	Indicates the logon account that is required to access the remote computer. You can add a new logon account or edit an existing account.

- 4 To save changes, click **Apply**.

Changing the location of a deduplication disk storage device

You can change the location of an existing deduplication disk storage device. You can move the deduplication disk storage device to another volume on the same Backup Exec server. If the Central Admin Server Option is installed, you can move the deduplication disk storage device to another Backup Exec server. Only one deduplication disk storage device is supported per Backup Exec server.

Symantec recommends that you use the same name for the deduplication disk storage device when you change the location, but it is not required.

To change the location of a deduplication disk storage device

- 1 Ensure that no jobs are running or are scheduled to run until you have completed the process of changing the location of a deduplication disk storage device.
- 2 Document the current path of the deduplication disk storage device by viewing the properties of the device.
- 3 On the **Storage** tab, right-click the deduplication disk storage device, and then click **Disable**.
- 4 Right-click the deduplication disk storage device again, and then click **Delete**.
- 5 When you are prompted to delete the device, click **Yes**.

- 6 Click the Backup Exec button, select **Configuration and Settings**, select **Backup Exec Services**, and then click **Stop all services**.
- 7 In Windows Explorer, copy the deduplication disk storage device to the new path or volume.

Note: This step can take a long time as all of the data within the device is copied to the new location.

- 8 On the Backup Exec Administration Console, click the Backup Exec button, select **Configuration and Settings**, select **Backup Exec Services**, and then click **Restart all services**.
- 9 On the **Storage** tab, click **Configure Storage**, click **Disk-based storage**, and then click **Next**.
- 10 Click **Deduplication disk storage**, and then click **Next**.
- 11 Enter a name and description for the deduplication disk storage device, and then click **Next**.
- 12 Do one of the following:

If the deduplication disk storage device was created in Backup Exec 2012

Click **Create a new deduplication disk storage device**.

If the deduplication disk storage device was created in a version prior to Backup Exec 2012

Click **Import an existing deduplication disk storage device**, and then enter the path to which you moved the deduplication disk storage device.

- 13 Click **Next**.
- 14 Specify the logon account that was used for the original deduplication disk storage device, and then click **Next**.
- 15 Specify if you want to enable encryption, and then click **Next**.
- 16 Specify the number of concurrent operations to run on the deduplication disk storage device, and then click **Next**.
- 17 Review the summary, and then click **Finish**.
- 18 When the deduplication disk storage device comes online, you can delete the original files.

Sharing a deduplication device between multiple Backup Exec servers

If you use the Backup Exec Central Admin Server Option, you can select which Backup Exec servers can share a deduplication disk storage device or an OpenStorage device. When you add a deduplication disk storage device or an OpenStorage device, the Backup Exec server that you used to add the device is automatically selected for sharing.

Note: To share a deduplication disk storage device, you must add it as an OpenStorage device on all Backup Exec servers that you want to access the device, except for the Backup Exec server that was used to create it.

This type of sharing is not the same as direct access sharing. With direct access sharing, a remote computer bypasses the Backup Exec server to directly access storage devices that are hosted by the Backup Exec server.

See [“Sharing storage devices”](#) on page 448.

See [“Selecting storage devices for direct access sharing”](#) on page 817.

How to use client-side deduplication

Client-side deduplication enables a remote computer to send data directly to an OpenStorage device or a deduplication disk storage device. By using client-side deduplication, the Backup Exec server is bypassed, which leaves the Backup Exec server free to perform other operations. If your deduplication device supports client-side deduplication, a remote computer deduplicates data and then sends only the unique data directly to a deduplication disk storage device or an OpenStorage device. Client-side deduplication is available for Windows computers and Linux computers.

Note: Client-side deduplication may increase the CPU utilization on the remote computer if your deduplication device supports client-side deduplication.

When you create a backup job with client-side deduplication, keep in mind the following items:

- The backup job can include resources from only one remote computer.
- The Agent for Windows is required on the remote Windows computer to perform Windows client-side deduplication. The Agent for Linux is required on the Linux computer to perform Linux client-side deduplication.

Note: A maximum of 64 remote agents with client-side deduplication enabled are allowed per Backup Exec server.

- The remote computer must be pingable.
- The remote computer cannot be a Backup Exec server.
- A deduplication disk storage device or an OpenStorage device must be used for the backup job.
- The option **Client-side deduplication** must be enabled on the properties for the storage device.
- The option **Enable the remote computer to directly access the storage device and to perform client-side deduplication, if it is supported** must be selected in the **Storage** options for the backup job. This option is selected by default when you select a deduplication disk storage device or an OpenStorage device as the storage for a backup job.

If you do not configure the remote computer to use client-side deduplication, then the data from the remote computer is sent to the Backup Exec server to be deduplicated. Then, the deduplicated data is backed up to the deduplication disk storage or the OpenStorage device. This process increases the CPU utilization on the Backup Exec server. However, this process is useful if you are backing up older remote computers.

See [“About the Deduplication Option”](#) on page 795.

See [“How to set up backup jobs for deduplication”](#) on page 822.

See [“Editing the properties of an OpenStorage device”](#) on page 806.

See [“Editing the properties of a deduplication disk storage device”](#) on page 811.

How to set up backup jobs for deduplication

Set up a backup job for deduplication by selecting the option **Back Up to Deduplication Disk Storage**. Then, on the **Storage** settings, select either an OpenStorage device or a deduplication disk storage device as the destination device, and then select the deduplication method to use.

The following deduplication methods are available:

- If you want to enable client-side deduplication, select the option **Enable the remote computer to directly access the storage device and to perform client-side deduplication, if it is supported**. This is the default option. If the storage device that you select for the job does not support client-side

deduplication, then either Backup Exec server-side deduplication or appliance deduplication is used.

- If you want to enable Backup Exec server-side deduplication, select the option **Enable the remote computer to access the storage device through the Backup Exec server and to perform Backup Exec server-side deduplication, if it is supported**. If the storage device that you select for the job does not support server-side deduplication, then appliance deduplication is used.

See [“How to use client-side deduplication”](#) on page 821.

Using optimized duplication to copy deduplicated data between OpenStorage devices or deduplication disk storage devices

Backup Exec supports optimized duplication, which enables deduplicated data to be copied directly from one OpenStorage device to another OpenStorage device from the same vendor. Both devices must be attached to a single Backup Exec server. For example, you can copy data from one Symantec PureDisk device to another Symantec PureDisk device. Because the data is deduplicated, only unique data is copied between the devices.

To copy data between OpenStorage devices or deduplication disk storage devices, you must create a job to duplicate backup sets. The destination device for the duplicate job must be the same type of device from the same vendor as the device that was used in the source backup job. No additional settings are required; optimized duplication occurs automatically when you set up a duplicate backup job between appropriate devices. You can restore data from either device.

Optimized duplication can be performed on backup sets that were enabled for Granular Recovery Technology (GRT). However, only deduplication disk storage devices and PureDisk devices support optimized duplication for GRT-enabled backup sets.

You can find a list of compatible types of storage devices in Backup Exec's Hardware Compatibility List:

<http://entsupport.symantec.com/umi/V-269-2>

Note: The OpenStorage devices must be from the same vendor. You cannot perform optimized duplication between OpenStorage devices from different vendors. If you attempt to copy deduplicated data between OpenStorage devices from different vendors, regular duplication is performed instead of optimized duplication.

Using optimized duplication with the Central Admin Server Option

If you use the Central Admin Server Option (CASO), the functionality of optimized duplication is expanded to let you do the following:

- Copy data from a deduplication disk storage device on one Backup Exec server to a deduplication disk storage device on another Backup Exec server.
- Copy data from an OpenStorage device that is attached to a Backup Exec server to another OpenStorage device that is attached to a different Backup Exec server.

To use optimized duplication with CASO, the following requirements must be met:

- You must have a license for the Enterprise Server Option. CASO is installed as part of the Enterprise Server Option.
- All Backup Exec servers that you use with CASO as either a central administration server or as managed Backup Exec servers must use the 64-bit version of Windows.
- You must have a central administration server and at least one managed Backup Exec server in your CASO environment.
- For client-side deduplication and Backup Exec server-side deduplication, you must configure one deduplication disk storage on the Backup Exec server from which you want to copy the deduplicated data. You must also configure one deduplication disk storage on the Backup Exec server to which you want to copy the deduplicated data.
- For appliance deduplication, the Backup Exec server from which you want to copy deduplicated data must have the appropriate plug-in for the OpenStorage device and a properly configured OpenStorage device. In addition, the Backup Exec server to which you want to copy the deduplication data must have the appropriate plug-in for the OpenStorage device and a properly configured OpenStorage device.
- You must share deduplication devices between the Backup Exec servers.
- You must inventory and catalog the media on the destination server before you recover any files from the duplicated backup set. You must do this regardless of how the catalog sharing option is configured for CASO.

Using optimized duplication to transfer backup data to a remote location

You can use optimized duplication to transfer backup data to a remote location over a WAN connection. You should prepopulate the destination deduplication disk storage device with a full backup of the servers. This prepopulation task is also

known as seeding the destination deduplication disk storage device. Seeding helps avoid the time-consuming and bandwidth-intensive process of transmitting large amounts of backup data over the low bandwidth WAN connection.

To seed a device, you can store a full backup to an external storage device such as a USB drive. You can then transport the USB drive to the remote location where the Backup Exec server and the destination deduplication disk storage device are kept, and duplicate the backup data to the deduplication disk storage device.

How to set up optimized duplication

Set up a duplicate backup job to perform optimized duplication.

Table B-5 How to set up optimized duplication

Step	For more information
<p>If you are using CASO, do the following:</p> <ul style="list-style-type: none"> ■ Verify that you have one central administration server and at least one managed Backup Exec server. ■ Verify that the Backup Exec server from which you want to copy the deduplicated data has a deduplication disk storage device (for client-side or Backup Exec server-side deduplication) or an OpenStorage device (for appliance deduplication). Also verify that the Backup Exec server to which you want to copy the deduplicated data has a deduplication disk storage device (for client or Backup Exec server-side deduplication) or an OpenStorage device (for appliance deduplication). ■ Verify that the Backup Exec servers are enabled for sharing. <p>Note: This information applies only to CASO. If you do not have CASO, skip this step.</p>	<p>See “Configuring an OpenStorage device” on page 804.</p> <p>See “Creating or importing deduplication disk storage” on page 809.</p>
Create a backup job that uses an OpenStorage device or a deduplication disk storage device as the destination.	See “Backing up data” on page 158.

Table B-5 How to set up optimized duplication (*continued*)

Step	For more information
Create a job to duplicate backup sets and select the appropriate OpenStorage device or deduplication disk storage as the destination. Note: The destination device for the duplicate job must be the same type of device from the same vendor as the device that was used in the source backup job.	See “Duplicating backup sets or a job history manually” on page 221.

See [“Sharing a deduplication device between multiple Backup Exec servers”](#) on page 821.

Copying deduplicated data to tapes

Backup Exec lets you copy deduplicated data from an OpenStorage device to tape for long-term or off-site storage. When data is copied to tape, it is rehydrated. In other words, the files are reassembled into their original form and are not deduplicated.

To copy deduplicated data to tapes, you must create a duplicate backup job that copies the backup sets from the OpenStorage device to a tape device.

See [“Duplicating backup sets or a job history manually”](#) on page 221.

Using deduplication with encryption

You should not use the Backup Exec encryption options for backup jobs that deduplicate data. Data cannot be deduplicated when the Backup Exec encryption options are used.

If you want deduplicated data to be encrypted on a deduplication disk storage device, you can enable the encryption property on the deduplication disk storage device.

See [“Encryption key management”](#) on page 586.

Restoring a deduplication disk storage device or deduplicated data

You can restore a deduplication disk storage by running the Restore Wizard. When you restore a deduplication disk storage device, the original folder is deleted and then replaced by the restored folder.

Note: You cannot redirect the restore of a deduplication disk storage device. You must restore a deduplication disk storage device to its original location.

To restore deduplicated data, you can create a regular restore job. No additional settings are required.

See [“Restoring data from a server, a backup set, a backup job, or a storage device”](#) on page 234.

To restore a deduplication disk storage device

- 1 On the **Backup and Restore** tab, right-click the computer for which you want to restore the deduplication disk storage device, and then click **Restore**.
- 2 Select **Shadow Copy Components** and then click **Next**.
- 3 Follow the **Restore Wizard** prompts to restore the data.

See [“Disaster recovery of deduplication disk storage devices ”](#) on page 827.

Disaster recovery of deduplication disk storage devices

A deduplication disk storage device is stored on the Backup Exec server. If your Backup Exec server experiences a disaster, then the data from the deduplication disk storage device is lost. Therefore, you should take steps to prepare for recovery from a system failure. To prepare for a disaster, Backup Exec lets you take a snapshot of a deduplication disk storage device. The snapshot includes the folder and the contents of the folder. You can store the snapshot on tape, which you can then use to recover your deduplication disk storage after you recover the Backup Exec server.

When you restore data from the snapshot, the following processes occur:

- Backup Exec stops the deduplication services if they are running. The deduplication services are separate from the Backup Exec services, so the Backup Exec services are not affected.
- Backup Exec deletes any files that are present in the deduplication disk storage.

- The deduplication disk storage is restored to its original location, along with the contents of the folder.
- The deduplication services are restarted.

Note: If you use Backup Exec Simplified Disaster Recovery (SDR) to recover the Backup Exec server, SDR does not recover the deduplication disk storage during the recovery of the Backup Exec server.

Preparing for disaster recovery of a deduplication disk storage device

To prepare for a disaster, Backup Exec lets you take a snapshot of a deduplication disk storage device. The snapshot includes the folder and the contents of the folder. You can store the snapshot on tape, which you can then use to recover your deduplication disk storage after a disaster.

To prepare for disaster recovery of a deduplication disk storage device

- 1 On the **Backup and Restore** tab, right-click the server where the deduplication disk storage device is located.
- 2 Select **Backup**, and then select **Backup to Tape**.
- 3 In the **Selections** box, click **Edit**.
- 4 Expand **Shadow Copy Components**, expand **User Data**, and then select **Backup Exec Deduplication Storage**.
- 5 Click **OK**.
- 6 Complete any additional options that you want to use.

Symantec recommends that you schedule this job to run just prior to the 12:20 a.m. and 12:20 p.m. deduplication maintenance times.
- 7 Click **OK** to create the job.

See [“Restoring a deduplication disk storage device or deduplicated data”](#) on page 827.

Disaster recovery of OpenStorage devices

The following disaster recovery scenarios are possible for OpenStorage devices:

- The device fails.
- The Backup Exec server that uses the device fails.

If the device fails, you should consult the documentation from the device's vendor. If the Backup Exec server fails and you need to reinstall Backup Exec on the Backup Exec server, you must reconfigure the device, and inventory and catalog the media from it after the Backup Exec server is recovered.

Symantec Backup Exec Agent for VMware

This appendix includes the following topics:

- [About the Agent for VMware](#)
- [Requirements for using the Agent for VMware](#)
- [Using the Agent for VMware with Windows Server 2012](#)
- [Using the Agent for VMware with vSphere 5.5](#)
- [About installing the Agent for VMware](#)
- [Adding VMware vCenter servers and ESX/ESXi hosts to the list of servers on the Backup and Restore tab](#)
- [Viewing details about VMware resources](#)
- [Installing the Agent for Windows on VMware virtual machines](#)
- [Push-installing the Agent for Windows to VMware virtual machines](#)
- [Backing up VMware data](#)
- [Restoring VMware virtual machines and vmdk files](#)

About the Agent for VMware

The Symantec Backup Exec Agent for VMware (Agent for VMware) lets you back up and restore virtual machines that use VMware ESX/ESXi/vSphere/vCenter 4.0, 4.1, 5.0, 5.1, and 5.5.

Backup Exec performs a single-pass backup to protect all guest virtual machines and the VSS-aware applications that are installed on the guest virtual machines. Backup Exec's Granular Recovery Technology (GRT) is enabled by default for files and folders. You can use a GRT-enabled backup to restore individual files and folders from a Windows guest virtual machine without restoring the entire virtual machine. In addition, you can restore individual items from Microsoft Exchange, SQL, SharePoint, and Active Directory applications that reside on guest virtual machines. Note that application-level GRT is not enabled by default. You must select the application-level GRT options that you want to use when you set up a backup job.

Additional features of the Agent for VMware let you do the following:

- Redirect the restore of data from a guest virtual machine to an alternate folder, datastore, host, or network.
- Back up to a disk device or to a tape device.
- Perform incremental and differential backup jobs. This feature is available only if your virtual machines are configured with hardware version 7 or later.

See [“Requirements for using the Agent for VMware”](#) on page 831.

See [“Using Granular Recovery Technology \(GRT\) with the Agent for VMware”](#) on page 847.

See [“Backing up VMware data”](#) on page 838.

See [“Restoring VMware virtual machines and vmdk files”](#) on page 852.

Requirements for using the Agent for VMware

Before you use the Agent for VMware, ensure that the following requirements are met:

- Install the license for the Agent for Hyper-V and VMware on the Backup Exec server. The Backup Exec server runs the backup and restore job.
- Determine if you are going to use Backup Exec's Granular Recovery Technology to recover individual items from Microsoft applications. If you plan to use GRT, install the Agent for Windows on your virtual machines that run Windows.
See [“Using Granular Recovery Technology \(GRT\) with the Agent for VMware”](#) on page 847.
See [“Installing the Agent for Windows on VMware virtual machines”](#) on page 836.
- Verify that your environment includes supported hardware and software by reviewing the Hardware Compatibility List and the Software Compatibility List. You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

- If you have virtual machines in an ESX 4.0 environment, Symantec recommends that you install the latest ESX 4.0 patch from VMware to avoid potential issues with quiesced snapshots.
- Verify that your virtual machine names do not contain any characters that VMware does not support. Only the following characters can be used in virtual machine names:
 - Uppercase and lowercase ASCII characters
 - Numbers
 - Period (.)
 - Hyphen (-)
 - Underscore (_)
 - Plus sign (+)
 - Left and right parentheses ()
 - SpacesUse of any characters that are not listed above may cause backup jobs to fail.

Using the Agent for VMware with Windows Server 2012

The following information applies to virtual machines that use Windows Server 2012 and VMware ESX/ESXi/vSphere/vCenter 5.0 and later:

- GPT disks are not supported.
- If the Backup Exec server is running an operating system prior to Windows 2012, Granular Recovery Technology (both application-level and file/folder level GRT) cannot be performed for a virtual machine that has a Windows 2012 Resilient File System (ReFS) volume. A backup job for a virtual machine that is configured with Windows 2012 ReFS and is enabled for GRT will complete with a status of success with exceptions.
- If the Backup Exec server is running an operating system prior to Windows 2012, Granular Recovery Technology (both application-level and file/folder level GRT) cannot be performed for the virtual machine that contains volumes on which Windows 2012 deduplication is enabled. A backup job for a virtual machine

that is configured with a Windows 2012 deduplication volume and is enabled for GRT will complete with a status of success with exceptions.

See [“About the Agent for VMware”](#) on page 830.

Using the Agent for VMware with vSphere 5.5

Backup Exec 2014 supports VMware vSphere 5.5. Backing up, restoring, and creating virtual machines with Hardware Version 10 is supported. However, the following features of vSphere 5.5 are not supported for backup or restore:

Table C-1 vSphere 5.5 support

vSphere 5.5 feature	How Backup Exec handles it
In-guest SATA drives	When Backup Exec encounters a virtual machine that has a SATA disk defined, the job fails and an error is listed in the job log.
Linked clones	Backup Exec treats linked clones as full virtual machines. When a linked clone is restored, it does not retain any links to the original parent. Note: This information also applies to vSphere 5.0 and 5.1.
VMDKs larger than 2 TB	Backup Exec does not support Granular Recovery Technology (GRT) for VMDKs that are larger than 2 TB. Jobs that are enabled for GRT fail and an error is listed in the job log. Non-GRT backups and restores of VMDKs are supported.
GPT disks	Backup Exec does not support Granular Recovery Technology (GRT) for GPT disks. Jobs that are enabled for GRT fail and an error is listed in the job log. Non-GRT backups and restores of GPT disks are supported.

About installing the Agent for VMware

The Agent for VMware is installed as part of the Agent for Hyper-V and VMware. Enter the license for the Agent for Hyper-V and VMware on the Backup Exec server

to enable the Agent for VMware. You do not have to install the agent on the host server.

See [“Installing additional Backup Exec options to the local Backup Exec server”](#) on page 76.

Adding VMware vCenter servers and ESX/ESXi hosts to the list of servers on the Backup and Restore tab

You can add VMware vCenter servers and ESX/ESXi hosts to the list of servers on the **Backup and Restore** tab so that these items can be selected for backup jobs.

To add a VMware vCenter server or ESX/ESXi host to the list of servers on the Backup and Restore tab

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Select **VMware vCenter server or ESX host**, and then click **Next**.
- 3 Select **Add a VMware vCenter server or ESX host to the list of servers**.
- 4 Enter the name or IP address of the server you want to add, and then add an optional description of the server.
- 5 If you want Backup Exec to install the Agent for Windows on any guest virtual machines for this host server, verify that the option **After adding the virtual host, install the Backup Exec Agent for Windows on the guest virtual machines** is selected.

To perform Granular Recovery Technology (GRT)-enabled backups of the Microsoft applications that are installed on the virtual machines, the Agent for Windows must be installed on the guest virtual machines. If you do not want to install the Agent for Windows on the guest virtual machines, clear the check box. You can install the Agent for Windows later if you decide that you want to use GRT.

- 6 Click **Next**.
- 7 Select the correct logon account for the server from the drop-down list.
If the logon account for the server is not in the list, click **Add/Edit** to add it to the list.
- 8 Click **Next**.
- 9 Review the summary information, and then click **Finish**.

See [“About the list of servers on the Backup and Restore tab”](#) on page 152.

See [“Installing the Agent for Windows on VMware virtual machines”](#) on page 836.

Viewing details about VMware resources

The **Backup and Restore** tab includes a **Details** pane on the lower portion of the screen. The **Details** pane provides additional details for any type of server that is selected from the list of servers in the upper portion of the screen. Additional details and functionality appear if a VMware host is selected. The **Details** pane for VMware resources includes details about the last 7 days of backup jobs, the date of the last backup, and the date of the next scheduled backup. In addition, it includes the ability to restore data and to filter the list of guest virtual machines.

The **Details** pane for VMware virtual machines includes the resources that may not appear in the list of servers, such as:

- Virtual machines that do not have the Agent for Windows installed.
- Non-Windows virtual machines, such as Linux.
- Virtual machines that do not have a GRT-enabled backup.

The **Details** pane for VMware virtual machines provides a combination of current information and historical information. The **Refresh** button enables Backup Exec to discover all server resources. However, if a virtual machine has been moved, deleted, or has had a credentials change, then the following occurs:

- If the virtual machine has never been backed up, Backup Exec removes it from the **Details** pane.
- If the virtual machine has been backed up, Backup Exec does not remove it from the **Details** pane.

If a virtual machine is renamed, then the following occurs in the **Details** pane:

- If the virtual machine has been backed up, the **Details** pane includes an entry for the new name and keeps the entry for the old name.
- If the virtual machine has never been backed up, the **Details** pane includes only the new virtual machine name.

To view details about VMware resources

- ◆ On the **Backup and Restore** tab, select a VMware host from the list of servers.
The resources for the selected host appear in the **Details** pane.

Installing the Agent for Windows on VMware virtual machines

To use Backup Exec's Granular Recovery Technology (GRT) for Microsoft application data, install the Agent for Windows on any virtual machines that run Windows. To perform this procedure, you must have already added the vCenter or ESX/ESXi host to the list of servers on the **Backup and Restore** tab.

Note: VMware Tools should be installed before you install the Agent for Windows on a guest virtual machine if you intend to use the Backup Exec VSS provider.

See [“Adding VMware vCenter servers and ESX/ESXi hosts to the list of servers on the Backup and Restore tab”](#) on page 834.

To install the Agent for Windows on VMware virtual machines

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Select **VMware vCenter or ESX host**, and then click **Next**.
- 3 Select **Install the Backup Exec Agent for Windows on the virtual machines of a VMware vCenter server or ESX host**.
- 4 Select the server from the drop-down list, and then click **Next**.
- 5 Check **Allow Backup Exec to establish a trust with the servers**, and then click **Next**.
- 6 Select the guest virtual machines that you want to install the Agent for Windows on, and then click **Next**.
- 7 Select the logon account for the guest virtual machines, and then click **Next**.

- 8 Select the following options, and then click **Next**.

Upgrade the Backup Exec Agent for Windows to the current version automatically

Select this option to install the most recent version of the Agent for Windows if an older version of the Agent for Windows is already installed on the selected virtual machines.

Note: If you simply want to reestablish the trust with the existing version of the Agent for Windows that is already installed on the selected virtual machines, you should uncheck this check box.

Restart the remote computer automatically after installing the Backup Exec Agent for Windows when a restart is required

Select this option to enable Backup Exec to automatically restart the remote computer, if required.

- 9 Review the summary, and then click **Install**.

Push-installing the Agent for Windows to VMware virtual machines

You use the **Add Server** wizard to push-install the Agent for Windows to the VMware virtual machines that you select. You can choose to install the Agent for Windows on all of the guest virtual machines that need it or on specific machines only. For example, if you know that some of your guest virtual machines are scratch machines and do not need to be backed up, you can exclude those virtual machines from the installation.

Note: VMware Tools should be installed before you install the Agent for Windows on a guest virtual machine if you intend to use the Backup Exec VSS provider.

To push-install the Agent for Windows to virtual machines

- 1 On the **Backup and Restore** tab, right-click the virtual host, and then select **Install Backup Exec Agent for Windows to guest virtual machines**.
- 2 Check the **Allow Backup Exec to establish a trust with the servers** option, and then click **Next**.
- 3 Select **Install the Backup Exec Agent for Windows on the guest virtual machines of a VMware vCenter or ESX server**.

- 4 In the **VMware vCenter or ESX server** field, select the VMware server that hosts the virtual machine.
- 5 Click **Next**.
- 6 Select the guest virtual machines on which you want to install the Agent for Windows, and then click **Next**.
- 7 Verify that the virtual machines you selected are online and select the appropriate logon account for those virtual machines, and then click **Next**.
- 8 Review the information on the **Summary** panel, and then click **Install**.

Backing up VMware data

When you create a backup job, you can select an entire vCenter server or ESX/ESXi host, datacenters, folders, or individual virtual machines. Additionally, Backup Exec can automatically back up new virtual machines and folders that are found when a backup job runs. If you select the vCenter server or the ESX/ESXi host for a backup job, all virtual machines are backed up. However, a backup of a vCenter server or ESX/ESXi host does not include independent disks or configuration files for the vCenter server or ESX/ESXi host.

Warning: Backup jobs fail for virtual machines that have Physical Raw Disk Mapping (RDM) devices.

Virtual compatibility mode RDM disks are automatically included in the backup of a guest virtual machine. However, Backup Exec can only restore the virtual compatibility mode RDM disks or file data on an RDM disk through redirected restore.

Note: If you want to back up an Exchange database availability group (DAG) on a virtual machine, use the Backup Exec Agent for Exchange. The Agent for VMware does not support backups of Exchange DAG.

Note: You cannot back up databases to devices that are attached to a computer on which the Remote Media Agent for Linux Servers is installed.

Note: If you back up a virtual machine that runs any supported version of Microsoft SharePoint, you may experience a processing time of more than 30 minutes.

To back up VMware data

- 1 On the **Backup and Restore** tab, right-click the VMware host server or virtual machine that you want to back up.

Note: You can select an individual virtual machine for backup or you can back up a virtual machine by selecting its host server for backup. If you back up a virtual machine through the vCenter server or ESX/ESXi host, you can use Granular Recovery Technology (GRT) to restore individual files and folders as well as items from Microsoft Exchange, SQL, SharePoint, and Active Directory. If you select an individual virtual machine for backup, you cannot enable the backup to use GRT.

- 2 Select **Backup**, and then select the type of backup you want to perform.
- 3 On the **Backup Definition Properties** dialog box, in the **Selections** box, click **Edit**.
- 4 On the **Backup Selections** dialog box, check the check boxes for the resources that you want to back up and uncheck the check boxes for the resources that you do not want to back up.

Note: If you selected a VMware host server in step 1, the server is not selected automatically for backup on the **Backup Selections** dialog box. You must manually select the server and any guest virtual machines that you want to back up.

- 5 Click **OK**.
- 6 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**.
- 7 On the **Backup Options** dialog box, in the left pane, select **Schedule**, and then select the schedule for this job.
- 8 On the **Backup Options** dialog box, in the left pane, select **Virtual Machines**.
- 9 Set any of the following options for this job:

Item	Description
Use the full backup method for virtual machines that do not support incremental or differential backups	Select this option if you want Backup Exec to perform a full backup if an incremental backup or a differential backup cannot be performed. If you do not select this option and Backup Exec cannot perform an incremental backup or a differential backup, then the job fails. In addition, if Backup Exec detects a configuration change, then a full backup must be performed. If a configuration change is detected and Backup Exec cannot perform a full backup, then the job fails if this option is not selected. This scenario applies only if a full backup and some incremental backups or differential backups have already been performed and the next scheduled job is for an incremental backup or a differential backup.
Back up virtual machines that are powered off	Select this option if you want Backup Exec to back up any virtual machines that are turned off.
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual files and folders from virtual machines	<p>Select this option to enable individual files and folders to be restored from the backup. This option is for the virtual machines that use a Windows operating system only.</p> <p>The vmdk file is not backed up if the virtual hard disk is configured as an Independent disk.</p> <p>Note: GRT is not meant for system recovery but only for the restore of individual files and folders on Windows computers.</p> <p>See “Using Granular Recovery Technology (GRT) with the Agent for VMware” on page 847.</p>
Enable GRT for Microsoft Active Directory objects on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual Active Directory objects on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft Active Directory is installed.
Enable GRT for Microsoft Exchange databases and mailbox items on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual Exchange databases and mailbox items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft Exchange is installed.

Item	Description
Enable GRT for Microsoft SQL (database-level only) on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual SQL database items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft SQL is installed.
Run a SQL Log backup after backing up the virtual machine	<p>Select this option to enable Backup Exec to back up the SQL logs for the databases that use logging. After the logs are backed up, the data from the logs is committed to the database and the log is emptied in order to receive new data.</p> <p>If this option is not selected, the SQL logs continue to grow until the disk is full or until you perform a manual backup job to back up the logs.</p>
Enable GRT for Microsoft SharePoint on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual SharePoint items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft SharePoint is installed.
Back up using Microsoft Volume Shadow Copy Service (VSS) snapshot provider	<p>Select this option to enable Backup Exec to create a quiesced snapshot for the virtual machine. If the quiesced snapshot cannot be created, then Backup Exec creates a crash-consistent snapshot instead.</p> <p>Note: Backup Exec automatically selects this option if you select any of the options for enabling GRT for Microsoft applications.</p>
Use VSS Copy backup	Select this option to enable Backup Exec to perform a VSS copy backup instead of a full backup. By default the Symantec VSS Provider initiates snapshots using the Full backup type setting. Each application responds differently to this request. In the case of Microsoft Exchange, the database logs are truncated. If you do not select this option, Backup Exec performs a full VSS backup

Item	Description
Transport mode priority list	<p>Select the method to transport the Virtual Machine Disk Format (vmdk) file from the ESX/ESXi host. You must select at least one of these options. If you select more than one option, the method is determined by the priority and the resources that are available. You can move the options up or down in the list to change the priority.</p> <p>The following methods are available:</p> <ul style="list-style-type: none">■ SAN - Use the SAN to move virtual disk data If you select this option, the virtual machine must reside on a SAN that the Backup Exec server can access. With this transport mode, the data is offloaded to the Backup Exec server so that the ESX/ESXi host is not affected.■ NBD - Do not encrypt the virtual disk data for over-the-network transfers Use this option if you do not use SSL for security and one of the following conditions exist:<ul style="list-style-type: none">■ The virtual machine is not located on the SAN.■ The Backup Exec server does not have access to the SAN.■ NBDSSL - Encrypt virtual disk data for over-the-network transfers Use this option if you use SSL for security and one of the following conditions exist:<ul style="list-style-type: none">■ The virtual machine is not located on the SAN.■ The Backup Exec server does not have access to the SAN.■ Hotadd - Use virtual disk files from the Backup Exec server on the virtual machine Use this option if you want to use the hotadd feature for ESX/ESXi. See your VMware documentation for more information about hotadd. <p>The vmdk file is not backed up if the virtual hard disk is configured as an Independent disk.</p>
vSphere Port Number	<p>Indicate the port that Backup Exec uses to connect to the vCenter Server. The default port is 902.</p>

Item	Description
Backup method	Select the backup method that you want to use for the backup jobs that are listed. You can change the names of the backup jobs or add more jobs from the Schedule properties. See “Configuring backup methods for backup jobs” on page 196.

10 On the **Backup Options** dialog box, click any of the optional settings in the left pane that you want to set for this job.

11 Click **OK**.

12 On the **Backup Definition Properties** dialog box, click **OK**.

See [“How Backup Exec automatically backs up new VMware virtual machines during a backup job”](#) on page 847.

See [“Using Granular Recovery Technology \(GRT\) with the Agent for VMware”](#) on page 847.

Setting default backup options for virtual machines

You can use the defaults that Backup Exec sets during installation for all VMware backup jobs, or you can choose your own defaults. You can override the default settings when you create individual jobs.

To set default backup options for virtual machines

1 Click the Backup Exec button, and then select **Configuration and Settings**.

2 Select **Job Defaults**, and then select a backup option.

For example, if you want to set up the default options for virtual machine backups to disk, select **Back Up to Disk**. The options that appear will vary depending on what types of storage devices you have configured. Different default options can be configured for backup jobs to different types of storage.

3 In the left pane, select **Virtual Machines**.

4 Select the appropriate options.

Item	Description
Use the full backup method for virtual machines that do not support incremental or differential backups	Select this option to enable Backup Exec to perform a full backup if an incremental backup or a differential backup cannot be performed. If you do not select this option and Backup Exec cannot perform an incremental backup or a differential backup, then the job fails. In addition, if Backup Exec detects a configuration change, then a full backup must be performed. If a configuration change is detected and Backup Exec cannot perform a full backup, then the job fails if this option is not selected. This scenario applies only if a full backup and some incremental backups or differential backups have already been performed and the next scheduled job is for an incremental backup or a differential backup.
Back up virtual machines that are powered off	Select this option to enable Backup Exec to back up virtual machines that are turned off.
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual files and folders from virtual machines	<p>Select this option to enable individual files and folders to be restored from the backup. This option is for virtual machines that use a Windows operating system only.</p> <p>The vmdk file is not backed up if the virtual hard disk is configured as an Independent disk.</p> <p>Note: GRT is not meant for system recovery but only for the restore of individual files and folders on Windows computers.</p>
Enable GRT for Microsoft Active Directory objects on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual Active Directory objects on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft Active Directory is installed.
Enable GRT for Microsoft Exchange databases and mailbox items on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual Exchange databases and mailbox items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft Exchange is installed.
Enable GRT for Microsoft SQL (database-level only) on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual SQL database items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft SQL is installed.

Item	Description
Run a SQL Log backup after backing up the virtual machine	<p>Select this option to enable Backup Exec to back up the SQL logs for databases that use logging. After the logs are backed up, the data from the logs is committed to the database and the log is emptied in order to receive new data.</p> <p>If this option is not selected, the SQL logs continue to grow until the disk is full or until you perform a manual backup job to back up the logs.</p>
Enable GRT for Microsoft SharePoint on virtual machines	<p>Select this option to enable Backup Exec to collect the information that is required to restore individual SharePoint items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Microsoft SharePoint is installed.</p>
Exclude virtual machines that must be put into a saved state for backup	<p>Select this option to exclude from the backup all offline virtual machines that do not support online backups and that are in a running state when the backup begins.</p> <p>This option applies only to the Agent for Hyper-V.</p>
Back up using Microsoft Volume Shadow Copy Service (VSS) snapshot provider	<p>Select this option to enable Backup Exec to create a quiesced snapshot for the virtual machine. If the quiesced snapshot cannot be created, then Backup Exec creates a crash-consistent snapshot instead.</p> <p>Note: Backup Exec automatically selects this option if you select any of the options for enabling GRT for Microsoft applications.</p>
Use VSS Copy backup	<p>Select this option to enable Backup Exec to perform a VSS copy backup instead of a full backup. By default the Symantec VSS Provider initiates snapshots using the Full backup type setting. Each application responds differently to this request. In the case of Microsoft Exchange, the database logs are truncated. If you do not select this option, Backup Exec performs a full VSS backup</p>

Item	Description
Transport mode priority list	<p>Select the method to transport the Virtual Machine Disk Format (vmdk) file from the ESX/ESXi host. You must select at least one of these options. If you select more than one option, the method is determined by the priority and the resources that are available. You can move the options up or down in the list to change the priority.</p> <p>The following methods are available:</p> <ul style="list-style-type: none">■ SAN - Use the SAN to move virtual disk data If you select this option, the virtual machine must reside on a SAN that the Backup Exec server can access. With this transport mode, the data is offloaded to the Backup Exec server so that the ESX/ESXi host is not affected.■ NBD - Do not encrypt the virtual disk data for over-the-network transfers Use this option if you do not use SSL for security and one of the following conditions exist:<ul style="list-style-type: none">■ The virtual machine is not located on the SAN.■ The Backup Exec server does not have access to the SAN.■ NBDSSL - Encrypt virtual disk data for over-the-network transfers Use this option if you use SSL for security and one of the following conditions exist:<ul style="list-style-type: none">■ The virtual machine is not located on the SAN.■ The Backup Exec server does not have access to the SAN.■ Hotadd - Use virtual disk files from the Backup Exec server on the virtual machine Use this option if you want to use the hotadd feature for ESX/ESXi. See your VMware documentation for more information about hotadd. <p>The vmdk file is not backed up if the virtual hard disk is configured as an Independent disk.</p>
vSphere Port Number	<p>Indicate the port that Backup Exec uses to connect to vCenter Server. The default port is 902.</p>

Item	Description
Backup method	<p>Select the backup method that you want to use for the backup jobs that are listed. You can change the names of the backup jobs or add more jobs from the Schedule properties.</p> <p>See “Configuring backup methods for backup jobs” on page 196.</p>

5 Click **OK**.

See [“Backing up VMware data”](#) on page 838.

How Backup Exec automatically backs up new VMware virtual machines during a backup job

Backup Exec's dynamic inclusion feature protects new virtual machines and folders that are found when a backup job runs. If new virtual machines are added between the time when the backup job is created and when the backup job runs, Backup Exec automatically backs up the new virtual machines. Because the backup job may include new virtual machines, the job may require more storage space and more time to run than you anticipated. The job history shows the number of virtual machines that were backed up.

If you select a VMware server, then dynamic inclusion is enabled automatically for all of the nodes below it that have a folder icon. If no virtual machines are located during a backup job, then the job fails.

See [“Backing up VMware data”](#) on page 838.

Using Granular Recovery Technology (GRT) with the Agent for VMware

Backup Exec's Granular Recovery Technology (GRT) lets you restore individual drives, files, and folders from backup sets without having to restore the entire virtual machine. It also lets you restore individual items from the following VSS-aware applications that reside on virtual machines:

Table C-2 Types of data that Backup Exec backs up for VSS-aware applications on virtual machines

Application	Types of data that Backup Exec backs up
Microsoft Exchange	Mailboxes, individual messages, calendar items, tasks, journal entries, and public folder data (disk-backups only)
Microsoft SQL	Databases
Microsoft Active Directory	Individual user accounts, printer objects, sites, and organizational units
Microsoft SharePoint	SharePoint data

GRT works only for the virtual machines that use a Windows operating system. GRT does not work for system recovery.

GRT must be enabled in the VMware backup job. When you create a backup job, Backup Exec automatically locates VSS-aware applications on virtual machines. During the backup job, Backup Exec backs up the data from the VSS-aware applications by using GRT. By default, Backup Exec enables GRT using the same credentials that were used to connect to the virtual machine. You can disable GRT for any of the VSS-aware application types.

Note: Backup Exec supports the granular recovery of individual Exchange and SQL items only in non-clustered and non-distributed configurations.

During the backup job, Backup Exec collects metadata from the applications. If Backup Exec is unable to collect the metadata, then you cannot restore individual items for the applications. However, the backup job may otherwise complete successfully.

Requirements for using GRT to back up Microsoft application data on virtual machines

The following items are required to back up data for Microsoft Exchange, SQL, SharePoint, and Active Directory on virtual machines:

- The virtual machine must be turned on.
- You must enter the appropriate credentials for the virtual machine. Ensure that the credentials for the virtual machine allow access to the VSS-aware applications.

- The Backup Exec server must be able to connect to the virtual machine using the network name or IP address.
- VMware Tools should be installed on the virtual machine.
- The Backup Exec Agent for Windows must be installed on the virtual machine. Be sure that the VMware Tools are installed before you install the Agent for Windows.

Note: If you installed the VMware Tools after you installed the Agent for Windows, you should copy the file `freeze.bat` from Backup Exec RAWs Install Path\VSS Provider to VMware Tools install path\backupscripsts.d, for example \\Program Files\VMware\VMware Tools\backupscripsts.d. If the backupscripsts.d directory does not exist, you must create it manually. Alternatively, instead of moving the `freeze.bat` file, you can repair the Agent for Windows, which automatically places all missing files in the appropriate location.

- The Backup Exec Agent for Applications and Databases must be installed on the Backup Exec server.
- The correct number of licenses must be entered for the applications that you want to protect on the virtual machines.
- The operating system on the virtual machine must support VSS.
- The options for application-level GRT must be selected in backup jobs. These options are not selected by default.

Unsupported configurations for GRT

Before you create a GRT-enabled backup job for VMware resources, review the following information to understand what configurations are not supported for GRT.

Table C-3 Unsupported configurations for GRT

Unsupported items	Details
Virtual machines that have a combination of GPT and MBR disks.	Backup Exec does not support GRT of GPT disks. If a virtual machine has both GPT and MBR disks, then GRT cannot be enabled for either type of disk. The virtual machine is backed up, but the job completes with a status of success with exceptions.
Virtual machines that have a combination of independent and non-independent disks	Backup Exec does not support GRT of independent disks.

Table C-3 Unsupported configurations for GRT (*continued*)

Unsupported items	Details
Virtual machines that have a virtual RDM disk	If you try to restore a backup of a virtual machine that has a virtual RDM disk, the vmdk that corresponds to the virtual RDM disk cannot be created or restored. The restore job fails with the error "Unable to open a disk of the virtual machine". Only non-virtual RDM disks can be restored.
Virtual machines that have RAID 5 disks	Backup Exec does not support file/folder-GRT of RAID 5 disks.
Virtual machines that have NTFS with unnamed mount points	Backup Exec does not support file/folder-GRT of NTFS with unnamed mount points. The restore job fails with the error "Unable to attach to a resource. Ensure that the selected resource exists and is online, and then try again."
Virtual machines that have utility partitions	File/folder-level GRT restore fails with the error "Unable to attach to a resource. Ensure that the selected resource exists and is online, and then try again." Backup Exec does not support backups of virtual machines that have utility partitions. Note that this is not the same as unnamed partitions.
Virtual machines that have one or more disks larger than 2 TB	<p>Backup Exec does not support GRT for VMware virtual machines with disks that are 2 TB or larger. Backup Exec automatically disables the GRT settings for virtual machines with disks that are 2 TB or larger. Therefore, even if GRT settings are enabled in a backup job, a non-GRT backup is performed for the virtual machines with disks that are 2 TB or larger. GRT is performed for the other virtual machines in the same backup job if they do not have disks that are 2 TB or larger.</p> <p>In addition, Backup Exec does not support the option to restore a virtual disk to a different path if the virtual disk is 2 TB or larger. Virtual disks that are 2 TB or larger can be restored to the original location or to a different vCenter server or ESX/ESXi host.</p>

Table C-3 Unsupported configurations for GRT (*continued*)

Unsupported items	Details
Restores of full and incremental backup sets from different storage devices	Backup Exec does not support restores from mixed media if GRT was enabled in the backup job. For example, if the full backup is on tape and the incremental backup is on a disk storage device, the restore job will fail. Restores from mixed media types are supported if GRT is not enabled.

See [“Backing up VMware data”](#) on page 838.

How delayed cataloging works with VMware virtual machine backups

When you enable Granular Recovery Technology (GRT) for a backup job of a virtual machine, you can choose to run the catalog job for GRT as part of the backup job, as a separate job immediately after the backup job completes, or according to a schedule. By default, the catalog operation runs immediately after the backup job completes.

Note: The delayed catalog feature is not supported for backups to tape.

The catalog operation can be time consuming. It requires access to the storage device that is used for the backup. You may want to schedule the catalog operation to run outside of your backup window so that it does not interfere with backup jobs. If the catalog operation is scheduled, it runs only for the most recent backup set since the last catalog operation. In this situation, only the most recent backup set since the last catalog operation can be used for granular recovery on VMware virtual machines. GRT data is not available until the delayed catalog job completes, but a full restore of a virtual machine can be performed before the delayed catalog job completes.

For example, if you set up incremental backups to run every 11 hours and set up the catalog operation to run at midnight, you would have the following backup sets:

- Full (11:00 A.M.)
- Incremental 1 (10:00 P.M.)
- Catalog 1 (Midnight). This job catalogs Incremental 1.
- Incremental 2 (9:00 A.M.)
- Incremental 3 (8:00 P.M.)

- Catalog 2 (Midnight). This job catalogs Incremental 3. Incremental 2 is not cataloged.
- Incremental 4 (7:00 A.M.)
- Incremental 5 (6:00 P.M.)
- Catalog 3 (Midnight). This job catalogs Incremental 5. Incremental 4 is not cataloged.
- Incremental 6 (5:00 A.M.). This backup is not cataloged.

In the example above, you can perform a granular recovery from Incremental 5, Incremental 3, or Incremental 1. You can perform a full restore from any of the above backup sets.

How byte count is calculated for delayed catalog operations

On the **Job Monitor** and **Job History**, the byte count that displays for a catalog operation may differ from the byte count that displays for the corresponding backup job. The byte count for a catalog job may be larger than the byte count for a backup job. The way in which Backup Exec catalogs the data affects the byte count that appears for the catalog job.

- When a catalog operation is performed for a full backup, the data is read on a file-by-file basis and the byte count is calculated accordingly. During the full backup job, the data is read in terms of the number of sectors and the byte count is calculated based on the number of sectors. Therefore, the byte count for the catalog job may be larger than the byte count for the backup job.
- When the catalog operation is performed for an incremental backup, all files on the virtual disk are cataloged instead of only the changed files. Therefore, the byte count for the catalog job takes into account both the full backup and the incremental backup.

Restoring VMware virtual machines and vmdk files

You can use the Restore Wizard to restore the following:

- A complete virtual machine.

Note: To restore any virtual machines that were backed up with Backup Exec 12.5, the VMware Converter (4.01 or later) must be installed on the Backup Exec server.

- The Virtual Machine Disk Format (vmdk) file for a virtual machine.

- Individual files and folders that were backed up from inside the vmdk file and individual items from Microsoft SharePoint, Exchange, SQL, or Active Directory. The ability to restore individual files and folders is available only if Granular Recovery Technology (GRT) was selected for the backup job.

Note: GRT cannot restore system state files such as the active registry.

Backup Exec lets you restore VMware data to either the location from which the data was backed up or to a different location. Restoring data to a different location is referred to as a redirected restore.

Note: The steps below apply to virtual machines that were backed up through the virtual host with the Agent for VMware. If you did not back up the virtual machine through the virtual host with the Agent for VMware, follow the steps for restoring a non-virtual backup.

See [“Restoring data from a server, a backup set, a backup job, or a storage device”](#) on page 234.

See [“To restore VMware virtual machines or vmdk files to the same location from which they were backed up”](#) on page 853.

See [“To redirect the restore of VMware virtual machines or vmdk files to a different host”](#) on page 855.

See [“To restore to a different path”](#) on page 858.

To restore VMware virtual machines or vmdk files to the same location from which they were backed up

- 1 On the **Backup and Restore** tab, do one of the following:

To restore individual files and folders from a GRT-enabled backup Do the following in the order listed:

- In the **Details** pane at the bottom of the screen, select the virtual machine.
- Click **Restore**, and then select **Restore GRT-enabled data**.
- In the **Restore Wizard**, select **Files, folders, or volumes**, and then click **Next**.

To restore the entire virtual machine or virtual disks

Do the following in the order listed:

- In the **Details** pane at the bottom of the screen, select the virtual machine.
- Click **Restore**, and then select **Restore virtual machine from the host**.
- In the **Restore Wizard**, select **VMware data**, and then click **Next**.

- 2 On the **Resource View** tab, expand the virtual host server name or the virtual machine name, select the items that you want to restore, and then click **Next**.
- 3 If the **Where do you want to restore the data from** panel appears, verify that you want to restore from the selected storage or select different storage to restore from, and then click **Next**.
- 4 Select **To the original location**, and then click **Next**.
- 5 Select one or more transport modes to transport the vmdk file to the host.

You must select at least one of the transport mode options. If you select more than one option, the method is determined by the priority and the resources that are available. Click **Move Up** or **Move Down** to arrange the transport modes in the priority that you want to use.

NBD-Do not encrypt the virtual disk data for over-the-network transfers

Use this option if you do not use SSL for security and one of the following conditions exists:

- The virtual machine is not located on the SAN.
- The Backup Exec server does not have access to the SAN.

NBDSSL-Encrypt virtual disk data for over-the-network transfers

Use this option if you use SSL for security and one of the following conditions exists:

- The virtual machine is not located on the SAN.
- The Backup Exec server does not have access to the SAN.

Hotadd-Use virtual disk files from the Backup Exec server on the virtual machine

Use this option if you want to use the hotadd feature for ESX/ESXi. The hotadd feature lets you use a virtual machine as your proxy server. See your VMware documentation for more information about hotadd .

- 6 Complete any of the following optional fields that apply to your environment, and then click **Next**.

Delete existing virtual machines prior to restore Select this option if the virtual machine that you are restoring already exists on the server. If you select this option, the virtual machines may be deleted even if the restore job fails. You cannot restore a virtual machine if it already exists on the virtual server unless you select this option.

Power on virtual machine after restore Select this option if you want Backup Exec to turn on the restored virtual machine after the restore job completes.

vSphere port number Enter the port that Backup Exec uses to connect to vCenter Server. The default port is 902.

- 7 If you want to run a command before or after the restore or enable notification for this restore job, complete the fields on the **What additional tasks do you want to perform before and/or after a restore** panel, and then click **Next**.
- 8 Enter a name for this restore job and select the schedule for the job, and then click **Next**.
- 9 Review the job summary, and then click **Finish**.

To redirect the restore of VMware virtual machines or vmdk files to a different host

- 1 On the **Backup and Restore** tab, do one of the following:

To restore individual files and folders from a GRT-enabled backup Do the following in the order listed:

- In the **Details** pane at the bottom of the screen, select the virtual machine.
- Click **Restore**, and then select **Restore GRT-enabled data**.
- In the **Restore Wizard**, select **Files, folders, or volumes**, and then click **Next**.

To restore the entire virtual machine or virtual disks

Do the following in the order listed:

- In the **Details** pane at the bottom of the screen, select the virtual machine.
 - Click **Restore**, and then select **Restore virtual machine from the host**.
 - In the **Restore Wizard**, select **VMware data**, and then click **Next**.
- 2 On the **Resource View** tab, expand the virtual host server name or the virtual machine name, select the items that you want to restore, and then click **Next**.
 - 3 Select **To a different vCenter or ESX server**, and then click **Next**.
 - 4 Type the name of the vCenter server or ESX/ESXi host that you want to restore to, or click **Browse** to search for the server.
 - 5 Select the correct logon account for the server that you want to restore the data to.
 - 6 Click **Browse** next to **Virtual machine folder** to select the folder to which you want to restore.
 - 7 Click **Browse** next to **Resource pool** to select the resource pool to which you want to restore.
 - 8 If you want to create a new name for the virtual machine, type it in the **Virtual machine name** field. You might want to enter a new name if a virtual machine with the same name already exists on the server.
 - 9 Select the network that the new virtual machine should use after the restore job completes.

- 10 Complete any of the following optional fields that apply to your environment, and then click **Next**.

Use the original disk datastore selections if available on the selected host

Check this check box to use the original datastore selections on the virtual server. If the original datastore selections do not exist, then the datastore selections from the backup data are used.

Restore virtual machine to the most recent hardware version that the destination environment supports

Check this check box to restore the virtual machine with VMware hardware version 7 or later.

Restore virtual clients with thin provisioning

Check this check box to restore the virtual machine with thin provisioning. Thin provisioning can help you more efficiently dedicate storage capacity in your VMware ESX Server version 4.0 environment.

- 11 Select one or more transport modes to transport the vmdk file to the host.

You must select at least one of the transport mode options. If you select more than one option, the method is determined by the priority and the resources that are available. Click **Move Up** or **Move Down** to arrange the transport modes in the priority that you want to use.

NBD-Do not encrypt the virtual disk data for over-the-network transfers

Use this option if you do not use SSL for security and one of the following conditions exists:

- The virtual machine is not located on the SAN.
- The Backup Exec server does not have access to the SAN.

NBDSSL-Encrypt virtual disk data for over-the-network transfers

Use this option if you use SSL for security and one of the following conditions exists:

- The virtual machine is not located on the SAN.
- The Backup Exec server does not have access to the SAN.

Hotadd-Use virtual disk files from the Backup Exec server on the virtual machine

Use this option if you want to use the hotadd feature for ESX/ESXi. The hotadd feature lets you use a virtual machine as your proxy server. See your VMware documentation for more information about hotadd.

- 12 Complete any of the following optional fields that apply to your environment, and then click **Next**.

Delete existing virtual machines prior to restore

Select this option if the virtual machine that you are restoring already exists on the server. If you select this option, the virtual machines may be deleted even if the restore job fails. You cannot restore a virtual machine if it already exists on the virtual server unless you select this option.

Power on virtual machine after restore

Select this option if you want Backup Exec to turn on the restored virtual machine after the restore job completes.

vSphere port number

Enter the port that Backup Exec uses to connect to vCenter Server. The default port is 902.

- 13 If you want to run a command before or after the restore or enable notification for this restore job, complete the fields on the **What additional tasks do you want to perform before and/or after a restore** panel, and then click **Next**.
- 14 Enter a name for this restore job and select the schedule for the job, and then click **Next**.
- 15 Review the job summary, and then click **Finish**.

To restore to a different path

- 1 On the **Backup and Restore** tab, do one of the following:

To restore individual files and folders from a GRT-enabled backup Do the following in the order listed:

- In the **Details** pane at the bottom of the screen, select the virtual machine.
- Click **Restore**, and then select **Restore GRT-enabled data**.
- In the **Restore Wizard**, select **Files, folders, or volumes**, and then click **Next**.

To restore the entire virtual machine or virtual disks Do the following in the order listed:

- In the **Details** pane at the bottom of the screen, select the virtual machine.
- Click **Restore**, and then select **Restore virtual machine from the host**.
- In the **Restore Wizard**, select **VMware data**, and then click **Next**.

- 2 On the **Resource View** tab, expand the virtual host server name or the virtual machine name, select the items that you want to restore, and then click **Next**.
- 3 Select **To a different path**.
- 4 Enter the drive and path that you want to restore to, and then click **Next**.
- 5 If you want to run a command before or after the restore or enable notification for this restore job, complete the fields on the **What additional tasks do you want to perform before and/or after a restore** panel, and then click **Next**.
- 6 Enter a name for this restore job and select the schedule for the job, and then click **Next**.
- 7 Review the job summary, and then click **Finish**.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

Symantec Backup Exec Agent for Microsoft Hyper-V

This appendix includes the following topics:

- [About the Agent for Microsoft Hyper-V](#)
- [Requirements for using the Agent for Microsoft Hyper-V](#)
- [About installing the Agent for Microsoft Hyper-V](#)
- [Using the Agent for Hyper-V with Windows Server 2012/2012 R2](#)
- [Disk space optimization with the Agent for Hyper-V](#)
- [Adding a Hyper-V host to the list of servers on the Backup and Restore tab](#)
- [Viewing details about Hyper-V resources](#)
- [Installing the Agent for Windows on Hyper-V virtual machines](#)
- [Push-installing the Agent for Windows to Hyper-V virtual machines](#)
- [Backing up Microsoft Hyper-V virtual machines](#)
- [Restoring Microsoft Hyper-V virtual machines](#)

About the Agent for Microsoft Hyper-V

The Symantec Backup Exec Agent for Microsoft Hyper-V (Agent for Hyper-V) lets you back up and restore the following resources:

- Microsoft Windows Server 2008/2008 R2/2012/2012 R2 Hyper-V hosts.
- All virtual machines that reside on the Hyper-V hosts.

- Clustered Hyper-V hosts, including the virtual machines that reside on cluster shared volumes (CSV).

Backup Exec performs a single-pass backup to protect the host configuration data, all virtual machines, and VSS-aware applications that are installed on the virtual machines. Backup Exec's file/folder-level Granular Recovery Technology (GRT) is enabled by default for backup jobs. You can use a GRT-enabled backup to restore individual files and folders from a Windows virtual machine without restoring the entire virtual machine. In addition, you can restore individual items from Microsoft Exchange, SharePoint, and Active Directory applications that reside on virtual machines. You can also restore individual databases from Microsoft SQL when it resides on virtual machines.

Note: You must have the appropriate Backup Exec agent for Microsoft Exchange, SQL, SharePoint, or Active Directory on the virtual machine to perform GRT.

Backup Exec can back up the virtual machines that are online or that are in an offline state or a saved state. The virtual machines that use Microsoft Windows 2003 or later (with Hyper-V Integration Services) can be backed up while they are online. You can include both online and offline virtual machines in the same backup job. During the backup of an online virtual machine, Backup Exec takes a snapshot backup of the Hyper-V host. The host in turn takes a snapshot of the virtual machines on the host. This process enables Backup Exec to back up virtual servers without any downtime. If an online backup cannot be performed and the option **Exclude virtual machines that must be put in a saved state to back up** is selected for the backup job, then an offline backup is performed. With an offline backup, the virtual machine is placed briefly in a saved state. However, the virtual machine does not remain in the saved state for the entire backup job.

The amount of downtime for a saved state backup job depends on the following:

- The amount of memory that is allocated to the virtual machine.
- The current load on the host's operating system.

See [“Requirements for using the Agent for Microsoft Hyper-V”](#) on page 862.

See [“Backing up Microsoft Hyper-V virtual machines”](#) on page 873.

See [“Restoring Microsoft Hyper-V virtual machines”](#) on page 885.

Requirements for using the Agent for Microsoft Hyper-V

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

The following items are required:

Table D-1 Requirements for the Agent for Microsoft Hyper-V

Software	Installed on
Microsoft Windows Server 2008/2008 R2 Hyper-V or Windows Server 2012/2012 R2 Hyper-V	Microsoft Hyper-V host
Backup Exec	Backup Exec server Note: Symantec recommends that the Backup Exec server should run a version of Windows that is equal to or greater than the highest version of Windows that the virtual machines in your environment run. For example, if your virtual machines run Microsoft Windows Server 2012, the Backup Exec server should also run Microsoft Windows Server 2012.
Agent for VMware and Hyper-V	Microsoft Hyper-V host Note: If you use Hyper-V in a cluster environment, you must install the Agent for VMware and Hyper-V on each node in the cluster. If you do not, you may not be able to see all of the clustered backup sources that are available for backup.

Table D-1 Requirements for the Agent for Microsoft Hyper-V (*continued*)

Software	Installed on
VHDMount	<p>Backup Exec server (if the Backup Exec server is not the virtual server).</p> <p>Note: VHDMount is required only if the Backup Exec server runs Microsoft Windows 2003/2008 without the Hyper-V role installed. You can install the VHDMount component from Microsoft Virtual Server 2005 R2 SP1. In addition, VHDMount is required for Granular Recovery Technology (GRT)-enabled backups. If you do not plan to use GRT, then VHDMount is not required.</p>

Requirements for online backups

To run an online backup, the following requirements must be met:

- Microsoft Windows Server 2012/2008/2003 SP2/Vista SP1/XP SP3 is installed on the virtual machine.
- Hyper-V Integration Services with Backup (Volume snapshot) is installed.

Note: Using an incorrect version of Hyper-V Integration Services may lead to a virtual machine being backed up offline instead of online. For example, if you copy a virtual machine from a Windows 2012 Hyper-V host to a Windows 2012 R2 Hyper-V host, the Hyper-V Integration Services are not updated automatically.

- The virtual machine is in a running state.

If these conditions are not met, the virtual machine is placed in a saved state if it is running. If the virtual machine is turned off, then that virtual machine is backed up only if you select the option **Back up virtual machines that are powered off** on the **Virtual Machines** options dialog box.

Requirements for collecting catalog data for Microsoft applications

To enable Backup Exec to collect catalog data for Microsoft Exchange, SharePoint, Active Directory, and SQL on the virtual machine, the following items are required on the virtual machine:

- A licensed version of the Backup Exec agent for the application.
- The Agent for Windows.

The Agent for Microsoft Hyper-V includes a license for the Agent for Windows. The agents for Microsoft Exchange, Active Directory, and SQL also include a license for the Agent for Windows. No separate license is required for the Agent for Windows.

- The virtual machine must be capable of being backed up online.
- The credentials that you use to access the virtual machine must also have access to the application.

Requirements for virtual machines

The following items are required for virtual machines:

- Virtual machines must have unique names. Backup Exec does not support virtual machines that have duplicate names.
- You must assign shadow storage to each NTFS volume that exists on a virtual machine if the virtual machine uses the Windows 2008 operating system. Otherwise, the backup jobs fail for that virtual machine.
- The Agent for Windows must be installed on the virtual machine to enable individual files and folders to be restored back to the original virtual machine.
- The Agent for Windows must be installed on the virtual machine to enable individual SQL databases, Exchange items, SharePoint items, and Active Directory objects to be restored back to the original virtual machine.
- Do not use special characters, such as the percent sign (%) or a forward slash (/) in Hyper-V virtual machine names. Special characters may cause redirected restores to fail.

See [“About the Agent for Microsoft Hyper-V”](#) on page 860.

About installing the Agent for Microsoft Hyper-V

The Symantec Backup Exec Agent for Microsoft Hyper-V is installed as part of the Agent for VMware and Hyper-V. The Agent for Microsoft Hyper-V is installed on the Microsoft Hyper-V host. If your Backup Exec server is also your Microsoft Hyper-V host, you can install the Agent for Microsoft Hyper-V when you install Backup Exec. Or, you can install it after Backup Exec has been installed.

If Backup Exec is not installed on your Microsoft Hyper-V host, you must push-install the Agent for Windows to your Microsoft Hyper-V host. You do not need to install the Agent for Microsoft Hyper-V on virtual machines. However, a license is required on the Backup Exec server for the Agent for Microsoft Hyper-V. The Agent for Windows is included with the Agent for Microsoft Hyper-V.

See [“Installing additional Backup Exec options to the local Backup Exec server”](#) on page 76.

See [“Push-installing the Agent for Windows to remote computers”](#) on page 86.

Using the Agent for Hyper-V with Windows Server 2012/2012 R2

The Backup Exec 2014 Agent for Hyper-V lets you back up and restore Hyper-V virtual machines that are hosted on Microsoft Windows Server 2012/2012 R2. This version of Backup Exec supports the Windows Server 2012 vhd file format and Microsoft incremental backups.

This topic includes the following information:

[Using GRT with vhd files when the Backup Exec server runs Windows Server 2008 R2 and earlier](#)

[Using Microsoft incremental backups with Backup Exec](#)

[Additional notes about using the Agent for Hyper-V with Windows Server 2012 Windows Server 2012 R2 and the Agent for Hyper-V](#)

Using GRT with vhd files when the Backup Exec server runs Windows Server 2008 R2 and earlier

Virtual machine file/folder-level GRT and application-level GRT are handled in the following ways for vhd files if the Backup Exec server runs Windows Server 2008 R2 or earlier. Note that these limitations do not apply if the Backup Exec server runs Windows Server 2012 or later.

- If a virtual machine has only vhd files, then both file/folder-level GRT and application-level GRT are fully supported.
- If a virtual machine has only vhd files, file/folder-level GRT and application-level GRT are supported if the vhd has a capacity of less than 2040 GB and the logical sector size within the vhd is 512 bytes. If the vhd files have a capacity larger than 2040 GB, then file/folder-level GRT and application-level GRT are not supported. The job completes with a status of success with exceptions, but no GRT restore capability will be available. A full restore of the virtual machine will be available.
- If a virtual machine has a mixture of vhd and vhd files, but all of the vhd files have a capacity of less than 2 TB, then both file/folder-level GRT and application-level GRT are fully supported.

- If a virtual machine has a mixture of vhd and vhdx files, but one or more vhdx files have a capacity of greater than 2 TB, then file/folder-level GRT and application-level GRT are not supported. The job completes with a status of success with exceptions, but no GRT restore capability will be available. A full restore of the virtual machine will be available.
- Application-level GRT is supported for Microsoft Active Directory 2012, SQL 2012, SharePoint Server 2013, and Exchange Server 2013.
- File/folder-level GRT is not supported for any volumes that use the Windows 2012 Resilient File System (ReFS). In addition, application-level GRT cannot be performed for the virtual machine on which the ReFS volume is located.

Note: File/folder-level GRT is supported for any volumes that are configured with NTFS.

- File/folder-level GRT is not supported for any volumes that have Windows 2012 deduplication enabled. In addition, application-level GRT cannot be performed for the virtual machine on which Windows 2012 deduplication is enabled.

Note: File/folder-level GRT is supported for any volumes that are configured with NTFS.

Using Microsoft incremental backups with Backup Exec

Microsoft's incremental backups for Hyper-V differ from Backup Exec's incremental backups. No options are available within Backup Exec to configure the Microsoft incremental backup method for a virtual machine. Instead, you must use PowerShell on the Hyper-V host to configure the Microsoft incremental backup. The Microsoft incremental backup setting is applied as a property of each individual virtual machine, so some virtual machines may use the Microsoft incremental backup method while other virtual machines may not use it. You must configure a virtual machine to use the Microsoft incremental backup method before you run a backup job. For the virtual machines that do not have the Microsoft incremental backup method enabled, Backup Exec incremental backups are performed. For the virtual machines that do have the Microsoft incremental backup method enabled, the Microsoft incremental backup is performed instead of the Backup Exec incremental backup. The Backup Exec job log lists the incremental backup method that was used. A Backup Exec differential backup of virtual machines that have Microsoft incremental backup enabled is not supported, so a job with that configuration fails.

Note: To use Microsoft's incremental backups, the Hyper-V host must have Windows Server 2012 installed and the virtual machine must use the Windows Server 2012 version of the Hyper-V Integration Services.

Note: A one-time backup job may cause the virtual machines that have Microsoft incremental backups enabled to perform a full backup during the next scheduled full or incremental job.

Additional notes about using the Agent for Hyper-V with Windows Server 2012

The following limitations apply to using the Agent for Hyper-V with Windows Server 2012:

- Any virtual machines that are configured with Fibre Channel HBAs are skipped during backup jobs. The job log lists a message to indicate that the virtual machine was skipped.
- The deduplication stream handler for vhd files is not supported.
- Remote VSS is not supported. In other words, virtual machines with vhd or vhdx files that are located on shares are not supported.
- A VSS Copy backup of a Windows 2012 Hyper-V virtual machine is not supported.
- A virtual machine that was backed up from a Hyper-V 2012 host cannot be restored to a virtual machine with an earlier version of Hyper-V, such as 2008 or 2008 R2.
- Any virtual machines that have storage spaces are not supported because Microsoft Hyper-V VSS is unable to take snapshots of virtual machines with storage spaces.
- Hyper-V Replication virtual machine backups may result in redundant backups of the primary virtual machine and the replicated virtual machine.
- GRT is not supported for virtual machines that have vhdx image files with a logical sector size of 4096 bytes. This limitation applies to both application-level GRT and file/folder-level GRT. A full restore of a virtual machine with that configuration can be performed.
- Virtual machines that are configured with remote storage are skipped during backup jobs. The job log contains a message about the virtual machines that were skipped.

Windows Server 2012 R2 and the Agent for Hyper-V

In addition to the items in the list above, the following items also apply to using the Agent for Hyper-V with Windows Server 2012 R2:

- Virtual machines that have shared vhd files are excluded from a backup job and the job fails.
- The restore of a Gen2 virtual machine can be redirected only to another Hyper-V host that runs Windows Server 2012 R2.
- If you want to enable GRT for a Hyper-V virtual machine that runs Active Directory or Active Directory Lightweight on Windows Server 2012 R2 and back up that virtual machine to tape, both the Hyper-V host and the Backup Exec server must run Windows Server 2012 R2.

See [“Requirements for using the Agent for Microsoft Hyper-V”](#) on page 862.

Disk space optimization with the Agent for Hyper-V

Backup Exec attempts to perform disk space optimization during Hyper-V backup jobs. Disk space optimization is performed whenever possible.

However, disk space optimization cannot be performed in the following situations:

- The file system is not NTFS. For example, if the file system is FAT, ReFS, or Linux, then disk space optimization is not performed.
- The disks are configured with Windows Storage Spaces.
- The volumes are configured with dynamic volumes, such as spanned, striped, mirrored, or RAID 5, for either MBR or GPT disks.
- The vhd file is configured with 4K logical sectors.
- The vhd file or vhd file is not consistent at the time of backup.

A file may not be consistent for the following reasons:

- If an offline backup of a running virtual machine is performed. This may occur if the Hyper-V Integration Services are not installed or if the version of Integration Services running on a virtual machine do not match the version running on a Hyper-V host. This may also occur if the virtual machine runs an operating system that is not supported.
- The vhd or vhd file is a direct parent of an AutoRecovery vhd or vhd file on Hyper-V 2012 R2.
- The vhd or vhd file is the direct parent of a ChildVHD vhd or vhd file on Hyper-V 2012 or later.
- The user-created checkpoints (snapshots) do not generate consistent disks.

Adding a Hyper-V host to the list of servers on the Backup and Restore tab

You can add a Hyper-V host to the list of servers on the **Backup and Restore** tab so that the host server and the virtual machines that it hosts can be selected for backup jobs.

To add a Hyper-V host to the list of servers on the Backup and Restore tab

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Select **Microsoft Hyper-V host**, and then click **Next**.
- 3 Check **Allow Backup Exec to establish a trust with the servers**, and then click **Next**.
- 4 Select **Add a Microsoft Hyper-V server to the list of servers**.
- 5 Type the name or IP address of the Hyper-V host that you want to add, and type an optional description.
- 6 If you want to install the Agent for Windows on all of the host's virtual machines, verify that the option **After adding the virtual host, install the Backup Exec Agent for Windows on the virtual machines** is selected.

If you do not want to install the Agent for Windows on the guest virtual machines, clear the check box.

- 7 Click **Next**.
- 8 Select the correct logon account for the server from the drop-down list.
If the correct logon account is not in the list, click **Add/Edit** to add it to the list.
- 9 Click **Next**.

10 Select any of the following options, and then click **Next**.

Upgrade the Backup Exec Agent for Windows to the current version automatically

Enables Backup Exec to install the most recent version of the Agent for Windows if an older version of the Agent for Windows is already installed on the selected virtual machines.

Restart the remote computer automatically after installing the Backup Exec Agent for Windows when a restart is required

Enables Backup Exec to automatically restart the remote computer, if required.

11 Review the summary information, and then click **Install**.

Viewing details about Hyper-V resources

The **Backup and Restore** tab includes a **Details** pane on the lower portion of the screen. The **Details** pane provides additional details for any type of server that is selected from the list of servers in the upper portion of the screen. Additional details and functionality appear if a Hyper-V host is selected. The **Details** pane for Hyper-V resources includes details about the last 7 days of backup jobs, the date of the last backup, and the date of the next scheduled backup. In addition, it includes the ability to restore data and to filter the list of guest virtual machines.

The **Details** pane for Hyper-V virtual machines includes the resources that may not appear in the list of servers, such as:

- Virtual machines that do not have the Agent for Windows installed.
- Non-Windows virtual machines, such as Linux.
- Virtual machines that do not have a GRT-enabled backup.

The **Details** pane for Hyper-V virtual machines provides a combination of current information and historical information. The **Refresh** button enables Backup Exec to discover all server resources. However, if a virtual machine has been moved, deleted, or has had a credentials change, then the following occurs:

- If the virtual machine has never been backed up, Backup Exec removes it from the **Details** pane.
- If the virtual machine has been backed up, Backup Exec does not remove it from the **Details** pane.

If a virtual machine is renamed, then the following occurs in the **Details** pane:

- If the virtual machine has been backed up, the **Details** pane includes an entry for the new name and keeps the entry for the old name.
- If the virtual machine has never been backed up, the **Details** pane includes only the new virtual machine name.

To view details about Hyper-V resources

- ◆ On the **Backup and Restore** tab, select a Hyper-V host host from the list of servers.

The resources for the selected host appear in the **Details** pane.

Installing the Agent for Windows on Hyper-V virtual machines

To use Backup Exec's Granular Recovery Technology (GRT) for Microsoft application data, install the Agent for Windows on any virtual machines that run Windows.

To install the Agent for Windows on Hyper-V virtual machines

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Select **Microsoft Hyper-V host**, and then click **Next**.
- 3 Check **Allow Backup Exec to establish a trust with the servers**, and then click **Next**.
- 4 Select **Install the Backup Exec Agent for Windows on the virtual machines of a Microsoft Hyper-V host**.
- 5 Select the host server from the drop-down list, and then click **Next**.
- 6 Select the virtual machines that you want to install the Agent for Windows on, and then click **Next**.
- 7 Select the logon account for the virtual machines, and then click **Next**.

- 8 Select any of the following options, and then click **Next**.

Upgrade the Backup Exec Agent for Windows to the current version automatically	Enables Backup Exec to install the most recent version of the Agent for Windows if an older version of the Agent for Windows is already installed on the selected virtual machines.
Restart the remote computer automatically after installing the Backup Exec Agent for Windows when a restart is required	Enables Backup Exec to automatically restart the remote computer, if required.

- 9 Review the summary, and then click **Install**.

Push-installing the Agent for Windows to Hyper-V virtual machines

You use the **Add Server** wizard to push-install the Agent for Windows to the Hyper-V virtual machines that you select. You can choose to install the Agent for Windows on all of the guest virtual machines that need it or on specific virtual machines only. For example, if you know that some of your guest virtual machines are scratch machines and do not need to be backed up, you can exclude those virtual machines from the installation.

To push-install the Agent for Windows to Hyper-V virtual machines

- 1 On the **Backup and Restore** tab, right-click the virtual host, and then select **Install the Agent for Windows to guest virtual machines**.
- 2 Check the option **Allow Backup Exec to establish a trust with the servers**, and then click **Next**.
- 3 Select **Install the Backup Exec Agent for Windows on the guest virtual machines of a Microsoft Hyper-V server**.
- 4 In the **Microsoft Hyper-V server** field, select the Hyper-V server that hosts the virtual machine.
- 5 Click **Next**.
- 6 Select the guest virtual machines on which you want to install the Agent for Windows, and then click **Next**.
- 7 Verify that the virtual machines you selected are online and select the appropriate logon account for those virtual machines, and then click **Next**.
- 8 Review the information on the **Summary** panel, and then click **Install**.

Backing up Microsoft Hyper-V virtual machines

The following backup selections are available for Microsoft Hyper-V:

Table D-2 Microsoft Hyper-V backup selections

Container name	Items in the container	What is included in the backup job
Microsoft Hyper-V	This item includes Initial Store and Virtual Machines .	If you select the Microsoft Hyper-V container for backup, the backup job includes the application configuration settings and all virtual machines.
Initial Store Note: This container does not appear for Hyper-V 2012 and later.	This item includes the virtual server application configuration settings.	If you select Initial Store for backup, the backup job includes a single XML file that contains the Hyper-V authorization configuration.
Host Component Note: This container appears only for Hyper-V 2012 and later.	This item includes the virtual server application configuration settings.	If you select Host Component for backup, the backup job includes multiple files that contain the Hyper-V authorization configuration.

Table D-2 Microsoft Hyper-V backup selections (*continued*)

Container name	Items in the container	What is included in the backup job
Virtual Machines	<p>This item includes each virtual machine that resides on the virtual server.</p> <p>Note: When you select an individual virtual machine, the files that are on that virtual machine appear in the results pane. However, you cannot select individual files to include in or exclude from the backup.</p>	<p>If you select an individual virtual machine, the backup is a full image backup of the entire virtual machine, which includes the following items:</p> <ul style="list-style-type: none">■ vhd files■ avhd files■ Memory snapshot files■ vsv files■ bin files■ vhdx files (Hyper-V 2012)■ avhdx files (Hyper-V 2012)■ The virtual machine's configuration *XML files■ Hyper-V managed snapshots <p>Note: Virtual machines that have remote vhd files are excluded from the backup job. You can use the Agent for Windows and the appropriate Backup Exec agent to protect virtual machines that have remote vhd files.</p> <p>Warning: Backup jobs fail for virtual machines that have pass thru disks. You can install the Agent for Windows and the appropriate Backup Exec agent on virtual machines that have pass thru disks and protect them as if they are physical computers.</p> <p>Note: You may experience a processing time of more than 30 minutes if you back up a virtual machine that runs any supported version of Microsoft SharePoint.</p>

Note: If you want to back up an Exchange database availability group (DAG) on a virtual machine, use the Backup Exec Agent for Exchange. The Agent for Hyper-V does not support backups of Exchange DAG.

To back up Hyper-V virtual machines

- 1 On the **Backup and Restore** tab, right-click the Hyper-V host server or virtual machine that you want to back up.

Note: You can select an individual virtual machine for backup or you can back up a virtual machine by selecting its host server for backup. If you back up a virtual machine through the Hyper-V host, you can use Granular Recovery Technology (GRT) to restore individual files and folders as well as items from Microsoft Exchange, SQL, SharePoint, and Active Directory. If you select an individual virtual machine for backup, you cannot enable the backup to use GRT.

- 2 Select **Backup**, and then select the type of backup that you want to perform.
- 3 On the **Backup Definition Properties** dialog box, in the **Selections** box, click **Edit**.
- 4 On the **Backup Selections** dialog box, check the check boxes for the resources that you want to back up and uncheck the check boxes for the resources that you do not want to back up.
- 5 Click **OK**.
- 6 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**.
- 7 On the **Backup Options** dialog box, in the left pane, select **Virtual Machines**.

8 Set any of the following options for this job.

Use the full backup method for virtual machines that do not support incremental or differential backups	Select this option to enable Backup Exec to run a full backup job if an incremental backup or a differential backup of the virtual machine cannot be performed. Backup Exec may not be able to perform an incremental backup or a differential backup for a number of reasons, such as if the snapshot configuration is altered or the configuration of the host server changed. If this option is not selected and an incremental backup or a differential backup cannot be performed, the job fails.
Back up virtual machines that are powered off	Select this option to enable Backup Exec to back up virtual machines when they are not powered on.
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual files and folders from virtual machines	Select this option to enable individual files and folders to be restored from the full backup. You must install the Agent for Windows on the virtual machine on which you want to restore the data. The Agent for Windows does not have to be installed on the virtual machine to back up the data.
Enable GRT for Microsoft Active Directory objects on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual Active Directory objects on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Active Directory is installed.
Enable GRT for Microsoft Exchange databases and mailbox items on virtual machines	Select this option to enable Backup Exec to collect the information that is required to restore individual Exchange databases and mailbox items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Exchange is installed.

Enable GRT for Microsoft SQL (database-level only) on virtual machines

Select this option to enable Backup Exec to collect the information that is required to restore individual SQL database items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which SQL is installed.

Run an SQL Log backup after backing up the virtual machines

Select this option to enable Backup Exec to back up the SQL log. This option applies to SQL databases that update files in a log instead of in the database file. After the SQL log is backed up, the data from the log files is committed to the database and the log files are emptied. If you do not select this option, the log file continues to grow until it becomes full or until you manually back it up.

Enable GRT for Microsoft SharePoint on virtual machines

Select this option to enable Backup Exec to collect the information that is required to restore SharePoint data on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which SharePoint is installed.

Exclude virtual machines that must be put in a saved state to back up

Select this option to exclude from the backup all offline virtual machines that do not support online backups and that are in a running state when the backup begins.

Backup method

Select this option to change the backup method for the backup jobs that are listed. You can change the names of the backup jobs or add more jobs from the Schedule properties.

- 9 On the **Backup Options** dialog box, click any of the optional settings in the left pane that you want to set for this job.
- 10 Click **OK**.
- 11 On the **Backup Definition Properties** dialog box, click **OK**.

Setting default backup options for Hyper-V

You can use the defaults that Backup Exec sets during installation for all Hyper-V backup jobs, or you can choose your own defaults. You can override the default settings when you create individual jobs.

To set default backup options for Hyper-V

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Job Defaults**, and then select a backup option.

For example, if you want to set up the default options for virtual machine backups to disk, select **Back Up to Disk**. The options that appear will vary depending on what types of storage devices you have configured. Different default options can be configured for backup jobs to different types of storage.

- 3 In the left pane, select **Virtual Machines**.

4 Select the default options that you want to use.

Use the full backup method for virtual machines that do not support incremental or differential backups	Select this option to enable Backup Exec to run a full backup job if an incremental backup or a differential backup of the virtual machine cannot be performed. Backup Exec may not be able to perform an incremental backup or a differential backup for a number of reasons, such as if the snapshot configuration is altered or the configuration of the host server changed. If this option is not selected and an incremental backup or a differential backup cannot be performed, the job fails.
Back up virtual machines that are powered off	Select this option to enable Backup Exec to back up virtual machines when they are not powered on.
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual files and folders from virtual machines	Select this option to enable individual files and folders to be restored from the full backup. You must install the Agent for Windows on the virtual machine on which you want to restore the data. The Agent for Windows does not have to be installed on the virtual machine to back up the data.
Enable GRT for Microsoft Active Directory objects on virtual machines	Enables Backup Exec to collect the information that is required to restore individual Active Directory objects on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Active Directory is installed.
Enable GRT for Microsoft Exchange databases and mailbox items on virtual machines	Enables Backup Exec to collect the information that is required to restore individual Exchange databases and mailbox items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which Exchange is installed.

Enable GRT for Microsoft SQL (database-level only) on virtual machines

Enables Backup Exec to collect the information that is required to restore individual SQL database items on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which SQL is installed.

Run an SQL Log backup after backing up the virtual machines

Enables Backup Exec to back up the SQL log. This option applies to SQL databases that update files in a log instead of in the database file. After the SQL log is backed up, the data from the log files is committed to the database and the log files are emptied. If you do not select this option, the log file continues to grow until it becomes full or until you manually back it up.

Enable GRT for Microsoft SharePoint on virtual machines

Enables Backup Exec to collect the information that is required to restore SharePoint data on the virtual machine. Backup Exec uses the logon credentials that were used for the virtual machine on which SharePoint is installed.

Exclude virtual machines that must be put in a saved state to back up

Excludes from the backup all offline virtual machines that do not support online backups and that are in a running state when the backup begins.

Backup method

Enables you to change the backup method for the backup jobs that are listed. You can change the names of the backup jobs or add more jobs from the Schedule properties.

5 Click **OK**.

How Backup Exec automatically protects new virtual machines during a backup job

Backup Exec's dynamic inclusion feature protects new virtual machines and folders that are found when a backup job runs. If new virtual machines are added between the time when the backup job is created and when the backup job runs, Backup Exec automatically backs up the new virtual machines. Because the backup job may include new virtual machines, the job may require more storage space and

more time to run than you anticipated. The job history shows the number of virtual machines that were backed up.

In the backup selection list, dynamic inclusion is enabled for the following Hyper-V nodes:

- Microsoft Hyper-V
- Virtual Machines under Microsoft Hyper-V
- The Hyper-V host node
If you select the host node, then dynamic inclusion is enabled automatically for the Microsoft Hyper-V node.
- Microsoft Hyper-V HA Virtual Machines
- The cluster name node
If you select the cluster name node, then dynamic inclusion is enabled automatically for the Microsoft Hyper-V HA Virtual Machines node.

See [“Backing up Microsoft Hyper-V virtual machines”](#) on page 873.

Using Granular Recovery Technology (GRT) with the Agent for Hyper-V

Backup Exec Granular Recovery Technology (GRT) lets you restore individual files and folders without having to restore the entire virtual machine. It also lets you restore individual items from the following VSS-aware applications that are installed on virtual machines.

Table D-3 Types of data that Backup Exec protects for VSS-aware applications on virtual machines

Application	Types of data that Backup Exec protects
Microsoft Exchange	Mailboxes, individual messages, calendar items, tasks, journal entries, and public folder data (disk-backups only)
Microsoft SQL	Databases
Microsoft Active Directory	Individual user accounts, printer objects, sites, and organizational units
Microsoft SharePoint	SharePoint databases

Note: GRT is not intended to be used for system recovery. However, you can perform a complete system recovery by selecting the entire virtual machine as a restore selection in a restore job.

When you create a backup job, Backup Exec automatically locates VSS-aware applications on virtual machines. During the backup job, Backup Exec backs up the data from the VSS-aware applications by using GRT. By default, Backup Exec enables GRT using the same credentials that were used to connect to the virtual machine. You can disable GRT for any of the VSS-aware application types.

To use GRT, you must select the individual files and folders that you want to restore from the list that appears when you expand the Netbios name or the computer name of the virtual machine. You cannot select individual folders and files from the virtual machines that appear when you expand the **Virtual Machines** node.

Note: Backup Exec supports the granular recovery of individual Exchange and SQL items only in non-clustered and non-distributed configurations.

During the backup job, Backup Exec collects metadata for the applications. If Backup Exec is unable to collect the metadata, then you cannot restore individual items for the applications. However, the backup job may otherwise complete successfully.

Backup Exec cannot collect metadata in the following situations:

- GRT is disabled for an application.
- Backup Exec cannot connect to the virtual machine.
- Incorrect credentials were entered for the virtual machine.

Note: Backup Exec uses the Microsoft Hyper-V writer during backups of VSS-aware applications on virtual machines. The Microsoft Hyper-V writer truncates application logs before data is moved to the storage device. Therefore, the application logs for the applications on the virtual machines are truncated if you use Microsoft Hyper-V.

GRT is not supported in the following situations:

- Restores from mixed media if GRT was enabled in the backup job. For example, if the full backup is on tape and the incremental backup is on a disk storage device, the restore job fails. Restores from mixed media types are supported if GRT is not enabled.
- Backups to tape of an Exchange 2003 virtual machine that runs on a Hyper-V host that is installed as a server role on a Server Core installation of Windows Server 2008 or later. However, backup to a disk storage folder is supported.

Requirements for using GRT to back up Microsoft application data on virtual machines

The following items are required to protect data for Microsoft Exchange, SQL, Active Directory, and SharePoint on virtual machines:

- The virtual machine must be turned on.
- You must enter the appropriate credentials for the virtual machine. Ensure that the credentials for the virtual machine allow access to the VSS-aware applications.
- The Backup Exec server must be able to connect to the virtual machine using the network name or IP address.
- The Backup Exec Agent for Windows must be installed on the virtual machine.
- The correct number of licenses must be entered for the applications that you want to protect on the virtual machines.
- The operating system on the virtual machine must support VSS.
- The virtual machine cannot use dynamic disks, such as spanned, mirrored, striped, or RAID 5.

About backing up and restoring highly available Hyper-V virtual machines

When virtual machines are configured for high availability, they appear in the **Highly Available Hyper-V Machines** node in the backup selection list. Virtual machines that are not configured for high availability remain in the **Microsoft Hyper-V** node. When you make a backup selection, Backup Exec checks for highly available virtual machines. If highly available virtual machines are discovered, Backup Exec reminds you to select those virtual machines for backup.

The restore selections are similar to the backup selections. You can restore a highly available virtual machine in the same way you restore any other virtual machine. The virtual machine maintains its high availability. However, if you redirect the restore to another Hyper-V host, then the virtual machine is no longer highly available when the restore job completes. You must reconfigure the virtual machine to be highly available.

See [“Backing up Microsoft Hyper-V virtual machines”](#) on page 873.

How delayed cataloging works with Hyper-V virtual machine backups

When you enable Granular Recovery Technology (GRT) for a backup job of a virtual machine, you can choose to run the catalog job for GRT as part of the backup job,

as a separate job immediately after the backup job completes, or according to a schedule. By default, the catalog operation runs immediately after the backup job completes.

Note: The delayed catalog feature is not supported for backups to tape.

The catalog operation can be time consuming. It requires access to the storage device that is used for the backup. You may want to schedule the catalog operation to run outside of your backup window so that it does not interfere with backup jobs. If the catalog operation is scheduled, it runs only for the most recent backup set since the last catalog operation. In this situation, only the most recent backup set since the last catalog operation can be used for granular recovery on Hyper-V virtual machines. GRT data is not available until the delayed catalog job completes, but a full restore of a virtual machine can be performed before the delayed catalog job completes.

For example, if you set up incremental backups to run every 11 hours and set up the catalog operation to run at midnight, you would have the following backup sets:

- Full (11:00 A.M.)
- Incremental 1 (10:00 P.M.)
- Catalog 1 (Midnight). This job catalogs Incremental 1.
- Incremental 2 (9:00 A.M.)
- Incremental 3 (8:00 P.M.)
- Catalog 2 (Midnight). This job catalogs Incremental 3. Incremental 2 is not cataloged.
- Incremental 4 (7:00 A.M.)
- Incremental 5 (6:00 P.M.)
- Catalog 3 (Midnight). This job catalogs Incremental 5. Incremental 4 is not cataloged.
- Incremental 6 (5:00 A.M.). This backup is not cataloged.

In the example above, you can perform a granular recovery from Incremental 5, Incremental 3, or Incremental 1. You can perform a full restore from any of the above backup sets.

How byte count is calculated for delayed catalog operations

On the **Job Monitor** and **Job History**, the byte count that displays for a catalog operation may differ from the byte count that displays for the corresponding backup job. The byte count for a catalog job may be larger than the byte count for a backup

job. The way in which Backup Exec catalogs the data affects the byte count that appears for the catalog job.

- When a catalog operation is performed for a full backup, the data is read on a file-by-file basis and the byte count is calculated accordingly. During the full backup job, the data is read in terms of the number of sectors and the byte count is calculated based on the number of sectors. Therefore, the byte count for the catalog job may be larger than the byte count for the backup job.
- When the catalog operation is performed for an incremental backup, all files on the virtual disk are cataloged instead of only the changed files. Therefore, the byte count for the catalog job takes into account both the full backup and the incremental backup.

See [“Configuring delayed catalog operations to improve backup performance”](#) on page 536.

Restoring Microsoft Hyper-V virtual machines

You can use the Restore Wizard to restore data from virtual machines in the following ways:

- Restore a complete virtual machine for disaster recovery purposes.
- Restore individual files or folders that were backed up from the virtual machine (if you selected the Granular Recovery Technology options for the backup job).
- Restore a virtual machine to a different Microsoft Hyper-V server.
- Redirect flat files from the virtual machine to any computer that has an Agent for Windows installed.

Notes about restoring Hyper-V virtual machines

- Linux virtual machines must be restored in their entirety at the vhd level.
- The restore of a Microsoft Hyper-V virtual machine that is created within a mount point fails if the mount point does not exist at the time of the restore. To avoid this issue, configure the virtual machine to use volume GUID paths with no mount points.
- Restores of virtual machines with pass-thru disks, fibre channel adapters, and shared vhdx files may fail if those items have been deleted or those items are unavailable at the time of restore. The restore job will succeed if you have not deleted the pass-thru disks, fibre channel adapter, or shared vhdx file.

Note: The steps below apply to virtual machines that were backed up through the virtual host with the Agent for Hyper-V. If you did not back up the virtual machine through the virtual host with the Agent for Hyper-V, follow the steps for restoring a non-virtual backup

See [“Methods for restoring data in Backup Exec”](#) on page 232.

To restore Hyper-V virtual machines

- 1 On the **Backup and Restore** tab, do one of the following.

To restore individual files and folders from a GRT-enabled backup Do the following in the order listed:

- In the **Details** pane at the bottom of the screen, select the virtual machine.
- Click **Restore**, and then select **Restore GRT-enabled backup**.
- In the **Restore Wizard**, select **Files, folders, or volumes**, and then click **Next**.

To restore the entire virtual machine or virtual disks

Do the following in the order listed:

- In the **Details** pane at the bottom of the screen, select the virtual machine.
- Click **Restore**, and then select **Restore virtual machine from the host**.
- In the **Restore Wizard**, select **Hyper-V data**, and then click **Next**.

- 2 Select the data that you want to restore, and then click **Next**.

Note: For virtual machines that have a System Reserved partition and are backed up with Granular Recovery Technology enabled, Backup Exec displays the virtual machine under a volume GUID instead of a drive letter. To restore data from these virtual machines, select one of the options to restore to another location.

3 Select the location where you want to restore the data, and then click **Next**.

To the original location

Select this option to restore the virtual machine to the same location from which it was backed up.

To another location; keep the same drive and path

Select this option to restore the virtual machine to a different server, but use the same drive and path name that were used in the original. You must enter the name of the server to which you want to redirect the restore, and you must also enter the logon account for that server

Example: The original virtual machine was on \\ServerA\D:\VMs\1.vhd and you want to restore it to \\ServerB\D:\VMs\1.vhd

To another location; change the drive, but keep the same path

Select this option to restore the virtual machine to a different server and change the drive, but use the same path name that was used in the original. You must enter the name of the server and the drive to which you want to redirect the restore. You must also enter the logon account for that server.

Example: The original virtual machine was on \\ServerA\D:\VMs\1.vhd and you want to restore it to \\ServerA\E:\VMs\1.vhd

To another location; change the drive and path, but include the original drive and path name in the new path

Select this option to restore the virtual machine to a different server, change the drive, and include the original drive and path names in the new path. You must enter the name of the server, the drive, and the path to which you want to redirect the restore. You must also enter the logon account for that server.

Example: The original virtual machines were on \\ServerA\D:\VMs\1.vhd and \\ServerA\E:\VMs\2.vhd, and you want to restore them to \\ServerB\Z:\ReplicatedVMs\D\VMs\1.vhd and \\ServerB\Z:\ReplicatedVMs\E\VMs\2.vhd

- 4 Select the additional options that you want to use for this restore job, and then click **Next**.

Overwrite powered on virtual machines Select this option to enable powered on virtual machines to be overwritten and restored. By default, virtual machines must be turned off before the restore job processes and the virtual machine is overwritten. If a virtual machine is running during a restore job, but this option is not selected, the job fails. You must manually turn off the virtual machine before you attempt to run the restore job again.

Do not overwrite powered on virtual machines Select this option to prevent powered on virtual machines from being overwritten and restored. By default, virtual machines are turned off before the restore job processes and the virtual machine is overwritten. If a virtual machine is running during a restore job, and this option is selected, the job fails. You must manually turn off the virtual machine before you attempt to run the restore job again.

Power on the virtual machine after restore and resume from the available saved state Select this option to turn on the virtual machine automatically after the restore job completes. The virtual machine resumes operations from the saved state from the time of the backup.

Note: This option applies only to virtual machines that are backed up using a saved state. Virtual machines that are backed up online do not have a saved state.

Power on the virtual machine after restore and discard the available saved state Select this option to turn on the virtual machine automatically after the restore job completes. The virtual machine discards the available saved state.

Note: This option applies only to virtual machines that are backed up using a saved state. Virtual machines that are backed up online do not have a saved state.

- 5 Enter a name for this restore job and select the schedule for the job, and then click **Next**.
- 6 Review the job summary, and then click **Finish**.

Symantec Backup Exec Agent for Microsoft SQL Server

This appendix includes the following topics:

- [About the Agent for Microsoft SQL Server](#)
- [Requirements for using the SQL Agent](#)
- [About installing the SQL Agent](#)
- [Backup strategies for SQL](#)
- [Adding SQL Servers to the list of servers on the Backup and Restore tab](#)
- [Configuring Backup Exec to run a consistency check before every SQL backup](#)
- [Using snapshot technology with the SQL Agent](#)
- [Using database snapshots for SQL Server](#)
- [Backing up SQL databases and transaction logs](#)
- [Restoring SQL databases and transaction logs](#)
- [Disaster recovery of a SQL Server](#)

About the Agent for Microsoft SQL Server

The Agent for Microsoft SQL Server (SQL Agent) enables network administrators to perform backup and restore operations on installations of SQL that are connected

to a network. SQL database backups can be integrated with network backups without separate administration or dedicated hardware.

The SQL Agent provides support for the following:

- Database, transaction log, and differential backups, as well as database recovery and replacement.
- An automated restore of the system databases.
- Simplified Disaster Recovery, which automates the disaster recovery process of SQL Servers.
- Restores of SQL databases to alternate locations.
- Hot backup copies of SQL databases during backup operations. This feature enables you to direct a copy of the actual data streams being sent to media by a SQL database to a local directory for later use.
- Backups of multiple instances.
- Standby database. If the primary SQL Server fails, or is shut down for maintenance, another database called a standby database can be brought online.
- Database Consistency Checks (DBCC) for each backup and restore job, including a fast database consistency check of only the physical consistency of the database.
- Full, bulk-logged, and simple recovery models. With the simple recovery model, copies of the transactions are not stored in the log file, which prevents transaction log backups from being run. Therefore, you can recover the database to the point of the last backup, but you cannot restore the database to the point of failure or to a specific point in time.
- Restores of transaction logs to a specific point in time or to a named transaction when log marks are used.
- Database snapshots.
- Maintenance replication settings during redirected restores.
- A Verify Only option for a restore job that determines both the validity of the SQL data on the media and the ability of the destination SQL database to accept this data before the database is deleted or overwritten during a restore job.
- Back up with checksum generation. This option is used as a redundancy check and works with the Verify Only option on a restore job.
- Continuation of restore jobs when errors are detected. This feature enables you to restore as much data as possible from a corrupt database backup.

- Copy-only one-time backups, which enable you to copy a database without affecting the full-differential-log restore sequence.
- In SQL Server 2008 or later editions that support compression, you can use SQL software compression for backup jobs.

See [“About installing the SQL Agent”](#) on page 893.

Requirements for using the SQL Agent

The following are required for the SQL Agent:

- Backup Exec must have access rights to read both of the following SQL registry keys:

- HKEY_LOCAL_MACHINE\Software\Microsoft\Microsoft SQL Server
- HKEY_LOCAL_MACHINE\Software\Microsoft\mssqlserver

If Backup Exec does not have access to these registry keys, a restore to the default directory may not work, and the Automate system database restore option on the Restore Job Properties for SQL dialog box will not work.

To ensure that Backup Exec has access rights, verify that the logon account used has administrator rights to the Windows server that the SQL instance is installed on.

- The Backup Exec server must have access to the SQL installation.
- The credentials stored in the Backup Exec logon account used for backing up and restoring SQL must have been granted the System Administrator role on the SQL instance.

When you test the logon account, apply the logon account to the Windows server that SQL is installed on, not to the actual SQL instance.

Note: SQL Server credentials are not supported.

- The Agent for Windows must be installed on any remote SQL Server that you want to back up.

For specific operating system requirements and supported SQL Server service packs for the SQL Agent, see the Backup Exec Software Compatibility List at the following URL:

<http://entsupport.symantec.com/umi/v-269-1>

See [“Testing logon accounts”](#) on page 620.

See [“Backup Exec logon accounts”](#) on page 610.

About installing the SQL Agent

The SQL Agent is installed as part of the Agent for Applications and Databases and can protect local or remote SQL Server databases.

See [“Installing additional Backup Exec options to the local Backup Exec server”](#) on page 76.

Backup strategies for SQL

Backup Exec incorporates online, nondisruptive SQL database protection as part of everyday backup routines, which increases the chance of data recovery and minimizes data loss without inhibiting daily database activity. Using database, differential, and log backups provides a good balance between backup windows and minimizes the amount of time that will be spent recovering a database if the need arises.

To decide which backup methods to use for the best data protection, consider the following for typical environments:

- In small environments, consider running a daily full database backup every evening and daily transaction log backups.
- In mid-sized environments, consider running a weekly full database backup and daily transaction log backups along with daily differential backups except on the day when the full backup is run.
- In large environments, consider running daily differential database backups, weekly full database backups, and transaction log backups as necessary. Many shops run full backups on a weekly basis, preferring to run differential backups throughout the week to keep backup run time to a minimum.

The trade-off with running fewer full backups and running more differential backups occurs at recovery time when you must recover using the full database backup as well as the last differential database backup, and all log backups made after the last differential database backup.

What will work best for you will be based on the size of your environment, the number of transactions processed each day, and the expectations of your users when a recovery is required.

When you develop a SQL backup strategy, consider the following:

Table E-1 Recommendations for backing up SQL

SQL Server backup strategies	Description
Protect the entire SQL Server.	To ensure SQL is completely protected, back up the following on a regular basis: <ul style="list-style-type: none">■ The system drive that SQL is on.■ The Windows registry and System State.■ Transaction logs.
When you upgrade, run new full database backups.	If you upgrade SQL, run new full database backups. You may not be able to restore backups from one version or service pack level of SQL to other versions.
Run consistency checks before backups.	Symantec recommends that you run a consistency check before a backup. If a database or transaction log contains errors when it is backed up, the backup will still contain the errors when it is restored, if it is restorable at all. See “Configuring Backup Exec to run a consistency check before every SQL backup” on page 895.
Back up your system databases regularly.	Back up the master database and service packs that are installed whenever procedures are run that change information in the database, especially after the following: <ul style="list-style-type: none">■ New databases are created.■ Files are added to an existing database.■ Usernames or passwords are added or changed. If changes are not backed up before the master database must be restored, the changes are lost.
Run one backup at a time.	Do not schedule more than one backup to occur simultaneously against a database or its transaction log.
Back up transaction logs on databases that are configured for the full recovery model.	Back up the transaction logs on databases because the transaction logs continue to grow if you do not back them up.

Adding SQL Servers to the list of servers on the Backup and Restore tab

You can add SQL Servers to the list of servers on the **Backup and Restore** tab so that you can select the SQL databases for backup.

To add a SQL Server to the list of servers on the Backup and Restore tab

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Select **Microsoft Windows computers and servers**, and then click **Next**.
- 3 Follow the **Add a Server** wizard prompts to add the SQL Server to the list of servers in the **Backup and Restore** tab.

See [“About the list of servers on the Backup and Restore tab”](#) on page 152.

Configuring Backup Exec to run a consistency check before every SQL backup

If you back up a database or transaction log that contains errors, these errors will still exist when the backup is restored. In some cases, this can prevent a successful restore. Backup Exec enables you to check the logical and physical consistency of the data before and after a backup. SQL reports any consistency check failures in the Backup Exec job log. Symantec strongly recommends that you always run a consistency check before the backup.

By default, the backup job default option **Consistency check before backup** is enabled for **Physical check only**.

Backup Exec's consistency check uses the following SQL consistency check utilities:

- CHECKDB
- CHECKCATALOG
- PHYSICAL_ONLY

CHECKDB, CHECKCATALOG, and PHYSICAL_ONLY are performed for database-related operations.

For more information concerning these utilities, see your Microsoft SQL documentation.

To run a consistency check before a SQL backup

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Job Defaults**.
- 2 Select a backup job type.

For example, if you want to set up the default options for SQL Server backups to disk, select Back Up to Disk. The options that appear vary depending on what types of storage devices that you configure. You can configure different default options for backup jobs that you send to different types of storage.
- 3 On the **Backup Job Defaults** dialog box, on the left pane, click **Microsoft SQL**.
- 4 Ensure that a consistency check is enabled in the field **Consistency check before backup**.
- 5 Click **OK**.

See [“Backing up SQL databases and transaction logs”](#) on page 899.

Using snapshot technology with the SQL Agent

Backup Exec uses snapshot technology by default for SQL Server backups. The SQL Agent supports full snapshot backups using Microsoft’s Volume Shadow Copy Service (VSS), a snapshot provider service only available on Windows 2003 or later. The use of snapshot technology can reduce both restore time and backup performance on the server.

When you submit a backup job that uses snapshot technology, a snapshot of each volume is created, providing a point-in-time record of the data. Backup Exec uses snapshot technology to momentarily suspend write activity to a volume so that a snapshot of the volume can be created. The data is then backed up from the snapshots, and then the snapshots are deleted.

Note: Use snapshot technology with jobs that use deduplication devices.

Before you use snapshot technology with the SQL Agent, review the following information:

- With snapshot technology, a point-in-time view of the SQL database is “snapped” and then backed up, leaving the actual SQL database open and available for users.
- SQL backups that use snapshot technology are considerably larger than regular SQL backups (also known as streaming backups).

- Performing consistency checks before backup is strongly recommended. See [“Configuring Backup Exec to run a consistency check before every SQL backup”](#) on page 895.
- The SQL Agent supports only full snapshot backups; transaction log snapshots and differential snapshots are not supported.
- With the SQL Agent, snapshot and streaming backups are interoperable when you restore SQL data.
- Performing database consistency checks both before and after backups affects the time required for the backup jobs.

The following SQL backup options are not supported with snapshot backups:

- **Use checksums on backup (SQL 2005 or later)**
This option is used as a redundancy check, and works with the **Run verify only and do not restore data** restore option.
- **SQL Server 2008 Enterprise Edition software compression**
- **Create on-disk copies of SQL backups to be placed on the SQL Server where the database is located**

Note: Microsoft Virtual Shadow Copy Service (VSS) snapshots are not the same as SQL database snapshots. VSS snapshots enable you to create point-in-time snapshots of disk volumes and shares; database snapshots enable you to create point-in-time copies of SQL databases.

See [“Using database snapshots for SQL Server”](#) on page 897.

Using database snapshots for SQL Server

SQL database snapshots enable you to quickly revert a database back to the state it was in when the database snapshot was created. When you use a database snapshot, a full restore of the host database is not required to revert the database. However, the changes that are made to the host between the time a database snapshot is created and the point at which it is reverted are lost.

The Backup Exec SQL Agent works with the SQL database to create database snapshots, which are read-only, point-in-time copies of an existing host database. When Backup Exec runs a SQL backup job using the Database Snapshot backup method, a request is sent to the host database instructing it to create a database snapshot.

Note: The snapshot backup method for SQL databases is only supported by SQL Server Enterprise Edition.

Database snapshots cannot be backed up to storage media. Rather, they are written to a SQL snapshot file on disk. After running the database snapshot job, Backup Exec creates history and job log information to indicate the job's status.

Because database snapshots cannot be backed up, all database snapshots will be lost if the disk where the host database is installed fails. Therefore, database snapshots should not be used as your sole database protection strategy. They should be used in conjunction with an overall Backup Exec database protection strategy that includes full, differential, and transaction log backups of the SQL database.

For more information, see your Microsoft SQL documentation.

Note: SQL database snapshots are not the same as Microsoft Virtual Shadow Copy Service (VSS) snapshots. Whereas VSS snapshots enable you to create point-in-time snapshots of disk volumes and shares, database snapshots enable you to create point-in-time copies of SQL databases.

Note: SQL database snapshot catalog information that refers to deleted database snapshots is periodically removed from the catalogs. If backup media is re-cataloged, the database snapshot catalog information will be periodically removed again.

SQL Server database snapshots created with Backup Exec can be used to revert a SQL database back to the state it was in at a previous point in time, without having to run a full database restore job.

When you view SQL database snapshots in the resource view in the Restore Wizard, they appear as backup sets, in chronological order with the most recent snapshot appearing first.

The following caveats apply when reverting a database:

- You cannot undo a SQL database that has been reverted.
- Before reverting a database, Backup Exec deletes all existing database snapshots with the exception of the snapshot used for the revert. After being deleted, the database snapshots cannot be recovered.
- You cannot redirect a database snapshot restore job.

Backing up SQL databases and transaction logs

Backup Exec includes three methods for backing up databases: Full, Differential, and Full Copy-only. The full method backs up the entire database including all system tables. The differential method backs up only the changes made to the database since the last full backup. The copy method works in the same manner as the full method, except that it does not affect future differential or log backups.

A differential backup is smaller and faster than a full backup, so differential backups can be run more often than full backups. Because differential backups allow the restore of a system only to the point that the differential backup was created, you should also create multiple log backups between the differential backups. Using transaction log backups allows you to recover the database to the exact point of failure.

Consider using differential backups when only a relatively small amount of data changes between full backups, or if the same data changes often. Differential backups may also work well in your environment if you are using the simple recovery model and need backups more often, but cannot spare the time to do frequent full backups. If you are using the full or bulk-logged recovery models, you can use differential backups to decrease the time it takes to roll forward log backups when restoring a database.

If you want to run database backups only, instead of a mix of database and log backups, use the simple recovery model for the database so that the transaction log is automatically truncated when a checkpoint occurs in the database. This helps prevent transaction logs from becoming full since with other recovery models the logs are not cleared after a database backup.

With the simple recovery model, copies of the transactions are not stored in the log file, which prevents transaction log backups from being run.

If you do not run transaction log backups, you can recover the database to the point of the last backup, but you cannot restore the database to the point of failure or to a specific point in time.

System databases can only be backed up with the full method; you cannot use the log or differential methods to back up the master database.

Note: You cannot back up databases to storage that is attached to a computer on which the Remote Media Agent for Linux Servers is installed.

The SQL Agent supports a mirrored SQL database configuration, although Microsoft places the following limitations on the mirroring of SQL databases:

- You cannot back up or restore a mirrored SQL database. If you attempt to back up or restore a mirrored database, the backup job or restore job fails.
- You cannot restore the primary SQL database while it is configured in a mirrored configuration. To restore the primary SQL database, you must stop database mirroring of the primary database.
- You can back up a primary SQL database and its transaction logs only if the backup job does not leave the database in a non-recovered state.

You can set backup job default options for all SQL backup jobs. Each time you create a backup job, the job uses the default options unless you change the options for that particular job.

Automatic exclusion of SQL data during volume-level backup

If you select a volume that contains SQL data for backup, the SQL Agent determines which SQL data should not be included in a volume-level backup. For example, .MDF and .LDF files should not be part of the backup because they are opened for exclusive use by the SQL system. These files will be automatically excluded for backup by a feature called Active File Exclusion. If this exclusion did not happen during a non-snapshot backup, these files would appear as in use - skipped. If this exclusion did not happen during a snapshot backup, the files would be backed up in a possible inconsistent state, which could create restore issues.

While it is not recommended, if you want to include SQL data in a volume-level backup, you must first dismount the database that you want to back up. Then, run the backup job.

How to back up SQL transaction logs

When you run log backups, you should use Backup Exec exclusively to perform log transaction backups.

Backup Exec provides the Log and Log No Truncate methods for backing up transaction logs.

Use the Log No Truncate method only when the database is corrupted or if database files are missing. This method backs up the transactions that you may not be able to access otherwise when the database is in this state. You can then use this transaction log backup along with the last database backup and any previous transaction log backups to restore the database to the point at which it failed. However, any uncommitted transactions are rolled back. The Log No Truncate method does not remove committed transactions after the log is backed up.

To use the Log No Truncate backup to restore a database, you should also have a database backup that was created before the Log No Truncate backup. The transaction log contains only the log files that are used in the restore process, which

alone are not sufficient to restore a complete database. You must have at least one database backup and a log backup of the database to restore the database.

Caution: Do not run a log backup using either method if the SQL database uses the simple recovery model. With the simple recovery model, you can recover data only up to the most recent full or differential backup. If you run a log backup on a database using the simple recovery completion state, the backup completes with exceptions.

To check the database properties, from the Database management tools on the SQL Server, right-click the database, click **Properties**, click the **Options** tab, and then view the configuration settings.

To back up SQL databases and transaction logs

- 1 On the **Backup and Restore** tab, right-click a SQL Server, and then right-click the selection.

To select multiple servers, Shift + click or Ctrl + click the server names, and then right-click one of the selected servers.
- 2 Select **Backup**, and then select the type of backup that you want to perform.
- 3 On the **Backup Definition Properties** dialog box, in the **Selections** box, click **Edit**.
- 4 On the **Backup Selections** dialog box, check the check boxes for the resources that you want to back up and uncheck the check boxes for the resources that you do not want to back up.

Note: You can select the SQL databases to back up on the **Browse** tab. In the right pane of the **Backup Selections** dialog box, you can view the Name, Size, Type, Modified Time, and Attributes for the selection. The Attributes provide the status of the database, so if there are any problems you can resolve them before you run the backup job. You can also include or exclude specific files or specific types of files using the **Selection Details** tab.

- 5 Click **OK**.
- 6 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**.
- 7 On the **Backup Options** dialog box, select the schedule for this job.
- 8 On the **Backup Options** dialog box, select the storage device for this job.
- 9 On the **Backup Options** dialog box, in the left pane, select **Microsoft SQL**.

10 Set any of the following options for this job:

Backup method

Select one of the following backup methods that you want to use for this job:

- **Full - Back up databases**
This option backs up the entire database. This option is selected by default.
- **Full Copy - Back up databases (copy)**
This option backs up the entire database without affecting future differential or log backups.
Unlike the Full backup method, the Full Copy backup method does not reset the SQL differential baseline that is used to indicate the database blocks that have changed since the last full backup.
After making a full backup, you can use the Full Copy backup method to make a copy of a SQL database without affecting the baseline backup set required to run future differential backups.

Backup method

Select one of the following SQL-specific backup methods that you want to use for this job:

- **Full - Back up databases**
This option backs up the entire database. This option is selected by default.
- **Full Copy - Back up databases (copy)**
This option backs up the entire database without affecting future differential or log backups.
Unlike the Full backup method, the Full Copy backup method does not reset the SQL differential baseline that is used to indicate the database blocks that have changed since the last full backup.
After making a full backup, you can use the Full Copy backup method to make a copy of a SQL database without affecting the baseline backup set that is required to run future differential backups.
- **Automatic - Backup up transaction log if enabled and then back up database changes since the last full or incremental**
This option lets you back up the entire SQL instance even though some databases may not support log backups. All of the databases are backed up using the Incremental (block level) backup method. In addition, the databases that support log backups are backed up using the Log backup method.
Note: If snapshot is not enabled, an Incremental (block level) backup method cannot be performed and the Differential backup method is used.
- **Log - Back up and truncate transaction log**
This option backs up only the data contained in the transaction log; it does not back up database data. After the transaction log is backed up, committed transactions are removed (truncated).
If the databases are configured for the SQL Server simple recovery model, log backups are not supported. To change the recovery model, use the SQL administration tools to set the recovery model to Full. You should run a new full backup if you change the recovery mode before a log backup is run.
Alternatively, you can run full backups only, or run full and differential backups of the SQL databases.
See [“Configuring Backup Exec to run a consistency check before every SQL backup”](#) on page 895.
- **Log No Truncate - Back up without truncating**

transaction log

This option backs up the database when it is corrupt or database files are missing. Since the Log No Truncate backup method does not access the database, you can still back up the transactions that you may not be able to access otherwise when the database is in this state. You can then use this transaction log backup along with the database backup and any previous transaction log backups to restore the database to the point at which it failed; however, any uncommitted transactions are rolled back. The Log No Truncate backup method does not remove committed transactions after the log is backed up.

- Differential - Backup up database changes since the last full

This option backs up only the changes made to the database or filegroup since the last full backup. Because differential backups allow the restore of a system only to the point in time that the differential backup was created, you should also create multiple log backups between the differential backups.

- Differential (block-level) - Back up database changes since the last full - use with convert to virtual machine job

This option backs up all of the blocks of data and logs that have been created or modified since the last full backup.

- Incremental (block-level) - Back up database changes since the last full or incremental - use with convert to virtual machine job

This option backs up all of the blocks of data and logs that have been created or modified since the last full or incremental backup.

- Database Snapshot - Read-only point-in-time copy of databases

This option creates a read only, point-in-time copy of another database.

See ["Using database snapshots for SQL Server"](#) on page 897.

Note: SQL differential or incremental backups are only supported for conversion to virtual when you use the Automatic, Differential (block-level), or Incremental (block-level) backup methods.

Additionally, Backup Exec runs a full backup when you select the Automatic or Log backup methods if a full backup was not previously run on the database. A full backup also runs for one of the following conditions:

- A new database is added or restored.
- Backup Exec did not run the last full backup.
- Only Full Copy and Incremental backups were run on the database instead of Full backups.

See [“Configuring backup methods for backup jobs”](#) on page 196.

Consistency check before backup

Select one of the following consistency checks to run before a backup:

- None.
This option does not run a consistency check before a backup. Symantec recommends that you always run a consistency check either before or after the backup.
- Full check, excluding indexes.
This option excludes indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough.
- Full check, including indexes.
This option includes indexes in the consistency check. Any errors are logged.
- Physical check only.
This option performs a low-overhead check of the physical consistency of the database. This option only checks the integrity of the physical structure of the page. This option is selected by default.

See [“Configuring Backup Exec to run a consistency check before every SQL backup”](#) on page 895.

Continue with backup if consistency check fails

Choose if you want to continue with the backup operation even if the consistency check fails. You may want to continue with the backup when the consistency check fails if you think that a backup of the database in its current state is better than no backup at all, or if you are backing up a very large database with only a small problem in a table.

Consistency check after backup

Select the consistency check to run after a backup. Because database transactions can occur during or after the consistency check, but before the backup runs, consider running a consistency check after the backup to ensure that the data was consistent at the time of the backup.

The following options are available:

- **None.**
This option does not run a consistency check after a backup. Symantec recommends that you always run a consistency check after the backup. This option is selected by default.
- **Full check, excluding indexes.**
This option excludes indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough.
- **Full check, including indexes.**
This option includes indexes in the consistency check. Any errors are logged.
- **Physical check only.**
This option performs a low overhead check of the physical consistency of the database. This option only checks the integrity of the physical structure of the page. This option is selected by default.

Use checksums on backup (SQL 2005 or later)

Choose to add the checksums to the SQL database data being backed up by Backup Exec. Adding checksums to the data being backed up is required if you want to use the restore option **Run verify only and do not restore data**. Use this option and the **Run verify only and do not restore data** option to ensure that you restore from a verified SQL backup during a restore of the SQL database.

Create on-disk copies of SQL backups to be placed on the SQL server where the database is located

Choose to create an on-disk copy of the SQL database that you want to back up. This option lets you simultaneously back up a SQL database to storage media while also writing a copy of the database to a disk path you specify in the **Save to path** option.

This option gives IT administrators the ability to back up SQL databases while also providing database administrators with copies of the database on disk, which can be used for such things as tests and restores.

Note: This option does not support snapshot technology.

Save to path

Specify a path in which to save on-disk copies of SQL backups.

**SQL Server 2008
Enterprise Edition
software compression**

Select a compression setting that you want to use for this backup job:

- None.
Do not use compression.
- Compress.
Use SQL Server 2008 or later compression if it is supported by the SQL Server instance that is installed.

SQL compresses the data on the computer on which SQL Server 2008 Enterprise Edition or later is installed. Therefore, faster SQL 2008 or later backups should occur if you use SQL compression.

You can find a list of compatible operating systems, platforms, applications, and supported service packs at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Symantec recommends that you do not use SQL 2008 or later software compression in a backup job that also uses Backup Exec-initiated software compression. Minimal benefits are gained when you enable Backup Exec compression. In fact, in jobs where both compression schemes are used, backup times may increase.

SQL 2008 or later software compression is not used if a backup job that includes SQL 2008 or later data uses Advanced Open File options.

Note: You cannot use this option for backup jobs that deduplicate data.

One-time backup method

Specifies one of the following methods for one-time backups:

- **Full - Back up databases**
This option backs up the entire database. This option is selected by default.
See [“Backing up SQL databases and transaction logs”](#) on page 899.
- **Full Copy - Back up databases (copy)**
This option backs up the entire database without affecting future differential or log backups.
Unlike the Full backup method, the Full Copy backup method does not reset the SQL differential baseline that is used to indicate the database blocks that have changed since the last full backup.
After making a full backup, you can use the Full Copy backup method to make a copy of a SQL database without affecting the baseline backup set required to run future differential backups.
- **Database Snapshot - Read-only, point-in-time copy of another database**
This option creates a read only, point-in-time copy of another database.
See [“Using database snapshots for SQL Server”](#) on page 897.
- **Log No Truncate - Back up without truncating transaction log**
This option backs up the database when it is corrupt or when database files are missing. Since the Log No Truncate backup method does not access the database, you can still back up transactions that you may not be able to access otherwise when the database is in this state. You can then use this transaction log backup along with the database backup and any previous transaction log backups to restore the database to the point at which it failed; however, any uncommitted transactions are rolled back. The Log No Truncate backup method does not remove committed transactions after the log is backed up.
- **Log - Back up and truncate transaction log**
This option backs up only the data contained in the transaction log; it does not back up database data. After the transaction log is backed up, committed transactions are removed (truncated).
If the databases are configured for the SQL Server simple recovery model, log backups are not supported.

To change the recovery model, use the SQL administration tools to set the recovery model to Full. You should run a new full backup if you change the recovery mode before a log backup is run.

Alternatively, you can run full backups only, or run full and differential backups of the SQL databases.

Warning: Data lifecycle management (DLM) deletes all expired backup sets that are created by a one-time backup job. DLM does not keep the last backup set after the retention date expires if the backup set is from a one-time backup. To prevent the backup sets from being automatically deleted, you can manually retain specific backup sets or you can change the expiration date of the backup set.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 335.

See [“Configuring Backup Exec to run a consistency check before every SQL backup”](#) on page 895.

- 11 On the **Backup Options** dialog box, click any of the optional settings in the left pane that you want to set for this job.
- 12 Click **OK**.
- 13 On the **Backup Definition Properties** dialog box, click **OK**.
See [“Adding a stage to a backup definition”](#) on page 219.
See [“Editing backup definitions”](#) on page 206.
See [“Changing default backup job settings”](#) on page 515.

Setting default backup options for SQL Server

You can use the defaults that Backup Exec sets during installation for all SQL Server jobs, or you can choose your own defaults. You can override the default settings when you create individual jobs.

To set default backup options for SQL Server

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Job Defaults**, and then select a backup option.

For example, if you want to set up the default options for SQL Server backups to disk, select Back Up to Disk. The options that appear vary depending on what types of storage devices you configure. You can configure different default options for backup jobs that you send to different types of storage.

- 3 In the left pane, select **Microsoft SQL**.
- 4 Select the appropriate options.

Backup method

Select one of the following backup methods that you want to use for this job:

- Full - Back up databases
This option backs up the entire database. This option is selected by default.
- Full Copy - Back up databases (copy)
This option backs up the entire database without affecting future differential or log backups.
Unlike the Full backup method, the Full Copy backup method does not reset the SQL differential baseline that is used to indicate the database blocks that have changed since the last full backup.
After making a full backup, you can use the Full Copy backup method to make a copy of a SQL database without affecting the baseline backup set required to run future differential backups.

Backup method

Select one of the following SQL-specific backup methods that you want to use for this job:

- **Full - Back up databases**
This option backs up the entire database. This option is selected by default.
- **Full Copy - Back up databases (copy)**
This option backs up the entire database without affecting future differential or log backups.
Unlike the Full backup method, the Full Copy backup method does not reset the SQL differential baseline that is used to indicate the database blocks that have changed since the last full backup.
After making a full backup, you can use the Full Copy backup method to make a copy of a SQL database without affecting the baseline backup set that is required to run future differential backups.
- **Automatic - Backup up transaction log if enabled and then back up database changes since the last full or incremental**
This option lets you back up the entire SQL instance even though some databases may not support log backups. All of the databases are backed up using the Incremental (block level) backup method. In addition, the databases that support log backups are backed up using the Log backup method.
Note: If snapshot is not enabled, an Incremental (block level) backup method cannot be performed and the Differential backup method is used.
- **Log - Back up and truncate transaction log**
This option backs up only the data contained in the transaction log; it does not back up database data. After the transaction log is backed up, committed transactions are removed (truncated).
If the databases are configured for the SQL Server simple recovery model, log backups are not supported. To change the recovery model, use the SQL administration tools to set the recovery model to Full. You should run a new full backup if you change the recovery mode before a log backup is run.
Alternatively, you can run full backups only, or run full and differential backups of the SQL databases.
See [“Configuring Backup Exec to run a consistency check before every SQL backup”](#) on page 895.
- **Log No Truncate - Back up without truncating**

transaction log

This option backs up the database when it is corrupt or database files are missing. Since the Log No Truncate backup method does not access the database, you can still back up the transactions that you may not be able to access otherwise when the database is in this state. You can then use this transaction log backup along with the database backup and any previous transaction log backups to restore the database to the point at which it failed; however, any uncommitted transactions are rolled back. The Log No Truncate backup method does not remove committed transactions after the log is backed up.

- Differential - Backup up database changes since the last full

This option backs up only the changes made to the database or filegroup since the last full backup. Because differential backups allow the restore of a system only to the point in time that the differential backup was created, you should also create multiple log backups between the differential backups.

- Differential (block-level) - Back up database changes since the last full - use with convert to virtual machine job

This option backs up all of the blocks of data and logs that have been created or modified since the last full backup.

- Incremental (block-level) - Back up database changes since the last full or incremental - use with convert to virtual machine job

This option backs up all of the blocks of data and logs that have been created or modified since the last full or incremental backup.

- Database Snapshot - Read-only point-in-time copy of databases

This option creates a read only, point-in-time copy of another database.

See ["Using database snapshots for SQL Server"](#) on page 897.

Note: SQL differential or incremental backups are only supported for conversion to virtual when you use the Automatic, Differential (block-level), or Incremental (block-level) backup methods.

Additionally, Backup Exec runs a full backup when you select the Automatic or Log backup methods if a full backup was not previously run on the database. A full backup also runs for one of the following conditions:

- A new database is added or restored.
- Backup Exec did not run the last full backup.
- Only Full Copy and Incremental backups were run on the database instead of Full backups.

See [“Configuring backup methods for backup jobs”](#) on page 196.

Consistency check before backup

Select one of the following consistency checks to run before a backup:

- None.
This option does not run a consistency check before a backup. Symantec recommends that you always run a consistency check either before or after the backup.
- Full check, excluding indexes.
This option excludes indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough.
- Full check, including indexes.
This option includes indexes in the consistency check. Any errors are logged.
- Physical check only.
This option performs a low-overhead check of the physical consistency of the database. This option only checks the integrity of the physical structure of the page. This option is selected by default.

See [“Configuring Backup Exec to run a consistency check before every SQL backup”](#) on page 895.

Continue with backup if consistency check fails

Choose if you want to continue with the backup operation even if the consistency check fails. You may want to continue with the backup when the consistency check fails if you think that a backup of the database in its current state is better than no backup at all, or if you are backing up a very large database with only a small problem in a table.

Consistency check after backup

Select the consistency check to run after a backup. Because database transactions can occur during or after the consistency check, but before the backup runs, consider running a consistency check after the backup to ensure that the data was consistent at the time of the backup.

The following options are available:

- **None.**
This option does not run a consistency check after a backup. Symantec recommends that you always run a consistency check after the backup. This option is selected by default.
- **Full check, excluding indexes.**
This option excludes indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough.
- **Full check, including indexes.**
This option includes indexes in the consistency check. Any errors are logged.
- **Physical check only.**
This option performs a low overhead check of the physical consistency of the database. This option only checks the integrity of the physical structure of the page. This option is selected by default.

Use checksums on backup (SQL 2005 or later)

Choose to add the checksums to the SQL database data being backed up by Backup Exec. Adding checksums to the data being backed up is required if you want to use the restore option **Run verify only and do not restore data**. Use this option and the **Run verify only and do not restore data** option to ensure that you restore from a verified SQL backup during a restore of the SQL database.

Create on-disk copies of SQL backups to be placed on the SQL server where the database is located

Choose to create an on-disk copy of the SQL database that you want to back up. This option lets you simultaneously back up a SQL database to storage media while also writing a copy of the database to a disk path you specify in the **Save to path** option.

This option gives IT administrators the ability to back up SQL databases while also providing database administrators with copies of the database on disk, which can be used for such things as tests and restores.

Note: This option does not support snapshot technology.

Save to path

Specify a path in which to save on-disk copies of SQL backups.

**SQL Server 2008
Enterprise Edition
software compression**

Select a compression setting that you want to use for this backup job:

- None.
Do not use compression.
- Compress.
Use SQL Server 2008 or later compression if it is supported by the SQL Server instance that is installed.

SQL compresses the data on the computer on which SQL Server 2008 Enterprise Edition or later is installed. Therefore, faster SQL 2008 or later backups should occur if you use SQL compression.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Symantec recommends that you do not use SQL 2008 or later software compression in a backup job that also uses Backup Exec-initiated software compression. Minimal benefits are gained when you enable Backup Exec compression. In fact, in jobs where both compression schemes are used, backup times may increase.

SQL 2008 or later software compression is not used if a backup job that includes SQL 2008 or later data uses Advanced Open File options.

Note: You cannot use this option for backup jobs that deduplicate data.

One-time backup method

Specifies one of the following methods for one-time backups:

- **Full - Back up databases**
This option backs up the entire database. This option is selected by default.
See [“Backing up SQL databases and transaction logs”](#) on page 899.
- **Full Copy - Back up databases (copy)**
This option backs up the entire database without affecting future differential or log backups.
Unlike the Full backup method, the Full Copy backup method does not reset the SQL differential baseline that is used to indicate the database blocks that have changed since the last full backup.
After making a full backup, you can use the Full Copy backup method to make a copy of a SQL database without affecting the baseline backup set required to run future differential backups.
- **Database Snapshot - Read-only, point-in-time copy of another database**
This option creates a read only, point-in-time copy of another database.
See [“Using database snapshots for SQL Server”](#) on page 897.
- **Log No Truncate - Back up without truncating transaction log**
This option backs up the database when it is corrupt or when database files are missing. Since the Log No Truncate backup method does not access the database, you can still back up transactions that you may not be able to access otherwise when the database is in this state. You can then use this transaction log backup along with the database backup and any previous transaction log backups to restore the database to the point at which it failed; however, any uncommitted transactions are rolled back. The Log No Truncate backup method does not remove committed transactions after the log is backed up.
- **Log - Back up and truncate transaction log**
This option backs up only the data contained in the transaction log; it does not back up database data. After the transaction log is backed up, committed transactions are removed (truncated).
If the databases are configured for the SQL Server simple recovery model, log backups are not supported.

To change the recovery model, use the SQL administration tools to set the recovery model to Full. You should run a new full backup if you change the recovery mode before a log backup is run.

Alternatively, you can run full backups only, or run full and differential backups of the SQL databases.

See [“Configuring Backup Exec to run a consistency check before every SQL backup”](#) on page 895.

5 Click **OK**.

See [“Backing up SQL databases and transaction logs”](#) on page 899.

Restoring SQL databases and transaction logs

The SQL Agent lets you restore SQL Server databases. You can restore the databases to their original location or you can redirect the restore to a new location. The number of jobs you decide on depends on the types of backup jobs that protect the database. If you use one job to restore a database, select all the backup sets that you want to apply. Include the full backup, any differential backups, and any log backups.

With very large databases this process can take several hours to complete. During this time Backup Exec reports that no data is being transferred, and the **Byte count** field in the Job Monitor view is not updated. When SQL has completed filling the files with zeros, the restore job continues. This occurs for all database restores but is noticeable only on very large databases.

Restoring encrypted SQL databases

SQL 2008 supports Transparent Database Encryption (TDE), which lets you encrypt SQL 2008 databases at the backup set level.

When you back up a database that uses TDE, Microsoft recommends that you back up the certificate keys and encryption keys with the database. If you do not include the certificate keys and encryption keys, you must perform all backup and restore operations within the selected SQL instance.

Note: Backup Exec can redirect the restore of the database data that used TDE only if the certificate keys and encryption keys are applied to the destination instance. If the certificate keys and encryption keys are not applied to the destination instance, an error appears stating that the certificate thumbprint cannot be found. See your Microsoft SQL 2008 documentation.

How to restore from SQL transaction logs up to a point in time

You can restore transactions from a transaction log up to and including a point in time in the transaction log. After the point in time is reached, recovery from the transaction log is stopped. To find dates and times of transactions, check your client application event log.

If the specified point in time is later than the time contained in the most recent transaction log being restored, then the restore operation succeeds, but a warning is generated and the database remains in an intermediate state. If the specified point in time is before the time contained in the transaction log or logs being restored, no transactions are restored.

How to restore from SQL transaction logs up to a named transaction

You can restore transactions from a transaction log up to and including a named transaction (or mark). After the named transaction is reached, recovery from the transaction log is stopped.

Since named transactions do not necessarily have unique names, you can also specify a date and time after which the restore operation is to search for the named transaction. For example, if you specify a restore from a log up to the named transaction AfternoonBreak, found after 6/02/2000, 12:01 p.m., then the restore operation will not search for AfternoonBreak until after that time. To find dates and times of named transactions, check your client application event log.

If the named transaction is not found, then the restore operation succeeds, but a warning is generated and the database remains in an intermediate state.

Note: The names of transactions are case-sensitive. Ensure you enter the correct upper- and lower-case characters when specifying a named transaction.

Redirecting restores of SQL

You can redirect the following:

- A database backup to a different server, database, or instance.
- Differential and log backups to wherever the associated database is restored.
- A database from a 32-bit or 64-bit platform to any other platform.

You can use both single-job restores and multiple-job restores in redirected restore operations.

To restore SQL databases and transaction logs

- 1 On the **Backup and Restore** tab, right-click the server for which you want to restore data, and then click **Restore**.
 - 2 Select **Microsoft SQL Server databases**, and then click **Next**.
 - 3 Follow the **Restore Wizard** prompts to restore the data.
- See [“Restoring the SQL master database”](#) on page 923.

Restoring the SQL master database

If the master database is damaged, symptoms may include the following:

- An inability to start SQL.
- Segmentation faults or input/output errors.
- A report generated by SQL Database Consistency Checker utility (DBCC).

If the master database is critically damaged and SQL cannot be started, rather than running the Rebuild Master utility or reinstalling SQL to be able to restart SQL, you can replace the corrupted or missing databases with the copies of the master and model databases that Backup Exec automatically creates and updates whenever backups of those databases are run. After SQL is running again, you can restore any other databases, if needed.

If copies of the master and model databases were not made, then you must use Microsoft's rebuildm.exe utility to rebuild the master database and start SQL.

Because all changes made to the master database after the last backup was created are lost when the backup is restored, the changes must be reapplied. If any user databases were created after the master database was backed up, those databases cannot be accessed until the databases are restored from backups or reattached to SQL.

Note: Symantec recommends that you restore the master database first in a separate job. The SQL Server is restarted during the master database restore. If other database restores are included in the same job, the restores will fail.

To restore the SQL master database

- 1 On the **Backup and Restore** tab, in the list of servers view, right-click the server for which you want to restore the SQL master database, and then click **Restore**.
- 2 Select **Microsoft SQL Server databases**, and then click **Next**.

- 3 Select the backup set that contains the last master database backup, and then click **Next**.
- 4 Use the defaults or select the appropriate options and continue through the remaining wizard panels.
- 5 When prompted to run a consistency check, ensure that a check is run after the restore.

After the restore, SQL restarts in multi-user mode.

- 6 Restore the remaining SQL databases.

See [“Restarting SQL using database copies”](#) on page 924.

Restarting SQL using database copies

You can restart SQL manually using copies of the database from previous backups and then restore the master database.

See [“Restoring the SQL master database”](#) on page 923.

Table E-2 Restarting SQL using database copies

Step	Action
Step 1	Ensure that the SQL services are not running. Refer to the SQL Server documentation for details.
Step 2	Verify that the database copies are present. If necessary, restore the master and model database copies from a backup set to the same directory that the original master and model databases are in.
Step 3	Using the Windows Explorer, browse to the default data directory and delete the following files: <ul style="list-style-type: none">■ master.mdf■ mastlog.ldf■ model.mdf■ modellog.ldf.

Table E-2 Restarting SQL using database copies (*continued*)

Step	Action
Step 4	Rename the copies of the databases back to their original names. Do not use read-only files. The SQL services will not start with read-only files.
Step 5	Use the SQL Service Control Manager to start SQL Server.
Step 6	Restore the latest changes to the master database.

The database copies are named master\$4idr, mastlog\$4idr, model\$4idr, and modellog\$4idr.

Table E-3 SQL database copy locations

SQL database copy	Location
An initial installation of SQL 2005 or later	C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\Data*.*
A second installed instance of SQL 2005 or later	C:\Program Files\Microsoft SQL Server\MSSQL.2\MSSQL\Data*.*
A default installation of SQL 2008	C:\Program Files\Microsoft SQL Server\MSSQL10.<instance name>\MSSQL\Data

The following table lists the copied database name and the original database name.

Table E-4 SQL database names

Copied database name	Original database name
master\$4idr	master.mdf
mastlog\$4idr	mastlog.ldf
model\$4idr	model.mdf
modellog\$4idr	modellog.ldf

Disaster recovery of a SQL Server

Backup Exec provides a quicker method for restoring SQL rather than running the Rebuild Master utility or reinstalling SQL to restart SQL. Using Backup Exec, you can replace the corrupted or missing databases with copies of the master and model databases that Backup Exec automatically creates and updates whenever backups of those databases are run.

If you use Simplified Disaster Recovery (SDR), then during an SDR recovery of drive C, it will automatically replace the damaged databases with the copies of the master and model databases. You can then restart SQL, and restore the latest master database backup and any other databases that are necessary.

The following topics are included in this section:

- See [the section called “How to prepare for disaster recovery of SQL”](#) on page 926.
- See [the section called “Requirements for SQL disaster recovery”](#) on page 927.
- See [the section called “Disaster recovery for an entire SQL Server or for SQL databases”](#) on page 927.

How to prepare for disaster recovery of SQL

Do the following to prepare for disaster recovery of SQL:

- Back up both system and user databases and transaction logs regularly.
Copies of the master and model databases are automatically created by Backup Exec whenever you back up the master and model databases. Backup Exec places these copies in the same directory that the databases are in, where they must remain in order to be updated.

The following table includes information about MS SQL database locations:

The copies of the master and model databases are named:

- Master\$4idr
 - Mastlog\$4idr
 - Model\$4idr
 - Modellog\$4idr
- Back up the system drives that contain SQL instances.
Whenever you back up the system drive that contains a SQL instance, copies of the master and model databases are backed up. Backing up the system drive that SQL is on also backs up all the executables and registry settings needed for SQL to run.

- Back up the master database whenever any changes are made to SQL.
- Keep records of any service packs that have been installed.
- Ensure you are prepared to recover the entire server, not just SQL.

Requirements for SQL disaster recovery

To perform a recovery, you will need the following items:

- The latest backup of the SQL directory (\Program Files\Microsoft SQL Server\MSSQL), and the Windows registry/System State.
- The SQL database backups, and differential and log backups.
- An Administrator logon account (or an Administrator equivalent) during the recovery.

Disaster recovery for an entire SQL Server or for SQL databases

You can restore either the entire server, including the SQL databases, from full system backups, or restore only the SQL databases to a newly installed or other available SQL Server.

Restoring the entire server, including the SQL databases, has the added benefit of recovering other applications and data that may have resided on the server at the time of failure, and can be accomplished using one of the following methods:

- Manual recovery of the Windows server, and then manual recovery of the SQL databases. This method involves manually restoring the Windows server from full system backups, and then recovering the SQL databases.
- Simplified Disaster Recovery. This option provides an automated method of restoring the Windows server as well as the SQL databases from full system backups.

See [“Recovery notes for using Simplified Disaster Recovery with Exchange, SQL, SharePoint, CASO, Hyper-V hosts, and the Deduplication Option”](#) on page 771.

To restore only the SQL databases, review the following:

- To restore only the SQL databases to a newly-installed or other available server, the server must be running on the same hardware platform (cross-platform restores are not supported), and the same version of SQL with the same service pack level as the original server.
- To restore SQL databases to an existing installation of SQL with other active databases, you should redirect the restore.
See [“Manual recovery of a SQL Server”](#) on page 928.

Manual recovery of a SQL Server

When you recover SQL manually, you must first restore the Windows server from full system backups. After recovery of the Windows computer is complete, or after the new server installation is available, you can recover the SQL databases.

For more information on how to run a manual disaster recovery, click the following URL:

<http://entsupport.symantec.com/umi/V-290-706>

In order to restore SQL databases, SQL must be running; however, SQL cannot be started unless the master and model databases are present.

You can restore the master and model databases and start SQL using one of the following methods:

- Rename the files created by Backup Exec that replace the master and model databases. After the master and model databases are present on SQL, you must start SQL, and then restore all other databases.
See [“Restarting SQL using database copies”](#) on page 924.
- Reinstall SQL.

This topic only details how to restart SQL by using the copies of the master and model databases made by Backup Exec. For more information on the Rebuild Master utility, or on reinstalling SQL, refer to your Microsoft SQL documentation.

If you are restoring to a new SQL installation, start with the restore of the master database.

See [“Restoring the SQL master database”](#) on page 923.

Symantec Backup Exec Agent for Microsoft Exchange Server

This appendix includes the following topics:

- [About the Backup Exec Exchange Agent](#)
- [Requirements for using the Exchange Agent](#)
- [Granting permissions on the Exchange Server to enable database backups and restores, and Granular Recovery Technology operations](#)
- [About installing the Exchange Agent](#)
- [Adding Exchange Servers and database availability groups to the list of servers on the Backup and Restore tab](#)
- [Managing preferred server configurations for Microsoft Exchange Database Availability Groups](#)
- [Recommended configurations for Exchange](#)
- [Requirements for accessing Exchange mailboxes](#)
- [Backup strategies for Exchange](#)
- [How Granular Recovery Technology works with the Exchange Information Store](#)
- [Snapshot and offhost backups with the Exchange Agent](#)
- [Backing up Exchange data](#)
- [Restoring Exchange data](#)

- [Disaster recovery of an Exchange Server](#)

About the Backup Exec Exchange Agent

The Symantec Backup Exec Agent for Microsoft Exchange Server (Exchange Agent) lets you integrate backups of Microsoft Exchange Server databases with network backups without separate administration or dedicated hardware.

The Exchange Agent provides the following features:

- Restore individual items from backups for which you enable Granular Recovery Technology.
- Restore to a PST file.
- Restore public folders.
- Restore a database to a specific drive and a path.
- Recreate a mailbox account.
- Search for a specific message to restore.
- Select storage groups for backup and restore jobs, or to select one or more databases within the storage group for backup and restore jobs.
- Restore individual databases or storage groups from non-snapshot backups by using the Recovery Storage Group feature in Exchange Server 2007 and Recovery Database feature in Exchange Server 2010/2013. You can restore snapshot backups to a recovery storage group or database.
- Seeding of an Exchange Server 2010/2013 database copy. Seeding adds a database copy to a location on another mailbox server in a database availability group (DAG).
- Snapshot backup and off-host backup on Exchange Server 2007 instances that run on Windows Server 2003.
- Off-host backup with Granular Recovery Technology (GRT) for Exchange Server.

See [“About installing the Exchange Agent”](#) on page 939.

See [“Backup strategies for Exchange”](#) on page 945.

See [“Recommended configurations for Exchange”](#) on page 943.

See [“Granular Recovery Technology”](#) on page 590.

See [“How to use off-host backup to process remote computer backups on the Backup Exec server”](#) on page 1163.

Requirements for using the Exchange Agent

The Backup Exec server must meet the following requirements:

Table F-1 Backup Exec server requirements for the Backup Exec Exchange Agent

Backup Exec server requirements	Description
To support the Exchange Agent	<ul style="list-style-type: none">■ A license for the Symantec Backup Exec Agent for Microsoft Exchange Server (Exchange Agent) must be entered on the Backup Exec server.■ The Backup Exec Agent for Windows must be installed on any remote Exchange Server that you want to back up.■ The Backup Exec server must have access to the Exchange Server. You can find a list of compatible operating systems, platforms, applications, and supported service packs at the following URL: http://entsupport.symantec.com/umi/V-269-1■ Symantec recommends that you use a Backup Exec services account that has domain and local administrator rights on the Exchange server. For Exchange 2010/2013, you can have full Exchange permissions or minimal permissions that enable database backups and restores and Granular Recovery Technology operations. See “Granting permissions on the Exchange Server to enable database backups and restores, and Granular Recovery Technology operations” on page 935.

Table F-1 Backup Exec server requirements for the Backup Exec Exchange Agent (*continued*)

Backup Exec server requirements	Description
To back up Exchange Servers and to support the Backup Exec Data Discovery feature, which allows detection of new backup content within your domain	<p>The Exchange Management Tools for Microsoft Exchange Server must be installed on the Backup Exec server. The management tools on the Backup Exec server must be the same version or later as the management tools that are on the Exchange Server.</p> <p>You can install the management tools when you do a custom install of Microsoft Exchange Server. If you install the management tools and Backup Exec together on a Backup Exec server, install the tools first. If you install Backup Exec before you install the management tools, you must restart the Backup Exec server when the tools installation is complete.</p>
To support Granular Recovery Technology (GRT) for the restore of individual items from Information Store backups	<ul style="list-style-type: none">■ A Microsoft Windows operating system that supports minifilter drivers must be installed for Microsoft Exchange. You can find a list of compatible operating systems, platforms, applications, and supported service packs at the following URL: http://entsupport.symantec.com/umi/V-269-1■ Storage that you use for GRT-enabled backups may have additional requirements.■ If Exchange log files are on an advanced format disk or 512e disk, the Backup Exec server must also have a similar local volume available to perform GRT operations. <p>See “Granular Recovery Technology” on page 590.</p>

The following are requirements for the Exchange Server with the Backup Exec Exchange Agent:

Table F-2 Exchange Server requirements

Exchange Server requirements	Description
To support Exchange Server 2007	<p>Download the Microsoft Exchange Server MAPI Client and Collaboration Data Objects package and install it on Exchange Server 2007.</p> <p>This package provides support for the following:</p> <ul style="list-style-type: none">■ The restore of individual mailboxes, mail messages, and public folders from an Information Store backup.■ The collection of catalog information for a backup for which the Granular Recovery Technology option is enabled and the destination device is tape. <p>You can find this package on the Microsoft website.</p>
For operations on all Exchange Servers	<p>The user account must be a member of the following groups:</p> <ul style="list-style-type: none">■ The Administrators group■ The Domain Admins <p>To support the Granular Recovery Technology option, you must use the appropriate Exchange Server management utility to assign the user account the Exchange Organization Administrators role (2007) or the Exchange Organization Management role (2010/2013).</p> <p>Symantec recommends that the user account have full Exchange permissions on the Exchange Server. If the user account cannot have full Exchange permissions for Exchange 2010/2013, you can grant minimal permissions that enable database backups and restores and Granular Recovery Technology operations</p> <p>See “Granting permissions on the Exchange Server to enable database backups and restores, and Granular Recovery Technology operations” on page 935.</p>

Table F-2 Exchange Server requirements (*continued*)

Exchange Server requirements	Description
To support snapshot backups	<p>Use Microsoft Exchange Server that runs on Windows Server 2003 or later.</p> <p>To select incremental or differential backup methods, Exchange Server 2007 or later must be installed.</p>
To back up and restore Exchange Server 2010/2013	<p>The following are requirements to back up and restore Exchange Server 2010/2013:</p> <ul style="list-style-type: none">■ You must install the Agent for Windows on all of the mailbox server nodes in the DAG to back up the databases on a database availability group (DAG).■ You must have local administrator rights on each node of a Database Availability Group and on the Microsoft Exchange mailbox server to back up and restore Microsoft Exchange database files.■ You can configure minimal permissions for a user account to perform database backups and restores and to support Granular Recovery Technology on Exchange Servers if you cannot have full permissions■ For Microsoft Exchange Server 2013, you must have Microsoft Exchange Server 2013 with Cumulative Update 1 or later installed. <p>See “Granting permissions on the Exchange Server to enable database backups and restores, and Granular Recovery Technology operations” on page 935.</p> <p>See “About the Agent for Windows” on page 777.</p>

For specific operating system requirements for the Exchange Agent, see the Backup Exec Software Compatibility List at the following URL:

<http://entsupport.symantec.com/umi/v-269-1>

See [“Configuring Backup Exec to discover data to back up”](#) on page 571.

Granting permissions on the Exchange Server to enable database backups and restores, and Granular Recovery Technology operations

A user account must access mailboxes on the Exchange Server to perform Backup Exec operations. To gain access to the Exchange Server, you must have full access permissions on the Exchange Server. If you want to restrict the access on the Exchange Server, you can grant minimal permissions that enable users to perform database backups and restores and Granular Recovery Technology (GRT) operations.

Ensure that the user account has local administrator's rights on the Exchange Server and then use one of the following methods to grant the permissions:

- Grant full access permissions at the Organization Administrators or Organization Management level. Symantec recommends that the user account have full Exchange permissions on the Exchange Server to perform Backup Exec operations.

Permissions are then propagated automatically to any new Exchange Servers that you add under the level at which the permissions are assigned.

Note: You must have Exchange administrative permissions to grant permissions to other accounts.

- If the user account cannot have full Exchange permissions for Backup Exec operations, you can grant minimal permissions. Minimal permissions let users perform database backups and restores and Granular Recovery Technology operations explicitly on each Exchange Server.

If you grant permissions explicitly and then add another Exchange Server, you must grant permissions explicitly on the added server as well.

[Granting minimal permissions for a user account to perform database backups and restores of Exchange Servers](#)

[Granting minimal permissions for a user account to support Granular Recovery Technology on Exchange Servers](#)

Granting minimal permissions for a user account to perform database backups and restores of Exchange Servers

You can grant minimal permissions for a user account that let you perform database backups and restores of an Exchange Server.

To grant full permissions for Microsoft Exchange 2010/2013, use an account with the Organization Management role.

To grant minimal permissions for a user account to perform database backups and restores of an Exchange Server

◆ Do one of the following:

To grant permissions for a user account using the Exchange Management Console in Microsoft Exchange 2010 or the Exchange Admin Center in Microsoft Exchange 2013

Add the user account to the following roles:

- Public Folder Management
- Recipient Management
- Server Management

To grant permissions for a user account using the Exchange Management Shell

Do the following in the order listed:

- Type the following command:

```
new-RoleGroup -Name <role
group name> -Roles
@("Database Copies",
"Databases", "Exchange
Servers", "Monitoring", "Mail
Recipient Creation", "Mail
Recipients", "Recipient
Policies" "Mail Enabled
Public Folders", "Public
Folders")
```

For example:

```
new-RoleGroup -Name
BackupExecRoles -Roles
@("Database Copies",
"Databases", "Exchange
Servers", "Monitoring", "Mail
Recipient Creation", "Mail
Recipients", "Recipient
Policies", "Mail Enabled
Public Folders", "Public
Folders")
```

- Type the following command:

```
Add-RoleGroupMember -Identity
<role group name> -Member
<name of the user account>
```

For example:

```
Add-RoleGroupMember -Identity
BackupExecRoles -Member
BackupExecUser
```

Granting minimal permissions for a user account to support Granular Recovery Technology on Exchange Servers

You can grant minimal permissions for a user account that let you support only Granular Recovery Technology (GRT) on an Exchange Server.

For more information about recipient scope, see the Microsoft Exchange documentation.

To grant permissions for a user account to support only Granular Recovery Technology on an Exchange Server using the Exchange Management Shell

1 Type the following command:

```
New-ManagementRole -Name "<management role name>" -Parent  
ApplicationImpersonation
```

For example:

```
New-ManagementRole -Name "SymantecEWSImpersonationRole" -Parent  
ApplicationImpersonation
```

2 Type the following command:

```
New-ManagementRoleAssignment -Role "<management role assignment  
name>" -User <user name> -Name "<assignment name>"
```

For example:

```
New-ManagementRoleAssignment -Role "SymantecEWSImpersonationRole"  
-User BackupExecUser -Name "BackupExecUser-EWSImpersonation"
```

3 Do one of the following:

For Exchange 2010 Type the following command:

```
New-ThrottlingPolicy -Name "<throttling policy
name>" - EWSMaxConcurrency $null
-PowerShellMaxConcurrency $null
-EWSMaxSubscriptions $null
```

For example:

```
New-ThrottlingPolicy -Name
"SymantecEWSRestoreThrottlingPolicy" -
EWSMaxConcurrency $null
-PowerShellMaxConcurrency $null
-EWSPercentTimeInCAS $null -EWSPercentTimeInAD
$null -EWSPercentTimeInMailboxRPC $null
```

For Exchange 2013 Type the following command:

```
New-ThrottlingPolicy -Name "<throttling policy
name>" -EwsCutoffBalance Unlimited -EwsMaxBurst
Unlimited -EwsMaxConcurrency Unlimited
-ExchangeMaxCmdlets Unlimited -MessageRateLimit
Unlimited -PowershellCutoffbalance Unlimited
-PowershellMaxBurst Unlimited
-PowershellMaxCmdlets Unlimited
-PowershellMaxConcurrency Unlimited
-PowershellMaxOperations Unlimited
-RecipientRateLimit Unlimited
-ThrottlingPolicyScope Regular
```

For example:

```
New-ThrottlingPolicy -Name
"SymantecEWSRestoreThrottlingPolicy"
-EwsCutoffBalance Unlimited -EwsMaxBurst
Unlimited -EwsMaxConcurrency Unlimited
-ExchangeMaxCmdlets Unlimited -MessageRateLimit
Unlimited -PowershellCutoffbalance Unlimited
-PowershellMaxBurst Unlimited
-PowershellMaxCmdlets Unlimited
-PowershellMaxConcurrency Unlimited
-PowershellMaxOperations Unlimited
-RecipientRateLimit Unlimited
-ThrottlingPolicyScope Regular
```

4 Type the following command:

```
Set-Mailbox -Identity <user name> -ThrottlingPolicy "throttling  
policy name"
```

For example:

```
Set-Mailbox -Identity BackupExecUser -ThrottlingPolicy  
"SymantecEWSRestoreThrottlingPolicy"
```

5 Type the following command:

```
Set-ThrottlingPolicyAssociation -Identity <user name>  
-ThrottlingPolicy "throttling policy name"
```

For example:

```
Set-ThrottlingPolicyAssociation -Identity BackupExecUser  
-ThrottlingPolicy "SymantecEWSRestoreThrottlingPolicy"
```

About installing the Exchange Agent

The Exchange Agent is installed as part of the Agent for Applications and Databases and can protect local or remote Exchange Server databases.

To protect Exchange Server 2010/2013, you must install Backup Exec on a Microsoft Windows 2008 SP2 or later 64-bit server, Microsoft Windows 2008 R2 64-bit server, Microsoft Windows Server 2012, or Microsoft Windows Server 2012 R2 64-bit.

Note: When you install Microsoft Exchange Tools and Backup Exec together on a server, Exchange Tools must be installed first. If you install Backup Exec before Exchange Tools, you must restart the Backup Exec server after you finish the Exchange Tools installation.

See ["Installing additional Backup Exec options to the local Backup Exec server"](#) on page 76.

Adding Exchange Servers and database availability groups to the list of servers on the Backup and Restore tab

You can add an Exchange Server and a database availability group (DAG) to the list of servers on the **Backup and Restore** tab so that these servers can be selected for backup jobs. When you select the **Add** option in the **Servers and Virtual Hosts**

group on the **Backup and Restore** tab, you can add Microsoft Windows computers and servers.

Note: When you add a Microsoft Exchange database availability group, Symantec recommends that you manually restart each Exchange Server after installing the Agent for Windows. If you select to automatically restart after install, all of your Exchange Servers in the DAG may restart at the same time.

To add an Exchange Server or DAG to the list of servers on the Backup and Restore tab

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Do one of the following:

To add an Exchange Server	Click Microsoft Windows computers and servers .
To add a DAG	Click Microsoft Exchange database availability group .
- 3 Click **Next**.
- 4 Follow the **Add a Server** wizard prompts to add the Exchange Server or DAG to the list of servers in the Backup and Restore tab.

See [“Backing up Exchange data”](#) on page 950.

Managing preferred server configurations for Microsoft Exchange Database Availability Groups

Preferred server configurations are collections of one or more servers and sites that you select as preferred backup sources. Preferred server configurations take priority as backup sources in instances where database copies are replicated between multiple servers. You can create preferred server configurations for Microsoft Exchange Database Availability Groups (DAG).

You do not have to create a preferred server configuration to back up replicated database copies. You can let Backup Exec choose the best server from which to back up the replicated database copies. Designating a preferred server configuration gives you more control over your backup jobs. For example, you can select a local preferred server configuration to avoid having to back up replicated data over your WAN.

Backup Exec automatically includes the children of any site or DAG that you select as part of the preferred server configuration. To ensure that a backup is performed locally, you can select the local site as the preferred server configuration. Backup Exec selects from any of the local servers that belong to that site during the backup job. To ensure that a specific server is used for the backup, select only that server as the preferred server configuration.

You can create preferred server configurations for Microsoft Exchange Database Availability Groups. Preferred server configurations give you more control over your backup jobs since you can specify a preferred server from which Backup Exec backs up replicated data.

This topic contains information about the following subjects:

- [Creating preferred server configurations](#)
- [Deleting preferred server configurations](#)
- [Designating a default preferred server configuration](#)
- [Removing the default status for a preferred server configuration](#)

Creating preferred server configurations

To create preferred server configurations

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Preferred Servers**.
- 2 Click **New**.
- 3 Select an Active Directory forest that contains the Exchange DAG members that you want to specify as preferred servers for backup.
- 4 Type a name for the group of preferred servers.
You must enter a name before the preferred server configuration can be created.
- 5 Click **New** to enable the lists of available and selected servers from which you designate the preferred server.
- 6 In the **Available Servers** list, select the servers and sites that you want to use in the preferred server configuration.
- 7 On the **Preferred Servers** dialog box, click **OK**.
- 8 On the **Manage Preferred Servers** dialog box, click **OK**.

Deleting preferred server configurations

To delete preferred server configurations

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Preferred Servers**.
- 2 Select the preferred server configuration you want to delete.
- 3 Click **Delete**.
- 4 Click **OK**.

Designating a default preferred server configuration

You can designate a preferred server configuration to use as the default for all of your backup jobs that contain the appropriate replication data. When you back up data from a Microsoft Exchange Database Availability Group, you can set up Backup Exec to use your preferred server configuration that you selected as the default. You can override the selected preferred server configuration for specific jobs in the backup job settings.

When you designate a preferred server configuration to use as the default, it is not applied to existing backup jobs. The configuration is used for any subsequent backup jobs that you create.

To designate a default preferred server configuration

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Preferred Servers**.
- 2 Select the preferred server configuration you want to set as the default.
- 3 Click **Set as Default**.
- 4 Click **OK**.

Removing the default status for a preferred server configuration

If you no longer want the preferred server configuration to be the default, you can remove its default status.

To remove the default status for a preferred server configuration

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Preferred Servers**.
- 2 Select the preferred server configuration from which you want to remove the default status.
- 3 Click **Remove as Default**.
- 4 Click **OK**.

Recommended configurations for Exchange

Before starting backups for Exchange, read the following recommendations for configuring Exchange to make it easier to restore from backups:

Table F-3 Recommended configurations for Exchange

Recommendation	Description
Put transaction log files on a separate physical disk from the database.	This is the single most important configuration affecting the performance of Exchange. This configuration also has recovery implications, since transaction logs provide an additional recovery resource.
Make Write Cache unavailable on the SCSI controller.	The Windows operating system does not use buffers, so when Exchange receives a write complete notice from Windows, the write-to-disk has been completed. If Write Cache is enabled, Windows responds as though a write-to-disk has been completed, and will provide this information to Exchange (or other applications) incorrectly. The result could be data corruption if there is a system crash before the operation is actually written to disk.
Make circular logging unavailable if possible.	Circular logging minimizes the risk that the hard disk will be filled with transaction log files. But, if a solid backup strategy is in place, transaction log files are purged during the backup, thus freeing disk space. If circular logging is enabled, transaction log histories are overwritten, incremental and differential backups of storage groups and databases are disabled, and recovery is only possible up to the point of the last full backup.
Avoid making the Exchange Server a domain controller.	For disaster recovery purposes, it is much easier to restore Exchange if you don't have to restore the Active Directory first.
Install Exchange into a domain that has at least two domain controllers.	Active Directory replication is not possible with only one domain controller in a domain. If the domain controller fails and corrupts the Active Directory, some transactions may not be recoverable if they were not included with the last backup. With at least two domain controllers in a domain, databases on the failed domain controller can be updated using replication to fill in missing transactions after the database backups have been restored.

See [“Requirements for accessing Exchange mailboxes”](#) on page 944.

Requirements for accessing Exchange mailboxes

Backup Exec must have access to a uniquely named mailbox within the Exchange organization for Information Store operations, depending on how the backup and restore jobs are configured.

Access to a uniquely named mailbox is required when you do the following:

- Configure a backup job that has all of the following settings:
 - A disk storage device other than a legacy backup-to-disk folder is the destination device.
 - The Granular Recovery Technology option is enabled.
 - A backup method other than a snapshot method is used.
- You restore mailboxes and public folders.

You must use a Backup Exec logon account to connect to the Exchange Server when you select mailboxes or public folders for backup. Backup Exec attempts to find a mailbox with the same name as the user name that is stored in the Backup Exec logon account.

If you use a Backup Exec logon account that stores a unique user name and has a corresponding mailbox with the same name, then you are not prompted for an additional logon account. Otherwise, you must choose or create a Backup Exec logon account that stores the name of a unique mailbox within the Exchange organization.

A unique name does not share the first five characters in another mailbox name. For example, if EXCH1 is entered as the mailbox name, and there is another mailbox name such as EXCH1BACKUP, then Backup Exec cannot accept the name. You are prompted to choose another mailbox name.

You can choose or create a logon account that meets any of the following requirements:

- A logon account for which the user name matches a unique mailbox name.
- A logon account that uses a unique alias to a mailbox. The user account that connects to the Exchange Server must also have access to this mailbox.
- A logon account that uses the full computer name for a mailbox. The user account that connects to the Exchange Server must also have access to this mailbox.
An example of a full computer name is:
`/O=Exchange_Organization/OU=Administrative_Group/CN=Recipients/CN=mailbox_name`

See [“Granular Recovery Technology”](#) on page 590.

See [“Backup Exec logon accounts”](#) on page 610.

Backup strategies for Exchange

Backup Exec incorporates online, nondisruptive Exchange database protection as part of everyday backup routines, which increases the chance of data recovery and minimizes data loss without inhibiting daily activity. Backup Exec protects Exchange data, including the individual storage group, database, mailbox, and public folder with full, copy, incremental, and differential backups.

To decide which backup methods to use, consider the following:

- In small office environments with relatively small numbers of messages passing through the system, a daily full backup provides good data protection and the quickest recovery. If log file growth becomes an issue, consider using incremental online backups at midday to provide an added recovery point and manage the log file growth for you automatically.
- In large environments, incremental backups should be used to provide more frequent recovery point options throughout the day and to manage log file growth. Many organizations run full backups on a weekly basis, preferring to run incremental backups throughout the week to keep backup run time to a minimum. The trade-off with this technique occurs at recovery time when you must recover from the full backup and from each incremental backup as well.

What works best for you is based on the size of your environment, the number of transactions processed each day, and the expectations of your users when a recovery is required.

Consider the following backup strategies:

- Run full backups with the option to enable the restore of individual items selected so that you can restore individual mail messages and folders without restoring the entire database.

Depending on your environment, run full backups as follows:

- As frequently as possible, no less than once a day.
 - Daily with differential backups used at regular periods throughout the day.
 - Every few days (no less than weekly) with frequent incremental backups in between each full backup.
- Run Exchange backup jobs separately from other backup jobs.

In addition to backing up Exchange storage groups or databases, you should also back up the following on a regular basis:

Table F-4 Backup selections for Exchange configuration data

Recommended backup selections for configuration data	Description
File system	<p>Back up folders and drives containing files for Windows and Exchange. Usually, this is the root drive C:\, but may be different in each environment.</p> <p>Note: Back up the C:\ drive, but do not back up the virtual drive that is created by Exchange, if this virtual drive exists in your environment. It is intended only to provide Windows Explorer access to the Exchange data, but all file system functions may not be replicated. Backup and restore operations are not recommended or supported.</p>
Windows registry	<p>Back up the registry by running a full backup.</p>
System State and/or Shadow Copy Components	<p>Select System State and run a full backup to back up the following:</p> <ul style="list-style-type: none">■ The Internet Information Service (IIS) metabase■ The Windows registry <p>See “About selecting data to back up ” on page 169.</p> <p>If the entire server must be restored, you must restore both System State and Shadow Copy Components.</p>
Active Directory	<p>To back up Active Directory, select System State on the domain controllers and run a full backup.</p> <p>When there are configuration changes on the Exchange Server database, such as when objects are added, modified, or deleted, back up the Active Directory on the domain controllers.</p> <p>Note: Spread multiple domain controllers throughout each domain for efficient Active Directory replication, and so that if one domain controller fails, redundancy is still provided.</p>

Note: Configure an Information Store backup for which the Granular Recovery Technology (GRT) option is enabled to restore individual mailboxes, mail messages, and public folders.

See [“Granular Recovery Technology”](#) on page 590.

See [“Disaster recovery of an Exchange Server”](#) on page 968.

How Granular Recovery Technology works with the Exchange Information Store

Backup Exec Granular Recovery Technology (GRT) lets you restore individual items from an Information Store backup without having to restore the whole backup. You should review the requirements for a GRT-enabled backup before you configure it.

See [“Granular Recovery Technology”](#) on page 590.

When you select items to restore from GRT-enabled backups, you cannot select the top level of the Information Store. To restore these items, you must restore the entire mailbox.

Note: When you perform a granular restore of a linked mailbox, shared mailbox, or site mailbox, do not select the **Recreate user accounts and mailboxes if they do not already exist on the destination server** option. These types of mailboxes must be created manually before you perform the restore. However, you can restore user accounts for linked mailboxes when you perform a granular restore of Active Directory from a backup of a domain controller. For more information about creating these types of mailboxes, see the Microsoft Exchange documentation.

You can also enable GRT when you create an off-host backup for the Information Store. Off-host backup lets Backup Exec move the backup process from the host computer to the Backup Exec server. The host computer is the remote computer that contains the volumes that you selected for backup. To run a GRT-enabled off-host backup, you must install the Backup Exec Advanced Disk-based Option on the Backup Exec server.

GRT and Microsoft Exchange Web Services

Backup Exec uses Microsoft Exchange Web Services (EWS) to support the Granular Recovery Technology option. EWS provides support for the restore of individual mailboxes, mail messages, and public folders from an Exchange Server 2010/2013 database backup.

Note: You do not need to install the MAPI Client and Collaboration Data Objects package if you use EWS.

To use EWS to restore individual items, Backup Exec disables the client throttling policy for the resource credentials you specify for the restore job. The client throttling policy enforces connection bandwidth limits on the Client Access server.

Backup Exec also creates an impersonation role and a role assignment for Exchange Impersonation. Exchange Impersonation role assignment associates the impersonation role with the Backup Exec resource credentials you specify for the restore job.

Backup Exec creates and assigns the following roles:

- SymantecEWSImpersonationRole
- SymantecEWSImpersonationRoleAssignment

See [“Granular Recovery Technology”](#) on page 590.

See [“Configuring delayed catalog operations to improve backup performance”](#) on page 536.

See [“About the Advanced Disk-based Backup Option”](#) on page 1155.

Snapshot and offhost backups with the Exchange Agent

The Exchange Agent supports the Microsoft Volume Shadow Copy Service (VSS), a snapshot provider service only available on Windows Server 2003 or later. Using VSS, a point-in-time view of the Exchange database is snapped and then backed up, leaving the actual Exchange database open and available for users.

Offhost backup enables the backup operation to be processed on a Backup Exec server instead of on the Exchange Server. Moving the backup from the Exchange Server to a Backup Exec server enables better backup performance and frees the remote computer as well.

If the Advanced Disk-based Backup Option (ADBO) is installed on the Backup Exec server, you can use the Backup Exec Granular Recovery Technology (GRT) option when you create an offhost backup for the Information Store.

See [“Setting default Granular Recovery Technology \(GRT\) options”](#) on page 596.

This topic includes information on the following subjects:

- [Configuring a snapshot backup for Exchange data](#)
- [Troubleshooting Exchange Agent snapshot and offhost jobs](#)

The Exchange Agent snapshot does not support the following:

- NAS configurations
- Mixing snapshot backups and non-snapshot backups
Due to a Microsoft Exchange limitation, if non-snapshot backups are run as part of a backup strategy, then snapshot backups should not be run. If snapshot backups are run, non-snapshot backups should not be done.

You can find a list of compatible storage at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

The type of backup method that is available when using VSS with the Exchange Agent depends on the version of Exchange Server:

Table F-5 Available backup methods for Exchange snapshot versions

Exchange version	Available backup methods
Exchange Server 2007/2010/2013	The following backup methods are available: <ul style="list-style-type: none">■ Full■ Copy■ Differential■ Incremental storage group level snapshot backup■ Individual database restore
Exchange Server 2007	LCR/CCR - Back up from the passive copy or the active copy. Note: You cannot back up the passive copy of the Standby Continuous Replication (SCR) database with Exchange Server 2007. The SCR is not available for backup selection.

Configuring a snapshot backup for Exchange data

You can run a snapshot backup for Exchange data.

Table F-6 Configuring a snapshot backup for Exchange data

Step	Action
Step 1	Symantec recommends that you perform consistency checks before running a snapshot backup. See "Backing up Exchange data" on page 950.
Step 2	Create an Exchange backup job. See "Backing up Exchange data" on page 950.

Table F-6 Configuring a snapshot backup for Exchange data (*continued*)

Step	Action
Step 3	<p>If data that is not supported for snapshot backup is included in the backup selections, check Process logical volumes for backup one at a time to allow the job to complete with errors.</p> <p>This option is on the Advanced Disk-based Backup options on the Backup Job Defaults dialog box and on the Options dialog box for a backup job.</p>
Step 4	<p>Schedule or start the backup job.</p> <p>See “Backing up data” on page 158.</p>

Troubleshooting Exchange Agent snapshot and offhost jobs

An Exchange Agent snapshot job fails on the following conditions:

- The Exchange Agent snapshot fails.
- If incremental or differential backup methods are selected, and Exchange Server 2007 or later is not installed.
- If circular logging is enabled, and incremental or differential backup methods are selected.
- If an unsupported version of Exchange is installed. To verify supported versions of Exchange, see the Backup Exec Software Compatibility List at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Backing up Exchange data

To back up Exchange data, you can select the following:

- Multiple storage groups
- Individual storage groups
- Individual databases (Exchange 2010/2013 only)
- Database availability groups (DAG)

You must enter an Exchange Agent license on the Backup Exec server for each Exchange Server in the DAG that you want to back up. You must then install the Agent for Windows on all the servers in the DAG.

Each database in the DAG must be backed up through the DAG container that displays in the list of servers on the **Backup and Restore** tab. The DAG container displays an Exchange logo on the server.

Note: Backup Exec does not support DAGs without cluster administrative access points.

Note: If you add Exchange databases or storage groups after you create a backup job, you must edit the backup job to include the new selections.

Symantec recommends that you select individual storage groups for backup rather than selecting individual databases in storage groups. Although you can select individual databases in a storage group for backup, the transaction logs for the entire storage group are backed up.

For example, if you back up a single database in a storage group that contains four databases, the entire collection of transaction logs for the storage group is also backed up. The transaction logs are not deleted until a full backup is run on every database in the storage group. You can still restore an individual database from a storage group backup.

Note: To perform incremental and differential backups of storage groups, ensure that the Microsoft Exchange setting for circular logging is not enabled on the storage group.

You can set backup job default options for all Exchange backup jobs. Each time you create a backup job, the job uses the default options unless you change the options for that particular job.

Automatic exclusion of Exchange data during volume-level backups

If you select a volume that contains Exchange data for backup, the Exchange Agent uses Active File Exclusion to automatically exclude Exchange data that should not be included in a volume-level backup. For example, .EDB and .STM files, as well as transaction log files, should not be part of a volume-level backup because they are opened for exclusive use by Exchange.

Without this exclusion, these files appear as in use - skipped during a non-snapshot backup. During a snapshot backup, these files may be backed up in an inconsistent state, which could create restore issues.

While it is not recommended, if you want to include Exchange data in a volume-level backup, you must first dismount the storage groups or databases that you want backed up, and then run the backup job.

To back up Exchange data

- 1 On the **Backup and Restore** tab, right-click the server that contains the Exchange data that you want to back up.

To back up multiple servers, Shift + click or Ctrl + click the server names, and then right-click one of the selected servers.
- 2 Select **Backup**, and then select the type of backup that you want to perform.
- 3 On the **Backup Definition Properties** dialog box, in the **Selections** box, click **Edit**.
- 4 On the **Backup Selections** dialog box, check the resources that you want to back up and uncheck the resources that you do not want to back up.
- 5 Click **OK**.
- 6 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**.
- 7 On the **Backup Options** dialog box, in the left pane, click **Schedule**, and then select the schedule for this job.
- 8 On the **Backup Options** dialog box, in the left pane, click **Storage**, and then select a storage device for this job.
- 9 On the **Backup Options** dialog box, in the left pane, select **Microsoft Exchange**.
- 10 Set any of the following options for this job:

Perform a consistency check before the backup when using Microsoft Volume Shadow Copy Service (VSS) snapshot provider	<p>Select this option to run a consistency check when the Microsoft Volume Shadow Copy Service option is selected. The option Microsoft Volume Shadow Copy Service is automatically used whenever a software backup is selected on the Advanced Disk-based Backup properties. You can also select the Microsoft Volume Shadow Copy Service on the Advanced Open File backup properties.</p> <p>The consistency check, which is run on the snapshot, determines if possible data corruption exists.</p> <p>If this option is selected, and the dependent option Continue with backup if consistency check fails is not selected, then data for specific Exchange objects that are determined to be corrupt are not backed up. All other non-corrupt Exchange objects are backed up.</p> <p>For example, if any transaction log file for a Storage Group is corrupt, then none of the transaction log files are backed up for that Storage Group when the Continue with backup if consistency check fails option is not selected. However, the Exchange database files are backed up if Backup Exec determines they are not corrupt. Similarly, if a specific Exchange database file is corrupt, then backup is skipped only for that corrupt database file. All other non-corrupt database files and transaction log files are backed up.</p> <p>When the option Continue with backup if consistency check fails is enabled, then all Exchange data is backed up regardless if corrupt files exist.</p> <p>See “Snapshot and offhost backups with the Exchange Agent” on page 948.</p>
Continue with the backup if consistency check fails	<p>Select this option to continue the backup job even if the consistency check fails. You may want the job to continue if you think a backup of the database in its current state is better than no backup at all, or if you are backing up a very large database that may have only a small problem.</p>

**High
Availability
Server
(Exchange
2007 or later)**

Specify one of the following backup sources for Exchange:

- **Back up from the passive copy and if not available, try the active copy (recommended)**

Lets you back up a passive copy of the database by default. For Exchange 2010/2013, Backup Exec selects the passive copy based on your selections in the **Preferred Server** settings. However, if the passive copy is not available, Backup Exec backs up the active copy of the database. During the backup, database performance degradation can occur if you have to back up the database over a WAN.

- **Back up from the passive copy only (job fails if not available)**

Lets you back up a passive copy of the database. If Backup Exec cannot access the passive copy, the job fails. In this case, neither the active nor the passive database is backed up. Select this option when you do not want to affect the performance of the active copy of the database.

For Exchange Server 2010/2013, Backup Exec selects the passive copy based on your selections in the **Preferred Server** settings.

Note: You must have the preferred server settings configured to use this option.

See ["Managing preferred server configurations for Microsoft Exchange Database Availability Groups"](#) on page 940.

- **Back up from the active copy only (job fails if not available)**

Lets you back up the active copy of the database. If Backup Exec cannot access the active copy, the job fails. Therefore, neither the active copy nor the passive copy is backed up.

The active copy contains newer information than the passive copy. When you back up the active copy, you have a backup of the most recent database data.

Note: You cannot back up the passive copy of the Standby Continuous Replication (SCR) database with Exchange Server 2007. The SCR is not available for backup selection.

Backup method

Specify one of the following backup methods that you want to use for this job:

- **Full - Back up databases and logs (truncate logs)**
This option backs up the databases as well as their associated transaction log files. After the databases and transaction logs are backed up, the transaction log files are deleted if all transactions are committed to the database.
- **Full Copy - Back up databases and logs**
This option backs up the databases as well as their associated transaction log files. However, the transaction logs are not deleted after they are backed up.
You can use this option to make a full backup of a database without disturbing the state of ongoing incremental or differential backups.

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual mailboxes, mail messages, and public folders from Information Store backups

Select this option to enable the restore of individual items from Information Store backups. Ensure that the backups meet the requirements for Granular Recovery Technology.

Note: Symantec recommends that you do not send an incremental GRT-enabled Exchange backup to a deduplication disk storage device. The transaction logs contain primarily unique data that does not deduplicate well. For best results, create a backup definition that runs a full backup of Exchange to a deduplication disk storage device, and then runs an incremental backup to a disk storage device.

See [“Granular Recovery Technology”](#) on page 590.

See [“Configuring delayed catalog operations to improve backup performance”](#) on page 536.

Backup method

Specify one of the following Exchange-specific backup methods that you want to use for this job:

- **Full - Back up databases and logs (truncate logs)**
This option backs up the databases as well as their associated transaction log files. After the databases and transaction logs are backed up, the transaction log files are deleted if all transactions are committed to the database.
- **Full Copy - Back up databases and logs**
This option backs up the databases as well as their associated transaction log files. However, the transaction logs are not deleted after they are backed up.
You can use this option to make a full backup of a database without disturbing the state of ongoing incremental or differential backups.
- **Differential - Back up logs**
This option backs up all of the transaction logs that have been created or modified since the last full backup. However, the transaction logs are not deleted after they are backed up.
To restore from differential backups, the last differential backup and the last full backup are required.
If circular logging is enabled, differential backups cannot be performed.
- **Incremental - Back up logs (truncate logs)**
This option backs up all of the transaction logs that have been created or modified since the last full or incremental backup, and then delete the transaction logs that have been committed to the database.
To restore from incremental backups, the last full backup and all the incremental backups done since the last full backup are required.
See [“Snapshot and offhost backups with the Exchange Agent”](#) on page 948.

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual mailboxes, mail messages, and public folders from Information Store backups	<p>Select this option to enable the restore of individual items from Information Store backups. Ensure that the backups meet the requirements for Granular Recovery Technology.</p> <p>Note: Symantec recommends that you do not send an incremental GRT-enabled Exchange backup to a deduplication disk storage device. The transaction logs contain primarily unique data that does not deduplicate well. For best results, create a backup definition that runs a full backup of Exchange to a deduplication disk storage device, and then runs an incremental backup to a disk storage device.</p> <p>See “Granular Recovery Technology” on page 590.</p> <p>See “Configuring delayed catalog operations to improve backup performance” on page 536.</p>
---	---

- 11 Set any additional options for this job.

See [“Backing up data”](#) on page 158.

- 12 Click **OK**.

See [“Adding Exchange Servers and database availability groups to the list of servers on the Backup and Restore tab”](#) on page 939.

See [“Changing default backup job settings”](#) on page 515.

See [“Editing backup definitions”](#) on page 206.

See [“Editing a stage”](#) on page 221.

Setting default backup options for Exchange Server

You can use the defaults that Backup Exec sets during installation for all Exchange Server jobs, or you can choose your own defaults. You can override the default settings when you create individual jobs.

To set default backup options for Exchange Server

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Job Defaults**, and then select a backup option.

For example, if you want to set up the default options for Exchange Server backups to disk, select **Back Up to Disk**. The options that appear vary depending on what types of storage devices you configure. You can configure different default options for the backup jobs that you send to different types of storage.

- 3 In the left pane, select **Microsoft Exchange**.
- 4 Select the appropriate options.

Perform a consistency check before the backup when using Microsoft Volume Shadow Copy Service (VSS) snapshot provider	<p>Select this option to perform a consistency check when the Microsoft Volume Shadow Copy Service option is selected. The option Microsoft Volume Shadow Copy Service is automatically used whenever a software backup is selected on the Advanced Disk-based Backup properties. You can also select the Microsoft Volume Shadow Copy Service on the Advanced Open File backup properties.</p> <p>The consistency check, which is run on the snapshot, determines if possible data corruption exists.</p> <p>If this option is selected, and the dependent option Continue with backup if consistency check fails is not selected, then data for specific Exchange objects that are determined to be corrupt are not backed up. All other non-corrupt Exchange objects are backed up.</p> <p>For example, if any transaction log file for a Storage Group is corrupt, then none of the transaction log files are backed up for that Storage Group when the Continue with backup if consistency check fails option is not selected. However, the Exchange database files are backed up if Backup Exec determines they are not corrupt. Similarly, if a specific Exchange database file is corrupt, then backup is skipped only for that corrupt database file. All other non-corrupt database files and transaction log files are backed up.</p>
--	--

When the option **Continue with backup if consistency check fails** is enabled, then all Exchange data is backed up regardless if corrupt files exist.

See [“Snapshot and offhost backups with the Exchange Agent”](#) on page 948.

Continue with the backup if consistency check fails Select this option to continue the backup job even if the consistency check fails. You may want the job to continue if you think a backup of the database in its current state is better than no backup at all, or if you are backing up a very large database that may have only a small problem.

High Availability Server (Exchange 2007 or later)

Specify one of the following backup sources for Exchange:

- Back up from the passive copy and if not available, try the active copy (recommended)
Lets you back up a passive copy of the database by default. For Exchange 2010/2013, Backup Exec selects the passive copy based on your selections in the **Preferred Server** settings. However, if the passive copy is not available, Backup Exec backs up the active copy of the database. During the backup, database performance degradation can occur if you have to back up the database over a WAN.
- Back up from the passive copy only (job fails if not available)
Lets you back up a passive copy of the database. If Backup Exec cannot access the passive copy, the job fails. In this case, neither the active nor the passive database is backed up. Select this option when you do not want to affect the performance of the active copy of the database. For Exchange Server 2010/2013, Backup Exec selects the passive copy based on your selections in the Preferred Server settings.

Note: You must have the preferred server settings configured to use this option.

See ["Managing preferred server configurations for Microsoft Exchange Database Availability Groups"](#) on page 940.

- Back up from the active copy only (job fails if not available)
Lets you back up the active copy of the database. If Backup Exec cannot access the active copy, the job fails. Therefore, neither the active copy nor the passive copy is backed up.
The active copy contains newer information than the passive copy. When you back up the active copy, you have a backup of the most recent database data.

Note: You cannot back up the passive copy of the Standby Continuous Replication (SCR) database with Exchange Server 2007. The SCR is not available for backup selection.

Backup method Specify one of the following backup methods that you want to use for this job:

- **Full - Back up databases and logs (truncate logs)**
This option backs up the databases as well as their associated transaction log files. After the databases and transaction logs are backed up, the transaction log files are deleted if all transactions are committed to the database.
- **Full Copy - Back up databases and logs**
This option backs up the databases as well as their associated transaction log files. However, the transaction logs are not deleted after they are backed up.
You can use this option to make a full backup of a database without disturbing the state of ongoing incremental or differential backups.

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual mailboxes, mail messages, and public folders from Information Store backups Select this option to enable the restore of individual items from Information Store backups. Ensure that the backups meet the requirements for Granular Recovery Technology.
Note: Symantec recommends that you do not send an incremental GRT-enabled Exchange backup to a deduplication disk storage device. The transaction logs contain primarily unique data that does not deduplicate well. For best results, create a backup definition that runs a full backup of Exchange to a deduplication disk storage device, and then runs an incremental backup to a disk storage device.
See [“Granular Recovery Technology”](#) on page 590.
See [“Configuring delayed catalog operations to improve backup performance”](#) on page 536.

Backup method

Select one of the following Exchange-specific backup methods that you want to use for this job:

- **Full - Back up databases and logs (truncate logs)**
This option backs up the databases as well as their associated transaction log files. After the databases and transaction logs are backed up, the transaction log files are deleted if all transactions are committed to the database.
- **Full Copy - Back up databases and logs**
This option backs up the databases as well as their associated transaction log files. However, the transaction logs are not deleted after they are backed up.
You can use this option to make a full backup of a database without disturbing the state of ongoing incremental or differential backups.
- **Differential - Back up logs**
This option backs up all of the transaction logs that have been created or modified since the last full backup. However, the transaction logs are not deleted after they are backed up.
To restore from differential backups, the last differential backup and the last full backup are required.
- **Incremental - Back up logs (truncate logs)**
This option backs up all of the transaction logs that have been created or modified since the last full or incremental backup, and then delete the transaction logs that have been committed to the database.
To restore from incremental backups, the last full backup and all the incremental backups done since the last full backup are required.
See [“Snapshot and offhost backups with the Exchange Agent”](#) on page 948.

If circular logging is enabled, incremental and differential backups cannot be performed.

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual mailboxes, mail messages, and public folders from Information Store backups	<p>Select this option to enable the restore of individual items from incremental Information Store backups. Ensure that the backups meet the requirements for Granular Recovery Technology.</p> <p>Note: Symantec recommends that you do not send an incremental GRT-enabled Exchange backup to a deduplication disk storage device. The transaction logs contain primarily unique data that does not deduplicate well. For best results, create a backup definition that runs a full backup of Exchange to a deduplication disk storage device, and then runs an incremental backup to a disk storage device.</p> <p>See “Granular Recovery Technology” on page 590.</p> <p>See “Configuring delayed catalog operations to improve backup performance” on page 536.</p>
---	---

- 5 Click **OK**.

Restoring Exchange data

The Exchange Agent lets you restore Exchange databases and storage groups as well as individual mailbox items. You can restore items to their original location or you can redirect the restore to a new location.

Note: The **Restore Wizard** displays only up to 30,000 items. If you want to restore data from a database or mailbox that contains more than 30,000 items, you must search for the item that you want to restore. You can use search criteria such as the type of data and the date of the original backup to help reduce the number of items that display.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

The requirements and procedures for restoring Exchange data vary depending on the backup strategy you used. Before you restore Exchange data, you should review the required configuration and tasks.

You can restore Exchange data in the following ways:

- Use the recovery storage group or recovery database to recover data from an older backup copy of the store without disturbing client access to current data.
- Restore Exchange data from snapshot backups.
- Restore individual Exchange items from a backup that is enabled for Granular Recovery Technology (GRT).
See [“How Granular Recovery Technology works with the Exchange Information Store ”](#) on page 947.
- Restore Exchange data to a server other than the one from which it was backed up.

This topic includes the following information:

- [Requirements for restoring Exchange](#)
- [Configuring a destination database for restore data in Exchange](#)
- [Restore data by using the Exchange 2007 recovery storage group or Exchange Server 2010/2013 recovery database](#)
- [About restoring Exchange data from snapshot backups](#)
- [Restoring individual Exchange public folder messages from tape by duplicating backup sets](#)
- [About redirecting Exchange restore data](#)
- [About redirecting Exchange mailbox items](#)

Requirements for restoring Exchange

Review the following before restoring Exchange:

- The storage groups and databases must already exist on the destination server, and must have the same names as the original storage groups or databases.
- The destination server must have the same Organization and Administrative Group name as the source server.
- Before you start the restore, configure the destination databases so that they can be overwritten by a restore.

Configuring a destination database for restore data in Exchange

Before you restore Exchange, you should configure the destination database.

To configure a database

1 Do one of the following:

For Exchange 2007/2010	Open the Exchange Management Console utility.
For Exchange 2013	Open the Exchange Admin Center.

2 Right-click the database that you want to overwrite.

3 Click **Properties**.

4 Do one of the following:

For Exchange 2007	On the Database tab, select This database can be overwritten by a restore .
For Exchange Server 2010/2013	On the Maintenance tab, select This database can be overwritten by a restore .

Restore data by using the Exchange 2007 recovery storage group or Exchange Server 2010/2013 recovery database

The Recovery Storage Group (RSG) feature in Exchange 2007 lets you mount a second copy of an Exchange mailbox store on any Exchange Server in the same Exchange Administrative Group as the original while the original store is still running and serving clients. This allows you to recover data from an older backup copy of the store without disturbing client access to current data.

Exchange Server 2010/2013 uses recovery databases instead of RSGs. Each server has a recovery database and there cannot be more than one mounted recovery database.

See your Microsoft Exchange documentation for more information about RSGs and recovery databases.

After the RSG or recovery database is created, you can restore online backup sets to it. Then you can use the version of the Exchange Management Shell in Exchange to extract mailbox data from the stores into .PST files, and optionally merge the extracted data back into the online stores.

If the RSG or recovery database resides on a different Exchange Server than the databases you are restoring, you should review the requirements for redirecting the restore of Exchange storage groups or recovery databases.

About redirecting Exchange restore data

Following are requirements for restoring data using the Exchange 2007 data Recovery Storage Group (RSG) or Exchange Server 2010/2013 recovery database:

- If multiple stores are selected for restore, mailbox stores in the RSG must come from the same storage group. You cannot add mailbox stores from different storage groups to the RSG at the same time.
- Public folder stores are not supported for restore using the RSG.
- Do not mount mailbox stores in the RSG before the restore. If you do mount the stores before the restore, then you must dismount them. Select the following option on the database property page in Exchange System Manager:
This database can be overwritten by a restore.
Then, delete any files created in the data path for the RSG and added stores prior to restoring them.
Any files created in the data path for the RSG and added store or stores should be deleted as well, if stores were mounted prior to the restore.
- On the server that hosts the RSG, there must be a storage group with the same name as the original storage group for the data you are restoring. If no such storage group exists on the server, then you can use that name for the RSG when you create it.
- The Active Directory topology of the Exchange system must be intact and in the same state it was in when the backup was made. You cannot restore mailbox stores that were deleted and recreated. In addition, you cannot recover mailboxes from stores if the mailboxes were deleted and purged from the system or moved to other servers or mailbox stores.
- When the RSG exists on a server, the mailbox stores that it contains are the only stores that can be restored on that server by default. Symantec recommends that you create the RSG only when you intend to recover data using it, and remove the RSG from the server after the data recovery is complete.
- You can have more than one recovery database; however, you can only mount one recovery database to recover data.
- Do not mount the recovery database before the restore. If you do mount the recovery database before the restore, you must dismount it. Select the **This database can be overwritten by a restore** option on the database property page in the Exchange Management Console utility.

Refer to your Microsoft Exchange Server documentation for more information on the requirements and restrictions of recovering Exchange data.

About restoring Exchange data from snapshot backups

Note the following when restoring Exchange data from snapshot backups:

- If circular logging is enabled, only point-in-time, loss restores are possible. Roll-forward, no-loss restores cannot be performed.
- For Exchange 2007, individual database restores cannot be performed. The job will fail.

See [“Snapshot and offhost backups with the Exchange Agent”](#) on page 948.

Restoring individual Exchange public folder messages from tape by duplicating backup sets

To restore individual public folder messages from tape, you must first duplicate the backup sets that contain the messages to disk storage. You can then restore the data from that disk storage.

The backup that you want to restore from must be a full backup or a copy backup. If there is an incremental or differential backup that is subsequent to the full backup, then you can restore individual items from the incremental or differential backup. The backup sets for the full backup and the incremental or differential backup must be on the same volume.

You cannot restore individual public folder messages from tape if the original backup is an incremental backup.

To restore individual Exchange public folder messages from tape by duplicating backup sets to disk

- 1 Insert the tape containing the required Exchange backup sets into a tape drive.
- 2 Create a duplicate backup sets job.

See [“Duplicating backup sets or a job history manually”](#) on page 221.

- 3 After the job completes, run a restore job to restore the individual public folder messages from the Exchange backup sets that are duplicated on the disk storage.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

About redirecting Exchange restore data

With Backup Exec, you can restore Exchange data to the server from which it was backed up or redirect the Exchange data to a different location. When redirecting Exchange data, the service pack on the Exchange Server where data is being redirected should be the same as the service pack on the original Exchange Server.

Following are requirements for redirecting Exchange storage group and database restores:

- The storage groups and databases must already exist on the destination server.
- The destination server must have the same Organization and Administrative Group name as the source server.
- The destination databases must be configured so that they can be overwritten. See [the section called “Configuring a destination database for restore data in Exchange”](#) on page 963.

You cannot redirect the restore of the following:

- A version of Exchange Server database to a different version of the database. Service packs for both Exchange Servers should also be the same.
- Site Replication Service (SRS) and Key Management Service (KMS). These services are dependent on the computer they reside on; redirection to another computer is not supported and could result in the loss of functionality of these services.

Note: KMS is not available in Exchange.

Before starting the redirected restore job, review information on finding and viewing specific data to restore, as well as for details on restore options and submitting restore jobs.

After completing the restore, it is recommended that a full backup of the restored databases be performed.

See [“Backing up Exchange data”](#) on page 950.

About redirecting Exchange mailbox items

With Backup Exec, you can restore mailbox items such as mailboxes and public folders to a different mailbox on the same server or to a different location.

You can also restore mailboxes or mailbox items to a .PST, which is a Microsoft Outlook data file that is compatible with Microsoft Outlook 2007 or 2010.

Following are requirements for redirecting Exchange mailbox and public folder items back to Exchange:

- The specified mailbox or public folder store must exist.
- The Backup Exec logon account must have rights to the destination mailbox.
- To overwrite existing public folder data, the logon account must have ownership of the public data. In this rare situation, Symantec recommends that when you restore public folder data, do not use the restore option **Overwrite existing mail messages and folders** on the Restore wizard panel titled **How do you want to restore the items?**

Following are requirements for restoring to a .PST file:

- The Agent for Windows must be installed on the destination server to which you redirect mailbox or public folder items.
- Microsoft Outlook 2007 or 2010 (32-bit only) must be installed on the destination server.

Note: The maximum size of the .PST file is 20 GB. If the restore exceeds the size limit, the data is spanned across multiple .PST files which are numbered consecutively.

Disaster recovery of an Exchange Server

A disaster preparation plan is an absolute necessity for restoring Exchange efficiently and effectively in the event of a catastrophic failure. Because Exchange uses Windows security for authentication, disaster recovery of Exchange cannot be separated from the disaster recovery of Windows.

Planning ahead reduces the time needed to recover.

It is critical to build a kit that includes the following items:

- An operating system configuration sheet
- A hard drive partition configuration sheet
- Any RAID configuration
- A hardware configuration sheet
- EISA/MCA configuration disks
- An Exchange configuration sheet
- A Windows emergency repair CD

To perform the actual recovery, you will need the following items:

- An installed copy of Backup Exec
- The latest full, incremental, and differential backups of the Exchange databases you want to recover
- The Microsoft Exchange Server Installation CD
- Any service packs that were applied to the original installation

You can use Simplified Disaster Recovery to recover the Exchange Server.

See [“Recovery notes for using Simplified Disaster Recovery with Exchange, SQL, SharePoint, CASO, Hyper-V hosts, and the Deduplication Option”](#) on page 771.

This procedure guides you through a complete restoration of Exchange using Backup Exec. You should have already performed all the appropriate preparation.

Always log on to Windows using the Administrator account (or an Administrator equivalent) during this procedure. Other requirements include:

- The storage groups and databases must already exist on the destination server, and have the same names as the original storage groups or databases.
- The destination server must have the same Organization and Administrative Group name as the source server.
- The destination databases must be configured so that they can be overwritten. See [“Restoring Exchange data”](#) on page 962.

To perform disaster recovery for Exchange

- 1 Recover the Windows server first.
Ensure that you restore the Exchange Server files that existed on all disk partitions.
- 2 From the Services applet, verify the Microsoft Exchange Information Store service is started.
- 3 Start Backup Exec.
- 4 Catalog the backup sets of the Exchange Server storage groups you want to recover.
See [“Backup sets ”](#) on page 342.
- 5 Run the Restore Wizard and select the latest full backup set of each storage group or database for restore.
See [“Methods for restoring data in Backup Exec”](#) on page 232.
- 6 If necessary, select all subsequent incremental storage group backup sets.
If differential backup sets are to be restored, only the most recent differential storage group backup set needs to be selected.
- 7 After completing the restore, it is recommended that a full backup of the restored databases be performed.

Symantec Backup Exec Agent for Microsoft SharePoint

This appendix includes the following topics:

- [About the Agent for Microsoft SharePoint](#)
- [About installing the Agent for Microsoft SharePoint](#)
- [Requirements for the Agent for Microsoft SharePoint](#)
- [Using the Agent for Microsoft SharePoint with SharePoint Server 2010/2013 and SharePoint Foundation 2010/2013](#)
- [Using the Agent for Microsoft SharePoint with SharePoint Server 2007 and Windows SharePoint Services 3.0](#)
- [Adding a Microsoft SharePoint server farm to the list of servers on the Backup and Restore tab](#)
- [Backing up Microsoft SharePoint data](#)
- [Restoring Microsoft SharePoint data](#)
- [Disabling or enabling communication between a Microsoft SharePoint web server and Backup Exec](#)
- [Viewing or changing SharePoint farm properties](#)
- [Disaster recovery of Microsoft SharePoint 2010/2013 data](#)
- [Disaster recovery of Microsoft SharePoint 2007 data](#)

About the Agent for Microsoft SharePoint

The Agent for Microsoft SharePoint (SharePoint Agent) is an optional, add-on component to Backup Exec. The SharePoint Agent enables network administrators to perform backup and restore operations on any supported Microsoft SharePoint installations that are connected to a network. SharePoint backups can be integrated with network backups without separate administration or dedicated hardware.

For specific platforms that the SharePoint Agent supports, see the Backup Exec Software Compatibility List at the following URL:

<http://entsupport.symantec.com/umi/v-269-1>

See “Using the Agent for Microsoft SharePoint with SharePoint Server 2010/2013 and SharePoint Foundation 2010/2013” on page 972.

See “Using the Agent for Microsoft SharePoint with SharePoint Server 2007 and Windows SharePoint Services 3.0” on page 973.

About installing the Agent for Microsoft SharePoint

Before you can back up Microsoft SharePoint server farms, you must install the Agent for Microsoft SharePoint (SharePoint Agent) on the Backup Exec server. The SharePoint Agent is installed as part of the Agent for Applications and Databases.

See “Installing additional Backup Exec options to the local Backup Exec server” on page 76.

See “Push-installing Backup Exec to remote computers” on page 78.

See “Requirements for the Agent for Microsoft SharePoint” on page 971.

Requirements for the Agent for Microsoft SharePoint

The Agent for Microsoft SharePoint (SharePoint Agent) has the following requirements:

- The SharePoint Agent must be installed on the Backup Exec server.
- The Agent for Windows must be installed on each remote SharePoint Server that you want to protect. In addition, the Agent for Windows must be installed on all servers in the server farm.
- You must use a logon account that has local administrative rights to back up and restore SharePoint data. The account should have local administrative rights on the servers where the SharePoint components are installed.

- To back up and restore the Single Sign-on database on SharePoint Server 2007, you must use a logon account with the appropriate credentials. The credentials must be either the account name or a member of the group that is specified in the Single Sign-on Settings in SharePoint.
 For more information on Single Sign-on Settings, see your SharePoint Server documentation.
- The logon account that you use to restore content into an existing site collection must have appropriate rights to create objects in that site collection. If you restore into a site collection that does not exist, the logon account becomes the primary site collection owner.
- Internet Information Services (IIS) rights can affect database backups and restores. Ensure that the logon account that you use for backup and restore has rights to access the IIS sites. Integrated Windows Security should be enabled within the IIS rights.

For specific operating system requirements for the SharePoint Agent, see the Backup Exec Software Compatibility List at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Using the Agent for Microsoft SharePoint with SharePoint Server 2010/2013 and SharePoint Foundation 2010/2013

The Agent for Microsoft SharePoint includes support for Microsoft Office SharePoint Server 2010/2013 and SharePoint Foundation 2010/2013.

SharePoint Server offers metadata features including tags, social bookmarks, and content ratings. These types of metadata are stored in the service applications that reside outside of the content database. For example, enterprise managed tags reside in the Managed Metadata Service application. You should back up all of your service applications to ensure that all metadata is protected.

Any metadata that is stored outside of the content database cannot be restored using Granular Recovery Technology (GRT). You can, however, use GRT to restore SharePoint data with metadata attached to it. As long as the metadata resides in the same service application, SharePoint maintains the link between the data and the metadata.

You can back up and restore the following types of SharePoint Server data:

- Web applications and their associated databases
- Individual documents and any pictures that are contained in libraries

- Sites and subsites
Individual objects and their versions can be restored from full database backups.
- Lists and list items
Individual objects and their versions can be restored from full database backups.
- Configuration database
A configuration database contains all of the configuration information for the entire SharePoint Server farm. Use caution when you restore this database. Any changes that you make to the farm topology before you restore from the backup are lost. The configuration database can be restored only to its original location.
- Service applications

See [“Disaster recovery of Microsoft SharePoint 2010/2013 data”](#) on page 985.

Using the Agent for Microsoft SharePoint with SharePoint Server 2007 and Windows SharePoint Services 3.0

The Agent for Microsoft SharePoint includes support for Microsoft Office SharePoint Server 2007 and Windows SharePoint Services 3.0.

You can back up and restore the following types of SharePoint Server 2007 and Windows SharePoint Services 3.0 data:

- Web applications and their associated databases
Symantec recommends that you restore all web application databases together to preserve the topology.
- Individual documents that are contained in libraries
- Sites and subsites
Individual objects and their versions can be restored from full database backups.
- Lists and list items
Individual items and their versions can be restored from full database backups.
- Configuration database
A configuration database contains all of the configuration information for the entire SharePoint Server farm. Use caution when you restore this database. Any changes that you make to the farm topology before you restore from the backup are lost. The configuration database can only be restored to its original location.
- Single Sign-on databases

Single Sign-on databases can be restored only to their original locations.

- **Shared Service Providers**
A Shared Service Provider is a grouping of shared services and related shared resources.

See [“Disaster recovery of Microsoft SharePoint 2007 data”](#) on page 989.

Adding a Microsoft SharePoint server farm to the list of servers on the **Backup and Restore** tab

Before you can back up Microsoft SharePoint data, you must add a SharePoint server farm to the list of servers on the **Backup and Restore** tab. If you add a single SharePoint server, Backup Exec adds the entire farm to which it belongs.

To add a Microsoft SharePoint server farm to the list of servers on the **Backup and Restore** tab

- 1 On the **Backup and Restore** tab, in the **Servers and Virtual Hosts** group, click **Add**.
- 2 Do one of the following:

To add a single SharePoint server	Click Microsoft Windows computers and servers .
To add a SharePoint server farm	Click Microsoft SharePoint server farm .
- 3 Click **Next**.
- 4 Follow the **Add a Server** wizard prompts to add the SharePoint server farm or server to the list of servers in the **Backup and Restore** tab.

See [“About the list of servers on the Backup and Restore tab”](#) on page 152.

Backing up Microsoft SharePoint data

The Agent for Microsoft SharePoint enables network administrators to perform backup operations on any Microsoft SharePoint installations that are connected to a network. SharePoint backups can be integrated with network backups without separate administration or dedicated hardware.

For more information about the specific types of SharePoint content that you can back up, see the following topics:

See [“Using the Agent for Microsoft SharePoint with SharePoint Server 2010/2013 and SharePoint Foundation 2010/2013”](#) on page 972.

See [“Using the Agent for Microsoft SharePoint with SharePoint Server 2007 and Windows SharePoint Services 3.0”](#) on page 973.

Backup Exec's dynamic inclusion feature automatically protects any new resources that were added after a backup job was created. If Backup Exec discovers that you added a new resource as a child to a protected resource, it automatically backs up the new resource. Because the backup job may include new resources, the job may require more storage space and more time to run than you anticipated.

You can set backup job default options for all SharePoint backup jobs. Each time you create a backup job, the job uses the default options unless you change the options for that particular job.

Note: Many backup options are available that are not described in the following procedure. For information on more backup options and methods, see the following topic:

See [“Backing up data”](#) on page 158.

To back up Microsoft SharePoint data

- 1 On the **Backup and Restore** tab, right-click the SharePoint server or farm that you want to back up.

To back up multiple servers, Shift + click or Ctrl + click the server names, and then right-click one of the selected servers.
- 2 Select **Backup**, and then select the type of backup that you want to perform.
- 3 On the **Backup Definition Properties** dialog box, in the **Selections** box, click **Edit**.
- 4 On the **Backup Selections** dialog box, check the resources that you want to back up and uncheck the resources that you do not want to back up.
- 5 Click **OK**.
- 6 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**.
- 7 On the **Backup Options** dialog box, in the left pane, click **Schedule**, and then select the schedule for this job.
- 8 On the Backup Options dialog box, in the left pane, click **Storage**, and then select the storage device that you want to use for the backup job.
- 9 On the **Backup Options** dialog box, in the left pane, click **Microsoft SharePoint**.

10 Set any of the following options for this job:

- | | |
|--|---|
| Perform a consistency check before the backup of any Microsoft SQL databases used by Microsoft SharePoint | Choose if you want to run a full consistency check (including indexes) of the Microsoft SQL databases that Microsoft SharePoint uses before you back up the databases. |
| Continue with the backup if consistency check fails | Choose to continue with the backup operation even if the consistency check fails. |
| Backup method | <p>Specify one of the following backup methods that you want to use for this job:</p> <ul style="list-style-type: none">■ Full - Back up databases
Backs up the entire database.■ Full Copy - Back up databases (copy)
This option backs up the entire database without affecting future differential or log backups.
Unlike the Full backup method, the Full Copy backup method does not reset the differential baseline that is used to indicate the database blocks that have changed since the last full backup.
After making a full backup, you can use the Full Copy backup method to make a copy of a database without affecting the baseline backup set required to run future differential backups. |
| Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual items from the database backup | <p>Choose if you want to enable the restore of individual documents, images, sites, subsites, lists, and list items from database backups.</p> <p>You must have a current version of the Agent for Windows on the SharePoint server when you run the GRT-enabled backup job.</p> |

Backup method

Specify one of the following SharePoint-specific backup methods that you want to use for this job:

- **Full - Back up databases**
Backs up the entire database.
- **Full Copy - Back up databases (copy)**
This option backs up the entire database without affecting future differential or log backups.
Unlike the Full backup method, the Full Copy backup method does not reset the differential baseline that is used to indicate the database blocks that have changed since the last full backup. After making a full backup, you can use the Full Copy backup method to make a copy of a database without affecting the baseline backup set required to run future differential backups.
- **Differential - Back up database changes since the last full**
Backs up only the changes made to the database since the last full backup.
- **Differential (block-level) - Back up database changes since the last full**
This option backs up all of the blocks of data and transaction logs that have been created or modified since the last full backup.
- **Incremental (block-level) - Back up database changes since the last full or incremental**
This option backs up all of the blocks of data and transaction logs that have been created or modified since the last full or incremental backup.
- **Log - Back up and truncate transaction log**
Backs up the data that is contained in the transaction log. This method does not back up database data. After the transaction log is backed up, committed transactions are removed (truncated).
If the databases are configured for the SQL Server simple recovery model, log backups are not supported. To change the recovery model, use the SQL administration tools to set the recovery model to Full. You should run a new full backup if you change the recovery model before a log backup is run.

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual items from the database backup For this job, choose if you want to enable the restore of individual documents, images, sites, subsites, lists, and list items from database backups. You must have a current version of the Agent for Windows on the SharePoint server when you run the GRT-enabled backup job.

See [“Granular Recovery Technology”](#) on page 590.

See [“Configuring delayed catalog operations to improve backup performance”](#) on page 536.

- 11 On the **Backup Options** dialog box, click any of the optional settings in the left pane that you want to set for this job.
 - 12 Click **OK**.
 - 13 In the **Backup Definition Properties** dialog box, click **OK**.
- See [“Changing default backup job settings”](#) on page 515.
- See [“Backing up data”](#) on page 158.
- See [“Editing backup definitions”](#) on page 206.

Setting default backup options for SharePoint

You can use the defaults that Backup Exec sets during installation for all SharePoint jobs, or you can choose your own defaults. You can override the default settings when you create individual jobs.

To set default backup options for SharePoint

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
 - 2 Select **Job Defaults**, and then select a backup option.
- For example, if you want to set up the default options for SharePoint backups to disk, select Back Up to Disk. The options that appear vary depending on what types of storage devices you configure. You can configure different default options for backup jobs that you send to different types of storage.
- 3 In the left pane, select **Microsoft SharePoint**.
 - 4 Select the appropriate options.

Perform a consistency check before the backup of any Microsoft SQL databases used by Microsoft SharePoint	Choose to run a full consistency check (including indexes) of the Microsoft SQL databases that Microsoft SharePoint uses before you back up the databases.
Continue with the backup if consistency check fails	Choose to continue with the backup operation even if the consistency check fails.
Backup method	<p>Specify one of the following backup methods that you want to use for this job:</p> <ul style="list-style-type: none">■ Full - Back up databases Backs up the entire database.■ Full Copy - Back up databases (copy) This option backs up the entire database without affecting future differential or log backups. Unlike the Full backup method, the Full Copy backup method does not reset the differential baseline that is used to indicate the database blocks that have changed since the last full backup. After making a full backup, you can use the Full Copy backup method to make a copy of a database without affecting the baseline backup set required to run future differential backups.
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual items from the database backup	Choose if you want to enable the restore of individual documents, images, sites, subsites, lists, and list items from database backups. You must have a current version of the Agent for Windows on the SharePoint server when you run the GRT-enabled backup job.

Backup method

Specify one of the following SharePoint-specific backup methods that you want to use for this job:

- **Full - Back up databases**
Backs up the entire database.
- **Full Copy - Back up databases (copy)**
This option backs up the entire database without affecting future differential or log backups.
Unlike the Full backup method, the Full Copy backup method does not reset the differential baseline that is used to indicate the database blocks that have changed since the last full backup.
After making a full backup, you can use the Full Copy backup method to make a copy of a database without affecting the baseline backup set required to run future differential backups.
- **Differential - Back up database changes since the last full**
Backs up only the changes made to the database since the last full backup.
- **Differential (block-level)**
This option backs up all of the blocks of data and transaction logs that have been created or modified since the last full backup.
- **Incremental (block-level)**
This option backs up all of the blocks of data and transaction logs that have been created or modified since the last full or incremental backup.
- **Log - Back up and truncate transaction log**
Backs up the data that is contained in the transaction log. This method does not back up database data. After the transaction log is backed up, committed transactions are removed (truncated).
If the databases are configured for the SQL Server simple recovery model, log backups are not supported. To change the recovery model, use the SQL administration tools to set the recovery model to Full. You should run a new full backup if you change the recovery model before a log backup is run.

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual items from the database backup For this job, choose if you want to enable the restore of individual documents, images, sites, subsites, lists, and list items from database backups. You must have a current version of the Agent for Windows on the SharePoint server when you run the GRT-enabled backup job.

See [“Granular Recovery Technology”](#) on page 590.

See [“Configuring delayed catalog operations to improve backup performance”](#) on page 536.

5 Click **OK**.

See [“Backing up Microsoft SharePoint data”](#) on page 974.

Restoring Microsoft SharePoint data

The Agent for Microsoft SharePoint lets you restore the following types of Microsoft SharePoint data to the original location or you can redirect the restore to a new location:

- SharePoint individual items such as documents, images, sites, subsites, lists, and list items.
- SharePoint web applications or portal sites and their associated content.
- SharePoint farm components such as configuration databases, service applications, shared service providers, or other components.
- SharePoint individual items such as documents, sites, and list items located through search.

You should keep in mind the following things when you restore SharePoint data:

- When you restore SharePoint data, individual SharePoint documents are always restored to SharePoint document libraries as checked out.
The documents are checked out using the same credentials as the logon account that you use for the restore. The documents must be checked in or published by that user before other users can use them.
- If you try to restore over a document that is checked out, the restore may fail. The restore fails if the document is checked out to a user that differs from the logon account credentials that are used for the restore.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

Restore individual items with Granular Recovery Technology (GRT)

Backup Exec also lets you restore individual documents, images, sites, subsites, lists, and list items from SharePoint database backups. To restore individual items from SharePoint database backups, ensure that the following Microsoft SharePoint option is selected during the backup job:

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual objects from the database backup

Note: Backup Exec does not support the restore of granular items from compressed databases or encrypted databases.

Redirected restores

Backup Exec lets you redirect the restore of the following:

- Redirect SharePoint file-based data, such as documents and images that have been uploaded to a document library or are attached to list items, to an NTFS file system for Windows SharePoint Services 3.0 and SharePoint Server 2007/2010, and SharePoint Foundation 2010.
- Redirect the restore of individual items such as documents, sites, and list items from one site to another site.

When you redirect a restore from one site to another site, the restored items inherit the security permissions of the parent item to which they are restored. If the site is in a different SharePoint farm, the versions of SharePoint must be the same.

- Restore SharePoint databases to an alternate SQL instance.

You can use the database to do the following:

- Manually harvest data using the SharePoint Central Administration console in SharePoint 2010/2013.
 - Manually attach to a SharePoint web application.
- Redirect SharePoint web application content databases to an alternate web application.

The destination for the web application must be online and be of the same topology as the web application that was backed up. If the web application is in a different SharePoint farm, the versions of SharePoint must be the same.

Caution: When you restore SharePoint Portal document library data, any documents that exist in the selected destination and that have the same name as the documents being restored may be overwritten. You can select whether they should be overwritten in the properties for the restore job.

For more information about the different types of SharePoint content, see the Microsoft SharePoint documentation.

To restore Microsoft SharePoint data

- 1 On the **Backup and Restore** tab, right-click the server for which you want to restore data, and then click **Restore**.
- 2 Select **Microsoft SharePoint**, and then click **Next**.
- 3 Follow the **Restore Wizard** prompts to restore the data.

See [“Using the Agent for Microsoft SharePoint with SharePoint Server 2010/2013 and SharePoint Foundation 2010/2013”](#) on page 972.

See [“Using the Agent for Microsoft SharePoint with SharePoint Server 2007 and Windows SharePoint Services 3.0”](#) on page 973.

Disabling or enabling communication between a Microsoft SharePoint web server and Backup Exec

Backup Exec communicates with the web servers that participate in Microsoft SharePoint server farms to discover the farm topology. This process may take some time if Backup Exec attempts to communicate with a web server that is unavailable. If you know that a particular web server in a farm is unavailable for a period of time, you can disable the communication between the web server and Backup Exec.

To disable or enable communication between a SharePoint web server and Backup Exec

- 1 On the **Backup and Restore** tab, double-click the SharePoint server farm to which the web server belongs.
- 2 In the left pane, click **Properties**.
- 3 Do one of the following:
 - To prevent Backup Exec from communicating with a SharePoint web server, clear the check box next to the web server name.
 - To let Backup Exec communicate with the SharePoint web server, select the check box next to the web server name.
- 4 Click **Apply**.

Viewing or changing SharePoint farm properties

You can view properties for any SharePoint farm that you monitor with Backup Exec, and you can change some properties. Backup Exec displays general and system information about the farm.

You can also enable or disable communication between a SharePoint web server and Backup Exec from the SharePoint farm properties dialog box.

See [“Disabling or enabling communication between a Microsoft SharePoint web server and Backup Exec”](#) on page 983.

To view or change SharePoint farm properties

- 1 On the **Backup and Restore** tab, double-click the SharePoint farm whose properties you want to view.
- 2 In the left pane, click **Properties**.
- 3 View or change the appropriate properties:

SharePoint farm name	Displays the name of the SharePoint farm.
-----------------------------	---

Description	Lets you enter a unique description to identify the farm in Backup Exec. The description is optional.
--------------------	---

Logon account	Lists the logon account that Backup Exec uses to access the farm.
----------------------	---

Click **Add/Edit** to add a new logon account or to edit an existing logon account.

Web Servers	Lists the web servers that belong to the farm. You can enable or disable communication between the web servers and Backup Exec.
--------------------	---

See [“Disabling or enabling communication between a Microsoft SharePoint web server and Backup Exec”](#) on page 983.

- 4 If you made changes, click **Apply**.

Disaster recovery of Microsoft SharePoint 2010/2013 data

You can use the Agent for Microsoft SharePoint to recover a Microsoft SharePoint server after a hard drive failure. Before you recover SharePoint data, you must recover the SharePoint server's operating system.

You can use Backup Exec's Simplified Disaster Recovery Option or you can manually recover the server's operating system.

For more information about manual recovery, click the following URL:

<http://entsupport.symantec.com/umi/V-290-706>

See “[About Simplified Disaster Recovery](#)” on page 736.

After the Windows server is recovered, you can recover SharePoint data. Complete the actions in the table in successive order to recover the SharePoint data.

Table G-1 To recover SharePoint data after a disaster

Step	Action	Notes
Step 1	Recover the master database and the model database for the SQL instances that SharePoint uses. You must perform this step if you manually recovered the server's operating system.	Skip this step if you used Backup Exec's Simplified Disaster Recovery Option to recover the server's operating system. See “ Manual recovery of a SQL Server ” on page 928.
Step 2	Inventory the media to be recovered.	See “ Inventorying a storage device ” on page 453.
Step 3	Catalog the media to be recovered.	See “ Cataloging a storage device ” on page 452.

Table G-1 To recover SharePoint data after a disaster (*continued*)

Step	Action	Notes
Step 4	Restore the msdb databases for any SQL instances that SharePoint uses.	<p>Select the backup sets that contain the msdb databases for the SQL instances that SharePoint uses.</p> <p>Configure the following Microsoft SQL restore options:</p> <ul style="list-style-type: none">■ Use the default settings for all restore jobs.■ Select Overwrite existing databases. <p>See “Methods for restoring data in Backup Exec” on page 232.</p>
Step 5	Restore all web applications.	<p>Select the backup sets for all SharePoint web applications.</p> <p>Select Yes, restore over existing databases.</p> <p>See “Methods for restoring data in Backup Exec” on page 232.</p>

Table G-1 To recover SharePoint data after a disaster (*continued*)

Step	Action	Notes
Step 6	Restore Shared Services Applications databases.	<p>Restore the following Shared Services Applications databases:</p> <ul style="list-style-type: none">■ Business Data Connectivity Service■ Managed Metadata Service■ PerformancePoint Service Application■ Search Service Application■ Secure Store Service■ User Profile Service Application■ Web Analytics Service Application■ Word Automation Services■ Services\State Services\Service DB 1 <p>Select Yes, restore over existing databases.</p> <p>See “Methods for restoring data in Backup Exec” on page 232.</p> <p>Note: Some of the remaining restore jobs may fail because communication with the SharePoint server has not been fully established yet. This behavior is expected. Proceed with the recovery process until all steps are complete.</p>

Table G-1 To recover SharePoint data after a disaster (*continued*)

Step	Action	Notes
Step 7	Restore search services.	<p>Restore the following services:</p> <ul style="list-style-type: none">■ SharePoint Foundation Help Search\Search Instance\Index Files 1■ Search-DB 1 <p>Select Yes, restore over existing databases.</p> <p>See "Methods for restoring data in Backup Exec" on page 232.</p> <p>Note: You may get a message in the job log to restart your computer. You can ignore the message.</p>
Step 8	Restore the SharePoint Configuration V4/V5-DB resource.	<p>Select the backup sets for the SharePoint Configuration V4-DB resource.</p> <p>Select Yes, restore over existing databases.</p> <p>See "Methods for restoring data in Backup Exec" on page 232.</p>
Step 9	Restart the SharePoint server.	<p>After the restore job is complete, restart the SharePoint server. Then proceed to the next step.</p>
Step 10	Restore remaining SharePoint resources.	<p>Select the backup sets for the SharePoint Global Settings resources, if necessary.</p> <p>See "Methods for restoring data in Backup Exec" on page 232.</p>

Table G-1 To recover SharePoint data after a disaster (*continued*)

Step	Action	Notes
Step 11 (For SharePoint 2013 multi-server farms)	Ensure that the servers in the SharePoint farm are configured correctly after the recovery.	Run the SharePoint Products Configuration Wizard.
Step 12	Back up the SharePoint server.	When the disaster recovery is complete, Symantec recommends that you perform a backup job as soon as possible. See “Backing up data” on page 158.

Disaster recovery of Microsoft SharePoint 2007 data

You can use the Agent for Microsoft SharePoint to recover a Microsoft SharePoint server after a hard drive failure. Before you recover SharePoint data, you must recover the SharePoint server's operating system.

You can use Backup Exec's Simplified Disaster Recovery Option or you can manually recover the server's operating system.

<http://entsupport.symantec.com/umi/V-290-706>

See [“About Simplified Disaster Recovery”](#) on page 736.

After the Windows server is recovered, you can recover SharePoint 2007 data. Complete the actions in the table in successive order to recover the SharePoint data.

Table G-2 To recover SharePoint 2007 data after a disaster

Step	Action	Notes
Step 1	Recover the master database and the model database for the SQL instances that SharePoint uses. You must perform this step if you manually recovered the server's operating system.	Skip this step if you used Backup Exec's Simplified Disaster Recovery Option to recover the server's operating system. See “Manual recovery of a SQL Server” on page 928.

Table G-2 To recover SharePoint 2007 data after a disaster (*continued*)

Step	Action	Notes
Step 2	Inventory the media to be recovered.	See "Inventorying a storage device" on page 453.
Step 3	Catalog the media to be recovered.	See "Cataloging a storage device" on page 452.
Step 4	Restore the msdb databases for any SQL instances that SharePoint uses.	<p>Select the backup sets that contain the msdb databases for the SQL instances that SharePoint uses.</p> <p>Configure the following Microsoft SQL restore options:</p> <ul style="list-style-type: none">■ Select Leave the database ready to use; additional transaction logs or differential backups cannot be restored.■ Select Overwrite existing databases. <p>See "Methods for restoring data in Backup Exec" on page 232.</p>
Step 5	Restore SharePoint farm components.	<p>Select the backup sets for the following SharePoint components:</p> <ul style="list-style-type: none">■ Help Search service■ WSS Administration■ Shared Service web application (if applicable) <p>Select Yes, restore over existing databases.</p> <p>See "Methods for restoring data in Backup Exec" on page 232.</p>

Table G-2 To recover SharePoint 2007 data after a disaster (*continued*)

Step	Action	Notes
Step 6	Restore SharePoint web applications.	Select the backup sets for any SharePoint web applications. See “Methods for restoring data in Backup Exec” on page 232.
Step 7	Restore the SharePoint Configuration V3-DB resource.	Select the backup sets for the SharePoint Configuration V3-DB resource. Select Yes, restore over existing databases . See “Methods for restoring data in Backup Exec” on page 232.
Step 8	Restart the SharePoint server.	After the restore job is complete, restart the SharePoint server. Then proceed to the next step.
Step 9	Restore remaining SharePoint resources.	Select the backup sets for the SharePoint Global Settings and Single Sign-on database resources, if necessary. See “Methods for restoring data in Backup Exec” on page 232.
Step 10	Back up the SharePoint server.	When the disaster recovery is complete, Symantec recommends that you perform a backup job as soon as possible. See “Backing up data” on page 158.

Symantec Backup Exec Agent for Oracle on Windows or Linux Servers

This appendix includes the following topics:

- [About the Backup Exec Oracle Agent](#)
- [About installing the Oracle Agent](#)
- [Configuring the Oracle Agent on Windows computers and Linux servers](#)
- [About authentication credentials on the Backup Exec server](#)
- [About Oracle instance information changes](#)
- [About backing up Oracle databases](#)
- [About restoring Oracle resources](#)

About the Backup Exec Oracle Agent

The Symantec Backup Exec Agent for Oracle on Windows or Linux Servers (Oracle Agent) uses Oracle's Recovery Manager (RMAN) to protect Oracle databases. RMAN is a tool that manages the backup and restore and recovery of Oracle databases.

The following features are available with the Oracle Agent:

- The ability to initiate backup and restore operations from Backup Exec or from the RMAN console as a Database Administrator (DBA).

Operations that the DBA performs on the RMAN console are referred to as DBA-initiated operations. You should refer to your Oracle documentation for information about RMAN.

- Multiple data stream support for increased performance during backup and restore operations.
- RMAN recovery catalog support to manage the backup, restore, and recovery of Oracle databases.
- Oracle Real Application Cluster (RAC) support.

The following are not supported:

- Tivoli Storage Manager (TSM) devices as storage for Oracle backup jobs.
- The Oracle Management Server.

See [“About installing the Oracle Agent”](#) on page 993.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 994.

About installing the Oracle Agent

The Oracle Agent is installed as part of the Agent for Applications and Databases and can protect local or remote Oracle instances.

To protect local or remote Oracle instances, you must install the following Backup Exec options:

- Backup Exec Agent for Windows on remote Windows computers.

Note: If you upgrade a previous version of the Agent for Windows on an Oracle server, you must restart the Oracle server after the upgrade. Backup Exec jobs cannot complete successfully until you restart the Oracle server.

See [“Methods for installing the Agent for Windows”](#) on page 86.

- Backup Exec Agent for Linux on remote Linux computers.
See [“About installing the Agent for Linux”](#) on page 1189.
- The Agent for Applications and Databases on the Backup Exec server.
See [“Installing additional Backup Exec options to the local Backup Exec server”](#) on page 76.

Configuring the Oracle Agent on Windows computers and Linux servers

Before you can back up or restore Oracle databases, you must do the following:

Table H-1 Configuring the Oracle Agent on Windows computers and Linux servers

Step	Action
Step 1	<div>Configure information about the Oracle instances for the Oracle Agent.</div> <div>See “Configuring an Oracle instance on Windows computers” on page 995.</div> <div>See “Configuring an Oracle instance on Linux servers” on page 1001.</div>

Table H-1 Configuring the Oracle Agent on Windows computers and Linux servers (*continued*)

Step	Action
Step 2	<p>Enable database access for the Backup Exec server.</p> <p>Whenever Oracle instance information changes or a new configuration is added, you must update the Backup Exec Agent Utility. If credential information is not updated, is incorrect, or the server is down, the error "Unable to attach to a resource..." may appear when you run a backup job. If this message appears, you must bring the server online and configure the information.</p> <p>For Oracle RAC, run the Backup Exec Agent Utility on each node and add information about the instances. When Oracle RAC nodes are added or removed, you must enter information about any changes to instances in the Backup Exec Agent Utility.</p> <p>Note: When you use the Backup Exec Agent Utility, the user account with which you are logged on should be a member of the Oracle DBA group.</p> <p>You must have administrator privileges to run the Backup Exec Agent Utility.</p> <p>See "Enabling database access for Oracle operations on Windows computers" on page 999.</p> <p>See "Enabling database access for Oracle operations on Linux servers" on page 1004.</p>
Step 3	<p>Set authentication credentials for Oracle.</p> <p>See "About authentication credentials on the Backup Exec server" on page 1005.</p>

Configuring an Oracle instance on Windows computers

You can use the Backup Exec Agent Utility to configure information about the Oracle instances for the Oracle Agent on Windows computers.

To configure an Oracle instance on Windows computers

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec Agent Utility**.

When the Backup Exec Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 On the **Oracle** tab, click **New**.

Any instances that currently exist on the computer appear on the tab.

- 3 Complete the appropriate options.

See [“Oracle Agent Configuration options”](#) on page 996.

- 4 Click **OK**.

Oracle Agent Configuration options

You can set the following Oracle Agent Configuration options.

See [“Configuring an Oracle instance on Windows computers”](#) on page 995.

See [“Editing an Oracle instance on Windows computers”](#) on page 998.

Table H-2 Oracle Agent Configuration options

Item	Description
Local instance name	<p>Displays the name of the Oracle instance. If you edit an instance, you cannot change the instance name.</p> <p>For Oracle RAC nodes, enter the fully qualified domain name of each physical node.</p> <p>The fully qualified domain name of the node appears in the list of servers on the Backup and Restore tab.</p> <p>The name is in the format RAC-<dbname>-<dbid>, where dbname is the database name, and dbid is the database ID.</p>
User name	<p>Displays the user name for the Oracle instance.</p> <p>If the credentials for the Oracle instance change, you must enter a user with SYSDBA rights to the Oracle instance.</p> <p>For Oracle RAC nodes, enter the same set of credentials for all of the nodes.</p>
Password	<p>Displays the password for the Oracle instance user name.</p>

Table H-2 Oracle Agent Configuration options (*continued*)

Item	Description
Confirm password	Displays the password again to confirm it.
Use recovery catalog	Indicates that you plan to use the Oracle recovery catalog. The Oracle Agent supports the use of the RMAN recovery catalog to manage the backup, restore, and recovery of Oracle databases. If you choose not to use a recovery catalog, RMAN uses the source database control file as the sole repository of metadata.
TNS name	Displays the Oracle Net Service name.
User name	Displays the user name for the Oracle recovery catalog.
Password	Displays the password for the Oracle recovery catalog.
Confirm password	Displays the password for the recovery catalog again to confirm it.
Backup Exec server name or IP address	Displays the name or IP address of the Backup Exec server where you want to send the DBA-initiated backup jobs. You must use the same form of name resolution for all operations.
Job template name	Displays the name of the Backup Exec job template that you want the DBA-initiated job to use for backup and restore operations. You create the job template on the DBA-initiated Job Settings dialog box on the Backup Exec server. If you do not specify a job template, the default job template is used. See “DBA-initiated job templates” on page 597.

Viewing an Oracle instance on Windows computers

You can use the Backup Exec Agent Utility to view information about the Oracle instances for the Oracle Agent on Windows servers.

To view an Oracle instance on Windows computers

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec Agent Utility**.
- 2 On the **Oracle** tab, view the instances that currently exist on the computer.
See [“Oracle options for the Backup Exec Agent Utility”](#) on page 998.
- 3 Click **OK**.

Oracle options for the Backup Exec Agent Utility

You can set the following Oracle options for the Backup Exec Agent Utility.

See [“Viewing an Oracle instance on Windows computers”](#) on page 997.

Table H-3 Oracle options for the Backup Exec Agent Utility

Item	Description
Instance	Displays the name of the Oracle instance.
User Name	Displays the user name for the Oracle instance.
Recovery Catalog	Displays the name of the recovery catalog.
Backup Exec Server	Displays the name or IP address of the Backup Exec server where you want to send the DBA-initiated backup jobs.
Job Template	Displays the name of the DBA-initiated template. See “About performing a DBA-initiated backup job for Oracle” on page 1010.
New	Lets you add an Oracle instance.
Edit	Lets you revise an Oracle instance.
Delete	Lets you remove an Oracle instance.

Editing an Oracle instance on Windows computers

You can use the Backup Exec Agent Utility to revise information about the Oracle instances for the Oracle Agent on Windows computers.

To edit an Oracle instance on Windows computers

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec Agent Utility**.

When the Backup Exec Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.

- 2 On the **Oracle** tab, click **Edit**.

Any instances that currently exist on the computer appear on the tab.

- 3 Edit the appropriate options.
See [“Oracle Agent Configuration options”](#) on page 996.
- 4 Click **OK**.

Deleting an Oracle instance on Windows computers

You can use the Backup Exec Agent Utility to remove an Oracle instance for the Oracle Agent on Windows computers.

To delete an Oracle instance on Windows computers

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec Agent Utility**.

When the Backup Exec Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.
- 2 On the **Oracle** tab, click **Delete**.

Any instances that currently exist on the computer appear on the tab.
- 3 Click **OK**.

Enabling database access for Oracle operations on Windows computers

You can use the Backup Exec Agent Utility to enable database access for the Windows computer after you configure an Oracle instance.

See [“About backing up Oracle databases”](#) on page 1008.

See [“About backing up Oracle RAC databases”](#) on page 1009.

See [“DBA-initiated job templates”](#) on page 597.

See [“Changing default backup job settings”](#) on page 515.

To enable database access for Oracle operations on Windows computers

- 1 On the computer on which the Agent for Windows is installed, on the taskbar, click **Start > All Programs > Symantec Backup Exec > Backup Exec Agent Utility**.

When the Backup Exec Agent Utility is running, an icon appears in the system tray. You can double-click this icon to view the utility.
- 2 Click the **Database Access** tab.

(Optional) The first time that you start the Backup Exec Agent Utility, click **Change Settings** to enable the options.

3 Complete the appropriate options to configure database access:

Enable the Backup Exec server to authenticate Oracle operations	Select this option to enable Oracle operations between the Backup Exec server and this computer.
User name	<p>Specify a user name that has administrative rights to this computer. This logon account is what the Backup Exec server uses when it connects to this computer.</p> <p>If you specify an IP address or a fully qualified computer name as part of the user name, the Backup Exec Agent Utility may not be able to verify the user account. If the credentials entered are incorrect, the error “cannot attach to a resource” may be displayed when you run a backup or restore job.</p> <p>You must add this computer name and logon account to the Backup Exec server's list of authentication credentials for Oracle servers. If the authentication fails when the Oracle resources are backed up, the backup job fails. If the authentication fails when you are browsing the backup sets for a restore job, then the backup sets become unavailable, and you must run a DBA-initiated restore job to restore data.</p>
Password	<p>Specify the password for this logon account.</p> <p>Note: For security reasons, the logon credentials are not stored on the remote computer.</p>
Confirm Password	Type the password again to confirm it.
Use a custom port to connect to the Backup Exec server during Oracle operations	<p>Select this option to change the port that is used for communications between this computer and the Backup Exec server during Oracle operations. By default, port 5633 is used.</p> <p>If you change the port number on this computer, you must also change it on the Backup Exec server, and then restart the Backup Exec Job Engine Service on the Backup Exec server.</p>
Port number	Type the port number that you want to use for communications between this computer and the Backup Exec server.

4 Click **OK**.

- 5 For Oracle RAC installations, type the fully qualified domain name that you want to publish to.

The Backup Exec server that you publish to lists the RAC databases in the list of servers on the **Backup and Restore** tab.

If you do not enter a fully qualified domain name to publish to, the RAC databases are not in the list of servers.

See [“About publishing the Agent for Windows to Backup Exec servers”](#) on page 784.

- 6 On the Backup Exec server, add the name of the Oracle server and the user name that you enabled for database access to the Backup Exec server's list of authentication credentials.

See [“About authentication credentials on the Backup Exec server”](#) on page 1005.

Configuring an Oracle instance on Linux servers

You can use the Backup Exec Agent Utility to configure information about the Oracle instances for the Oracle Agent on Linux servers.

To configure an Oracle instance on Linux servers

- 1 On the Linux server on which the Oracle instances are installed, open a Terminal window.
- 2 Change to the following directory:
cd /opt/VRTSralus/bin
- 3 Start the Backup Exec Agent Utility:
./AgentConfig
- 4 Type **2** to select Configure Oracle instance information, and then press **Enter**.
- 5 Type **1** to select the Add a new Oracle Instance option, and then press **Enter**.
- 6 Enter the name of the Oracle instance in upper case characters.
For example, ORACLENAME.

- 7 Enter the user name for the Oracle instance.

If the credentials for the Oracle instance are changed, you must update the credentials in this field. For Oracle RAC nodes, enter the same set of credentials for all of the nodes.

When you use the Backup Exec Agent Utility to enter the Oracle credentials for an instance, the credentials cannot be verified if the user account with which you are logged on is a member of the Oracle DBA group. If the credentials are incorrect, the error "Unable to attach to a resource..." may appear when you run a backup job.

- 8 To display the Oracle database in a Backup Exec server's list of servers on the **Backup and Restore** tab, type the Backup Exec server name or IP address to which you want the remote computer to publish to.

- 9 When prompted, specify if you want to use a recovery catalog.

The Oracle Agent supports the use of the RMAN recovery catalog to manage the backup, restore, and recovery of Oracle databases. If you choose not to use a recovery catalog, RMAN uses the source database control file as the sole repository of metadata.

If you specify a recovery catalog, any database that you want to back up must be registered in the recovery catalog before you can run backup jobs from the Backup Exec server.

- 10 To use a recovery catalog, type the recovery catalog name and a user name and password for the recovery catalog.
- 11 To use a customized DBA-initiated job settings template, type the name of the template.

See ["DBA-initiated job templates"](#) on page 597.

- 12 Do one of the following:

To commit the new entry to the configuration file Type **Y**, and then press **Enter**.

To cancel this entry Type **N**, and then press **Enter**.

Viewing an Oracle instance on Linux servers

You can use the Backup Exec Agent Utility to view information about the Oracle instances for the Oracle Agent on Linux servers.

The following information is listed:

- Name of the instance
- Logon name for the instance
- IP address of the default Backup Exec server name for DBA-initiated operations
- Name of the DBA-initiated job template

To view an Oracle instance on Linux servers

- 1 On the Linux server on which the Oracle instances are installed, open a Terminal window.

- 2 Change to the following directory:

```
cd /opt/VRTSralus/bin
```

- 3 Start the Backup Exec Agent Utility:

```
./AgentConfig
```

- 4 Type 4.

Editing an Oracle instance on Linux servers

You can use the Backup Exec Agent Utility to revise information about the Oracle instances for the Oracle Agent on Linux servers.

To edit an Oracle instance on Linux computers

- 1 On the Linux server on which the Oracle instances are installed, open a Terminal window.

- 2 Change to the following directory:

```
cd /opt/VRTSralus/bin
```

- 3 Start the Backup Exec Agent Utility:

```
./AgentConfig
```

- 4 Type **2** to select Configure Oracle Instance Information, and then press **Enter**.

Any instances that currently exist on the computer are discovered.

- 5 Type **2**.

- 6 Follow the prompts.

Deleting an Oracle instance on Linux servers

You can use the Backup Exec Agent Utility to remove an Oracle instance for the Oracle Agent on Linux servers.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 994.

To delete an Oracle instance for the Oracle Agent on Linux servers

- 1 On the Linux server on which the Oracle instances are installed, open a Terminal window.
- 2 Change to the following directory:

```
cd /opt/VRTSralus/bin
```
- 3 Start the Backup Exec Agent Utility:

```
./AgentConfig
```
- 4 Type **2** to select Configure Oracle Instance Information, and then press **Enter**.
Any instances that currently exist on the computer are discovered.
- 5 Type **3**.
- 6 Follow the prompts.

Enabling database access for Oracle operations on Linux servers

You can use the Backup Exec Agent Utility to enable database access for the Linux server after you configure an Oracle instance.

See [“Setting authentication credentials on the Backup Exec server for Oracle operations”](#) on page 1006.

See [“About backing up Oracle databases”](#) on page 1008.

See [“About backing up Oracle RAC databases”](#) on page 1009.

See [“DBA-initiated job templates”](#) on page 597.

See [“Changing default backup job settings”](#) on page 515.

To enable database access for Oracle operations on Linux servers

- 1 On the Linux server on which the Oracle instances are installed, open a Terminal window.
- 2 Change to the following directory:

```
cd /opt/VRTSralus/bin
```
- 3 Start the Backup Exec Agent Utility:

```
./AgentConfig
```
- 4 Type **1** to select Configure database access, and then press **Enter**.

- 5 Type the user name that is in the beoper group on the Linux system.

See [“About the Backup Exec operators group for the Agent for Linux”](#) on page 1192.

If the authentication fails when the Oracle resources are backed up, the backup job fails. If the authentication fails when you browse the backup sets for a restore job, then the backup sets become unavailable, and you must run a DBA-initiated restore job to restore data.

- 6 Type the password for this logon account, and then confirm it.

The logon credentials are not stored on this computer.

- 7 When prompted, specify if you want to use a custom port to connect to the Backup Exec server communications between this computer and the Backup Exec server during Oracle operations.

Port 5633 is used by default. If you change the port number on this computer, you must also change it on the Backup Exec server, and then restart the Backup Exec Job Engine Service on the Backup Exec server. If a Windows firewall is enabled, you must add this port as an exception.

See [“Changing network and security options for Backup Exec”](#) on page 576.

- 8 Do one of the following:

To commit the Oracle operation settings to the configuration file	Type Y , and then press Enter .
---	---

To cancel this entry	Type N , and then press Enter .
----------------------	---

About authentication credentials on the Backup Exec server

You must add the Oracle fully qualified domain name and the logon account name to the Backup Exec server's list of Oracle servers and authentication credentials. The Backup Exec server has database access for operations on Oracle instances that are included in the authentication list. Before you start any backup or restore operations, on the computer on which the Oracle instances are installed, ensure that you use the Backup Exec Agent Utility to configure instance information and database access.

The logon account name must have administrative rights to the Oracle server. If the user name is incorrect or is not provided, or if it does not have the appropriate rights, then you cannot perform Oracle backup or restore operations to that computer.

Note: For Oracle RAC nodes, enter the fully qualified domain name for the logon account name. You can view the fully qualified domain name of the node in the list of servers on the **Backup and Restore** tab. It is in the form RAC-<database name>-<database ID>.

See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 994.

See [“Setting authentication credentials on the Backup Exec server for Oracle operations”](#) on page 1006.

See [“Deleting an Oracle server from the Backup Exec server's list of authentication credentials”](#) on page 1007.

Setting authentication credentials on the Backup Exec server for Oracle operations

You must add the Oracle server to the list so that the Backup Exec server has database access for operations.

See [“About authentication credentials on the Backup Exec server”](#) on page 1005.

See [“About Oracle instance information changes ”](#) on page 1008.

To set authentication credentials on the Backup Exec server for Oracle operations

- 1 Click the Backup Exec button, and then select **Configuration and Settings** and then click **Backup Exec Settings**.

- 2 In the left pane, click **Oracle**.

- 3 Enter the name of the Oracle server on which the instance is installed.

The name of the Oracle server should match the name of the server that lists the Oracle resource. Symantec recommends that you enter the fully qualified domain name. For example, Servername.domain.com is the fully qualified domain name and Servername is the NETBIOS name. For Oracle RAC nodes, enter the RAC-<database name>-<database ID> for each node in the list.

- 4 Click **Add**.

5 To add the logon account name, do one of the following:

Click the arrow Select the logon account name that you want to add.

Click <new logon account> Enter the appropriate options.

Use the same logon account format that you use when you enter the logon account name on the **Database Access** tab in the Backup Exec Agent Utility. For example, if you entered Domainname\Username on the Backup Exec Agent Utility, use that same format on the list of authentication credentials.

6 Click **OK**.

Oracle job settings options

You can add the Oracle fully qualified domain name and the logon account name to the Backup Exec server's list of Oracle servers and authentication credentials.

See [“About authentication credentials on the Backup Exec server”](#) on page 1005.

Table H-4 Authentication credentials for Oracle servers options

Item	Description
Server name	Displays the name of the Oracle server.
Logon account	Displays the name of the logon account that has rights to the Oracle server.
Add	Lets you add the fully qualified domain name and logon account credentials to the list.
Delete	Lets you remove the fully qualified domain name and logon account credentials.

Deleting an Oracle server from the Backup Exec server's list of authentication credentials

You can delete an Oracle server name or logon account from a Backup Exec server's list of authentication credentials.

See [“About authentication credentials on the Backup Exec server”](#) on page 1005.

To delete an Oracle server from the Backup Exec server's list of authentication credentials

- 1 Click the Backup Exec button, and then select **Configuration and Settings** and then click **Backup Exec Settings**.
- 2 In the left pane, click **Oracle**.
- 3 Select the item that contains the server name or logon account that you want to delete.
- 4 Click **Delete**.
- 5 Click **OK**.

About Oracle instance information changes

Whenever information about the Oracle instance changes, such as the instance user name and password, you must update the Backup Exec Agent Utility.

When Oracle RAC nodes are added or removed, you must enter information about any changes to instances in the Backup Exec Agent Utility. After these changes are entered, the Backup Exec server discovers them.

If the changes are not updated in the Backup Exec Agent Utility, the error "Unable to attach to a resource..." may appear when you run a backup job.

See ["Configuring the Oracle Agent on Windows computers and Linux servers"](#) on page 994.

About backing up Oracle databases

Before you back up Oracle databases, review the following:

- You must run the Backup Exec Agent Utility on the Oracle server and add information about the instances before you can perform any backup or restore operations.

When Oracle instance information changes, you must update the Backup Exec Agent Utility. After these changes are entered, the Backup Exec server discovers them.

See ["Configuring the Oracle Agent on Windows computers and Linux servers"](#) on page 994.

- During a backup operation, the amount of data that is backed up may not equal the size of the Oracle files that are on the disk. This behavior is normal. Backup Exec backs up the selected data files as well as a copy of the control file.

- In a Central Admin Server Option environment, all backup jobs for a specific Oracle instance must be delegated to the same managed Backup Exec server. If you do not restrict the backup job to the same managed Backup Exec server, then before you can restore data, you must move the physical media that contains the backup sets to a single managed Backup Exec server.
See [“Selecting a Backup Exec server pool for backups”](#) on page 1137.

- If the Oracle database resides on the volumes that are configured with Oracle Automatic Storage Management (ASM), you cannot select these volumes as part of a file system backup.

The following message appears when you attempt to select the volumes:

```
An error was encountered while attempting to browse the
contents of <drive>. A device-specific error occurred.
```

- The database must be in a mounted or an open state before you can make backup selections.
- The database must be in ARCHIVELOG mode before the **Archived Logs** node can be displayed under the Oracle resource of the Oracle server.

You can add an Oracle database to the list of servers on the **Backup and Restore** tab so that the database can be selected for backup jobs. You can set backup job default options for all Oracle backup jobs. Each time you create a backup job, the job uses the default options unless you change the options for the particular job.

See [“About the list of servers on the Backup and Restore tab”](#) on page 152.

See [“Backing up data”](#) on page 158.

See [“Oracle backup options”](#) on page 1011.

See [“About backing up Oracle RAC databases”](#) on page 1009.

See [“About performing a DBA-initiated backup job for Oracle”](#) on page 1010.

About backing up Oracle RAC databases

Oracle Real Application Cluster (RAC) is an active and an active cluster with shared storage, in which multiple instances share a single physical database. Since all of the participating nodes can access the database, you can initiate backup, restore, or recovery from any node.

Oracle RAC Database resource container is added automatically to the Backup and Restore tab after

- You install the Remote Agent for Windows and configure the Backup Exec Agent Utility on a Windows Oracle server.

- You install the Remote Agent for Linux and UNIX and configure the AgentConfig utility on a Linux or Unix Oracle server.

Requirements for backing up Oracle RAC resources include the following items:

- You must run the Backup Exec Agent Utility on each node and add information about the instances before you can perform any backup or restore operations. When RAC nodes are added or removed, you must update the Backup Exec Agent Utility with information about the affected instances. After these changes are entered, the Backup Exec server discovers them.
See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 994.
- You must select the RAC fully qualified domain name when making backup selections.
Each node in the cluster uses the same fully qualified domain name. The fully qualified domain name of the node appears in the list of servers on the **Backup and Restore** tab. It is in the form RAC-<database name>-<database ID>.

Backing up Oracle RAC is similar to backing up standard Oracle databases.

You should be aware of the following differences:

- By default, each node in an Oracle RAC stores its archive logs locally. To have a meaningful backup of the archive logs, back up each archive log. Alternatively, you can move the archive logs to shared storage for backup.
- Each node that is part of the cluster is assigned a priority. For database backups, Backup Exec connects to the node that has the highest priority. Backup Exec uses the fully qualified domain name to connect to the node.

See [“About the list of servers on the Backup and Restore tab”](#) on page 152.

See [“Backing up data”](#) on page 158.

See [“Oracle backup options ”](#) on page 1011.

See [“About backing up Oracle databases”](#) on page 1008.

See [“About performing a DBA-initiated backup job for Oracle”](#) on page 1010.

About performing a DBA-initiated backup job for Oracle

A Database Administrator (DBA) can initiate a backup or restore operation for Oracle from the RMAN console. Example scripts for backup and restore operations that you can run from the RMAN console are installed to the following location:

```
\Program Files\Symantec\Backup Exec\scripts\Oracle
```

Refer to your Oracle documentation for more information on using the RMAN console.

Review the following notes before initiating backup jobs for Oracle from the RMAN console:

- Ensure that you have completed all of the preparations for configuring the Oracle Agent.
See [“Configuring the Oracle Agent on Windows computers and Linux servers”](#) on page 994.
- The channel is not released if the RMAN console is not exited, or if a new manual channel is not allocated on that console.
- The SKIP INACCESSIBLE option is available in RMAN to skip corrupt data and log files. Jobs that include this option may complete successfully, but it is likely that if this data is restored, the database will be in an inoperable state. The SKIP INACCESSIBLE option is not available for Backup Exec server operations. If a backup job encounters corrupt data or log files, the job fails. Symantec recommends that you do not use this option.
- In a CASO environment, the destination storage that you select in the DBA-initiated job template must be locally attached to the central administration server.
If the destination storage includes a storage device pool, all storage in the pool must be locally attached to the central administration server.

See [“About Oracle instance information changes ”](#) on page 1008.

See [“About Oracle instance information changes ”](#) on page 1008.

See [“About backing up Oracle databases”](#) on page 1008.

See [“DBA-initiated job templates”](#) on page 597.

Oracle backup options

The following options are available for Oracle backup jobs. These options appear when you select the Oracle option on the **Backup Job Defaults** dialog box and on the **Backup Options** dialog box for a backup job.

See [“About backing up Oracle databases”](#) on page 1008.

See [“Backing up data”](#) on page 158.

See [“Changing default backup job settings”](#) on page 515.

Table H-5 Oracle backup options

Item	Description
Maximum number of devices to use for resources that support multiple data streams	<p>Specifies the maximum number of devices that the backup job can use.</p> <p>If you specify more than one device, you must choose one of the following items as the destination storage for the backup job:</p> <ul style="list-style-type: none">■ A storage pool.■ A legacy backup-to-disk folder that has at least two concurrent operations enabled. <p>If there is only one storage device for the backup job to use, then the data streams from RMAN are backed up serially to the media.</p> <p>This feature is not available for DBA-initiated jobs.</p>
Minimum number of devices, terminate job if fewer devices are available	<p>Specifies the minimum number of storage devices that the job can use.</p> <p>If the job cannot acquire the minimum number of devices, the job fails.</p> <p>This feature is not available for DBA-initiated jobs.</p>
Delete backed up archived log files	<p>Lets you delete the archived log files automatically after the backup.</p>
Do not back up archived logfiles that have already been backed up	<p>Enables Backup Exec to skip any archived logfiles that have been backed up previously.</p>
Perform the backup offline	<p>Enables Backup Exec to take the database offline before you start the backup job. Backup Exec brings the database online after the backup job is complete.</p> <p>Select this option if the Oracle database is a non-archived logged database.</p>

About restoring Oracle resources

The Oracle Agent lets you restore Oracle databases, tablespaces, or datafiles. You can restore items to their original location or you can redirect the restore to a new location. The restore selections that you choose in Backup Exec are converted to a script. RMAN uses the script to determine what to restore from the Backup Exec backup set. After the data has been restored to the Oracle server, RMAN completes any requested recovery and restore operations. The option you select determine the recovery and restore operations.

Some recovery operations may not require media from the Backup Exec server. For example, the redo logs may still be on the Oracle server. During a restore operation, the amount of data that is restored may not be equal to the amount of data that is backed up. In some cases, the amount of data that is restored is listed as 0 bytes. This behavior is normal because Oracle might skip the datafiles that are already up-to-date on the disk.

If you perform a complete recovery on the whole database, or on a tablespace or datafile, you must restore a backup of the database or files that you want to recover. Then you must apply online or archived redo logs, or both. For the jobs that are initiated both from the Backup Exec server and from a DBA, RMAN determines the specific data that it requires from Backup Exec to complete the restore and the recovery that you request.

Note: Backup Exec does not support Oracle tablespace point-in-time restore (TSPITR) through server-initiated operations.

Backup Exec does not support restore from storage device media and from backup sets for the Oracle Agent because during restore, RMAN decides the backup sets needed for the restore job.

You can only choose Oracle restore selections from the **Resource View** in the Restore Wizard. The **Details View** displays backup sets, but you cannot browse or select the contents.

On the **Resource View**, you can make restore selections from the online database or from control files.

Table H-6 Restore selections for Oracle resources

View restore data in	Description
Online database	<p>Provides a view of the live database (if available). You can select an entire database or select individual tablespaces and datafiles.</p> <p>Note: For Oracle RAC, the Oracle database is listed under its fully qualified domain name. It is in the form RAC-<i>database name</i>-<database ID>.</p>

Table H-6 Restore selections for Oracle resources (*continued*)

View restore data in	Description
Control files	<p>Provides a list of all backed up control files. Each control file lists the date it was backed up and the control file's piece ID.</p> <p>You cannot select individual tablespaces or datafiles for restore.</p> <p>Caution: When you recover to a point in time by using a control file, ensure that the date of the control file backup is before the specified recovery point in time. There should not be any database structure changes between the two times. Additionally, when you restore a control file, the entire database reverts to the point in time of the restored control file.</p>

See [“Methods for restoring data in Backup Exec”](#) on page 232.

About DBA-initiated restore for Oracle

DBAs can initiate restore jobs directly from the RMAN console. For example, you can specify the resources you want restored, and the number of channels to allocate for the restore job. Refer to your Oracle documentation for more information on using the RMAN console.

All DBA-initiated restore jobs are deleted after the jobs have completed.

Note: If you attempt to use a DBA-initiated restore job to restore a datafile, a tablespace, or a database that is online, a message appears on the RMAN console. The message indicates that the restore cannot be performed because Oracle does not allow the restore of these items if they are online. However, this message is not reported to Backup Exec. Therefore, the DBA-initiated restore job is reported in Backup Exec as completing successfully.

About redirecting a restore of Oracle data

In Backup Exec, you can redirect an Oracle instance or its files by redirecting the following:

- An Oracle instance to another Oracle server.

Note: If you redirect the instance to a different Oracle server, ensure that an instance with the same name and database ID (DBID) is set up on that server. The database status should be Nomount. Refer to your Oracle documentation for details on creating an instance with the same name and database ID.

- An Oracle instance to another Oracle server and specifying alternate paths for the Oracle files.
- Tablespace, datafiles, and archive logs to an alternate location on the original server.

Symantec recommends that you select only one instance for each redirected restore operation.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

Requirements for recovering the complete Oracle instance and database using the original Oracle server

If you experience a complete loss, deletion, or destruction of the Oracle instance or database, you can use the same Oracle server for the recovery. You can also use these instructions when you configure a new physical server that uses the same server name and SID name.

To successfully complete the recovery using this scenario, you must have the following items:

Table H-7 Requirements when you recover using the original Oracle server

Item	Description
DBID	If you do not know the DBID, you can find it in the Backup Exec job log or in RMAN after you login.
ControlFile piece ID	You can identify the ControlFile piece ID in the Backup Exec restore view in the Control Files subnode under the Oracle node.
A full system Oracle backup	The full system Oracle backup must include the following: <ul style="list-style-type: none">■ controlfile■ datafiles■ archive logs
The original Oracle server	To successfully recover the Oracle system using disaster recovery scenario 1, you must restore to the original Oracle server.

Recovering the complete Oracle instance and database using the original Oracle server

You can use the same Oracle server for a recovery if you experience a complete loss, deletion, or destruction of the Oracle instance or database.

See [“Requirements for recovering the complete Oracle instance and database using the original Oracle server”](#) on page 1015.

To recover the complete Oracle instance or database using the original Oracle server

- 1 Recreate the Oracle database using the same name you used for the original database that was lost.
- 2 Find and rename the pwd<SID>.ora file.
- 3 Do the following in the order listed to create a new pwd<SID>.ora file:
 - Open a command prompt.
 - Type the following command:
`orapwd file=path\pwsid.ora password=<password>`
- 4 Type the following commands in the order listed:
 - `RMAN`
 - `CONNECT TARGET <sys/password@sid>;`
 - `SHUTDOWN ABORT;`
 - `STARTUP NOMOUNT;`
 - `SET DBID<dbid ID>;`
- 5 At the Backup Exec server, launch the Backup Exec Restore Wizard.
- 6 Select the appropriate ControlFile to restore.

The restore job will fail because the recovery portion encounters inconsistent archive logs. This is a normal occurrence during a disaster recovery.
- 7 After the restore job completes, exit Backup Exec.
- 8 At the Oracle server command prompt, type:
Alter database open resetlogs;
- 9 Close the command prompt.

Requirements for recovering the complete Oracle instance or database to a computer other than the original Oracle server

If you experience a complete loss, deletion, or destruction of the Oracle instance or database, you can restore the instance and database to a computer other than the original Oracle server.

See [“Recovering the complete Oracle instance and database using the original Oracle server”](#) on page 1016.

To successfully complete the recovery, you must have the following items:

Table H-8 Requirements when you recover using a new or alternate Oracle server

Item	Description
DBID	If you do not know the DBID, you can find it in the Backup Exec job log or in RMAN after you login.
ControlFile piece ID	You can identify the ControlFile piece ID in the Backup Exec restore view in the Control Files subnode under the Oracle node.
A full system Oracle backup	The full system Oracle backup must include the following: <ul style="list-style-type: none">■ controlfile■ datafiles■ archive logs

Recovering the complete Oracle instance or database to a computer other than the original Oracle server

You can restore an Oracle instance or database to a computer other than the original Oracle server.

See [“Requirements for recovering the complete Oracle instance or database to a computer other than the original Oracle server”](#) on page 1017.

To recover the complete Oracle instance and database to a computer other than the original Oracle Server

- 1 Recreate the Oracle instance using the same name you used for the original instance that was lost.
- 2 Find and rename the pwd<SID>.ora file.
- 3 Do the following in the order listed to create a new pwd<SID>.ora file:
 - Open a command prompt.

- Type the following command:
orapwd file=path\pwdsid.ora password=<password>
- 4 Type the following commands in the order listed:
 - **RMAN**
 - **CONNECT TARGET <sys/password@sid>;**
 - **SHUTDOWN ABORT;**
 - **STARTUP NOMOUNT;**
 - **SET DBID<dbid ID>;**
 - 5 At the Backup Exec server, launch the Backup Exec Restore Wizard.
 - 6 Select the appropriate ControlFile to restore.
 - 7 Select the option to restore **To a different Oracle server** and select the appropriate options.
 - 8 After the restore job completes, exit Backup Exec.

The restore job will fail because the recovery portion encounters inconsistent archive logs. This is a normal occurrence during a disaster recovery.
 - 9 Move to the Oracle server.
 - 10 Type **Alter database open resetlogs;**.
 - 11 Do one of the following:
 - If an error is encountered while Oracle tries to open the database Note the online redo log path and then update the path.
 - If an error does not occur Do nothing. The disaster recovery is complete.

Symantec Backup Exec Agent for Enterprise Vault

This appendix includes the following topics:

- [About the Agent for Enterprise Vault](#)
- [Requirements for the Enterprise Vault Agent](#)
- [About installing the Enterprise Vault Agent](#)
- [About backup methods for Enterprise Vault backup jobs](#)
- [About backing up Enterprise Vault components](#)
- [About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases](#)
- [About restoring Enterprise Vault](#)
- [Configuring Enterprise Vault to use the name of the new SQL Server that holds the Directory database](#)
- [Best practices for the Enterprise Vault Agent](#)
- [About the Backup Exec Migrator for Enterprise Vault](#)

About the Agent for Enterprise Vault

The Symantec Backup Exec Agent for Enterprise Vault (Enterprise Vault Agent) is installed as part of the Agent for Applications and Databases.

The Enterprise Vault Agent provides data protection for the following Enterprise Vault components:

- Sites

- Vault Store Groups
- Databases
- Indexes
- Vault partitions

The Enterprise Vault Agent can help provide a disaster recovery solution for the data that is archived with Enterprise Vault. Recovery of the archived data is not dependent on the archive source, such as Exchange Server or a specific file system.

The Enterprise Vault Agent lets you do the following:

- Back up and restore Enterprise Vault archives from open or closed vault store partitions.
- Back up and restore individual Enterprise Vault vault store groups from an Enterprise Vault site.
- Back up and restore Enterprise Vault sites, databases, and index locations.

When you back up Enterprise Vault servers the following Enterprise Vault components can be backed up along with the vault partitions:

- Enterprise Vault Directory and Monitoring databases
- Enterprise Vault Audit, FSA Reporting, and Fingerprint databases
- Enterprise Vault vault store databases
- Enterprise Vault indexing files

If you install the Enterprise Vault Compliance and Discovery Accelerator products, the following components can be backed up:

- Enterprise Vault Compliance Accelerator and Discovery Accelerator Configuration databases
- Enterprise Vault Compliance and Discovery Accelerator Customer databases
- Enterprise Vault Discovery Accelerator Custodian database

The Enterprise Vault Agent uses Enterprise Vault Backup mode to back up Enterprise Vault components. By using the Backup mode, the Enterprise Vault Agent can back up Enterprise Vault components without having to suspend Enterprise Vault archiving operations.

For example, when you select a vault store group or site for backup, the individual vault store or indexes are placed in Backup mode. Backup mode lets Enterprise Vault continue archiving operations in other vault store groups or sites. After the backup job successfully completes, the Enterprise Vault Agent takes the Enterprise Vault components out of Backup mode so that those components can continue archival operations.

While Enterprise Vault versions 8.x, 9.x, and 10.x all implement Backup mode, Enterprise Vault 9.x and 10.x offer you more flexibility with your vault store backup jobs. With Enterprise Vault 9.x and 10.x, you can run multiple backup jobs of the same Enterprise Vault 9.x and 10.x vault store simultaneously. With Enterprise Vault 8.x, multiple vault store backup jobs must run one at a time.

For example, you can create multiple backup jobs to back up a vault store. Each backup job includes in its selection list one or more unique partitions of the vault store. Under Enterprise Vault 9.x and 10.x, the partitions are backed up simultaneously when the different backup jobs access them at the same time. Under Enterprise Vault 8.x, the partitions are backed up in both backup jobs; however they are backed up sequentially. The first backup job must finish before the second job starts, or else a backup job failure occurs.

Note: With all versions of Enterprise Vault, the Enterprise Vault Agent automatically backs up the vault store database whenever an open partition is backed up.

The Enterprise Vault Agent backs up the Compliance Accelerator and Discovery Accelerator application databases while they are online. It does not place the databases in Read-only mode or Backup mode before it backs them up.

The Enterprise Vault Agent runs a physical check on each Enterprise Vault database before it backs them up. The Enterprise Vault Agent also runs a physical check on each database before you restore them.

Note: The Enterprise Vault Agent uses physical database consistency checks because physical checks consume less system resources than other types of Database Consistency Checking options.

Backing up and restoring Enterprise Vault databases and related components require specific user account credentials for each Enterprise Vault component you protect.

Table I-1 Supported the user accounts that are required to back up and restore Enterprise Vault components

Enterprise Vault components	User credentials
Enterprise Vault databases and components (vault store, indexes, partitions, vault store database, Directory, Monitoring, Fingerprint, FSA Reporting, and Audit databases)	<p>The following credentials are required:</p> <ul style="list-style-type: none">■ Vault Service account■ Domain Admin account with Role Based Admin privileges <p>You can also use any domain user account that meets the following requirements:</p> <ul style="list-style-type: none">■ The user account must be included in the Administrator group on all servers where Enterprise Vault partitions, and Enterprise Vault databases reside.■ The user account must have Backup related Role Based Admin privileges for the vault store and the index location. Backup-related Role Based Admin privileges include:<ul style="list-style-type: none">■ EVT Manage Vault Store Backup Mode■ EVT Manage Index Location Backup Mode <p>To configure Role Based Admin privileges for a Windows Domain Admin account, see your Enterprise Vault documentation.</p>
Compliance Accelerator and Discovery Accelerator	<p>The following credentials are required:</p> <ul style="list-style-type: none">■ Vault Service account■ Domain Admin account <p>A user account that is a member of the following groups:</p> <ul style="list-style-type: none">■ Administrator's group on those computers where the Compliance and Discovery Accelerator applications reside.

When you back up specific Enterprise Vault components, the other Enterprise Vault components are automatically backed up at the same time. Backup Exec includes these components to hasten an Enterprise Vault recovery.

Table I-2 Enterprise Vault databases that are automatically backed up

When you back up this	Backup Exec automatically backs up this	Description
Enterprise Vault site	Directory database	Backup Exec automatically backs up the Directory database that is associated with the Enterprise Vault site.
Open partition	Vault store database	Backup Exec automatically backs up the vault store database that is associated with the open partition.

Over time the amount of data that Enterprise Vault stores continues to grow. At some point, you may observe that as the data moves through its usage lifecycle, you no longer access it as frequently. You can use the Backup Exec Migrator for Enterprise Vault to automatically migrate the older Enterprise Vault data to the storage devices that Backup Exec manages.

See [“About the Backup Exec Migrator for Enterprise Vault”](#) on page 1041.

Requirements for the Enterprise Vault Agent

Review the following requirements before you use the Agent for Enterprise Vault (Enterprise Vault Agent).

- You must create at least one partition on an Enterprise Vault server before the Enterprise Vault server can publish itself to Backup Exec.
- You must install the Backup Exec Agent for Windows and license the Enterprise Vault Agent on any computer that hosts an Enterprise Vault component.

Note: The Enterprise Vault Agent uses the Agent for Windows to back up all NTFS shares on a remote computer that contains Enterprise Vault data. However, if the Agent for Windows is not installed, the Enterprise Vault Agent uses Microsoft's Common Internet File System (CIFS) to back up the data.

For a device or a filer that does not support the Agent for Windows, the Enterprise Vault Agent uses CIFS to back up the data. Symantec recommends that you create separate backup jobs when you want to do NDMP backups of Enterprise Vault data. You may see a significant performance improvement of NDMP backups with the Symantec Backup Exec NDMP Option.

About installing the Enterprise Vault Agent

The Agent for Enterprise Vault (Enterprise Vault Agent) is installed as part of the Agent for Applications and Databases. To back up all Enterprise Vault servers, the Enterprise Vault Agent must be installed on each Enterprise Vault server in your environment. In addition, the Enterprise Vault Agent must also be installed on any remote computer where Enterprise Vault components are installed. If the Compliance and Discovery Accelerators are installed on remote computers, the Enterprise Vault Agent must be installed on those computers too.

You can install the Enterprise Vault Agent in the following ways:

- Install it automatically from the Backup Exec server as part of a Agent for Windows installation to the local Enterprise Vault server. After you finish the installation, you may need to configure the Enterprise Vault Agent to publish itself to a Backup Exec server of your choice.
See [“About publishing the Agent for Windows to Backup Exec servers”](#) on page 784.
- Install the required Enterprise Vault Agent licenses on the Backup Exec server. After you install the licenses, you can push-install the Backup Exec Agent for Windows to all Enterprise Vault servers and the computers where other Enterprise Vault components are installed.

See [“Installing additional Backup Exec options to the local Backup Exec server”](#) on page 76.

See [“Push-installing the Agent for Windows to remote computers”](#) on page 86.

About backup methods for Enterprise Vault backup jobs

You can select a backup method that depends on the Enterprise Vault object that you want to backup.

The following table describes the type of Enterprise Vault backup jobs you can run. The table also describes the backup methods that are available for each type of backup job.

Table I-3 Backup methods to use with Enterprise Vault backup jobs

To back up	Select	Description
Directory and Monitoring databases Audit database and FSA Reporting database	Full, differential, or incremental backup method	Directory, Monitoring, Audit, and FSA Reporting database backups can use the full and incremental backup methods. These databases cannot be backed up using the differential backup method. If you select the differential backup method, Backup Exec does a full backup instead . Note: Selecting an incremental backup method backs up the database transaction logs and truncates them.
Vault database and Fingerprint database	Full, differential, or incremental backup method	Vault database and Fingerprint database backups can use all three backup methods: Full, differential, and incremental. Note: Selecting an incremental backup method backs up the database transaction logs and truncates them.
Vault partitions and index locations	Full, differential, or incremental backup methods.	You can use all of the backup methods that are available for standard file system backup jobs.

When you combine Enterprise Vault components in a backup job, each component may use a backup method that differs from what you selected for the overall job. For example, you create a job that uses the differential backup method to back up both a Directory database and a partition. However, because a Directory database cannot be backed up using the differential method, Backup Exec uses the full backup method to back up the Directory database. This results in fast and easy restores. After the Directory database is backed up, Backup Exec uses the differential backup method to back up the partition.

Use the following table as a guide.

Table I-4 Actual backup methods that are used for Enterprise Vault components

Enterprise Vault component	Full (F)	Differential (D)	Incremental (I)
Directory and Monitoring databases	F	F	I Always truncates the transaction logs
Vault store database	F	D	I Always truncates the transaction logs
Audit database	F	F	I Always truncates the transaction logs
FSAReporting database	F	F	I Always truncates the transaction logs
Fingerprint database	F	D	I Always truncates the transaction logs
Partition	F	D	I
Index root path	F	D	I
Compliance Accelerator/Discovery Accelerator Configuration database Note: Also includes the Compliance Accelerator and Discovery Accelerator databases that are installed with run-time versions of Enterprise Vault.	F	F	I Always truncates the transaction logs

Table I-4 Actual backup methods that are used for Enterprise Vault components (*continued*)

Enterprise Vault component	Full (F)	Differential (D)	Incremental (I)
Compliance Accelerator/Discovery Accelerator Customer database Note: Also includes the Compliance Accelerator and Discovery Accelerator databases that are installed with run-time versions of Enterprise Vault.	F	D	I Always truncates the transaction logs
Discovery Accelerator Custodian database Note: Also includes the Discovery Accelerator Custodian databases that are installed with run-time versions of Enterprise Vault.	F	D	I Always truncates the transaction logs

See [“Backup methods in Backup Exec”](#) on page 188.

See [“About backing up Enterprise Vault components”](#) on page 1028.

Enterprise Vault backup options

You can select a backup method that is based on the type of Enterprise Vault database that you want to back up.

See [“About backup methods for Enterprise Vault backup jobs”](#) on page 1024.

See [“Backing up data”](#) on page 158.

See [“About backing up Enterprise Vault components”](#) on page 1028.

About backing up Enterprise Vault components

You can select any or all of the Enterprise Vault components for backup when you create a backup job. If you select all of the components for backup in the same job, recovery time is faster. However, if you create multiple backup jobs for the components, the backup jobs run faster.

The Enterprise Vault components that you can select are described in the following table, along with recommendations for backup:

Table I-5 Enterprise Vault components

Enterprise Vault Component	Description
Directory database	<p>The Directory database is a Microsoft SQL Server database that contains configuration data.</p> <p>After the database is populated, the amount of data in the Directory database changes very little over time.</p> <p>You should back up the Directory database after you add or remove any Enterprise Vault component. You should also back up the Directory database if you change the location of any component. Configuration changes can include creating vault stores, creating vault store partitions, and changing vault store partition statuses.</p>

Table I-5 Enterprise Vault components (*continued*)

Enterprise Vault Component	Description
Monitoring database	<p>Enterprise Vault includes a Monitoring agent on each Enterprise Vault server. The Monitoring agent monitors the following:</p> <ul style="list-style-type: none">■ The status of Enterprise Vault services and tasks.■ Performance counters for vault stores, disk space, memory, and processors.■ The status of Exchange Server journal mailbox target archiving targets, including item counts for Inbox, Archive Pending, and failed operations such as Failed DL Expansion.■ The status of Lotus Domino Server journaling location archiving targets, including item counts for Inbox, Archive Pending, and failed operations. <p>The Monitoring agent collects monitoring data at scheduled intervals, typically every few minutes.</p> <p>All of the information that the Monitoring agent collects is stored in a Microsoft SQL Server database called the Monitoring database.</p>
Fingerprint databases	<p>The Fingerprint databases contains the single instance storage-related information for all of the vault stores in the vault store group.</p> <p>If you enable single instance storage of archived items, you should back up the Fingerprint databases on a regular basis.</p>
Index location	<p>The index location stores all of the archived data content that is indexed to enable fast searching and retrieval of archived items. The indexing data is stored in index files in the location that is specified when you install Enterprise Vault.</p> <p>You should back up the index location on regular basis.</p>

Table I-5 Enterprise Vault components (*continued*)

Enterprise Vault Component	Description
Vault store group	The vault store group is a logical entity. If you select it for backup, all of the vault databases, vault store partitions, and the Fingerprint databases are backed up. Because these components are closely related, you should consider selecting the vault store group to back up all of these components together.
Vault store	The vault store is a logical entity. If you select it for backup, all of the vault databases and the vault store partitions are backed up.
All partitions	<p>A vault store partition represents the physical location where the archived items are stored. A vault store can contain one or more vault store partitions. If you select All Partitions for backup, then all of the vault store partitions in the vault store are selected for backup.</p> <p>Note: When you back up an open partition, Backup Exec automatically backs up the vault store database.</p>
Site	An Enterprise Vault site is a logical representation of an installation of the Enterprise Vault. If you select this component for backup, the Directory database is also automatically backed up.
Compliance Accelerator database and Discovery Accelerator database	Installed as optional add-ons to Enterprise Vault, these databases are part of the Discovery Accelerator and Compliance Accelerator products.

See [“Backing up data”](#) on page 158.

See [“Editing backup definitions”](#) on page 206.

See [“Adding a stage to a backup definition”](#) on page 219.

About consistency checks for Enterprise Vault databases and Compliance and Discovery Accelerator databases

Backup Exec automatically checks the physical consistency of an Enterprise Vault database before a backup job and after a restore job. It also checks the consistency of the Compliance and Discovery databases before a backup job and after a restore job. Backup Exec uses Microsoft SQL Server's Physical Check Only utility for consistency checks of the databases. In the event a consistency check fails, Backup Exec continues with the job and reports the consistency check failures in the Backup Exec job log.

If consistency checks fail during a restore operation, Backup Exec continues the job and reports the consistency check failures in the Backup Exec job log.

For more information about the Physical Check Only utility, see the Microsoft SQL Server documentation.

About restoring Enterprise Vault

Review the following before you begin an Enterprise Vault restore operation.

- When you restore an Enterprise Vault installation, you should restore the Directory database in a separate restore job. After you successfully restore the Directory database, you can restore other Enterprise Vault components and partitions.
- When you restore Enterprise Vault databases, you can select the options that either leave databases in a ready-to-use state or in a non-operational state. The non-operational state options that you select apply to all Enterprise Vault databases except the vault store database. When you restore an Enterprise Vault vault store database, the Agent for Enterprise Vault (Enterprise Vault Agent) places the vault store database in Enterprise Vault Backup mode. If the vault store database remains in a non-operational state after the restore job completes, the Enterprise Vault Agent cannot remove it from Backup mode.

If you select the option that leaves the databases ready to use, the following applies:

- The Enterprise Vault Agent restores the vault store database in a ready-to-use, operational state. The vault store database's operational status is maintained even when you select additional backup sets for restore in the same vault store database restore job. Additional backup sets can include Full, Differential, and Incremental backup methods.

When you choose the option that leaves the databases in a nonoperational state, the following applies:

- The Enterprise Vault Agent prompts you to stop the **Enterprise Vault Storage Service** before you start the vault store database restore operation. You can restart the vault store restore operation again after the Enterprise Vault Storage Service stops.

As a best practice, Symantec recommends that you restore the vault store database in a ready-to-use state. When you restore the vault store database in a nonoperational state, Enterprise Vault cannot remove it from Backup mode after the restore operation finishes.

See [“Enterprise Vault restore options”](#) on page 1033.

- You can individually restore Enterprise Vault components. Before you begin the restore, the databases and other components may or may not exist on the destination Enterprise Vault server. If the databases do not exist, you can restore them using the Enterprise Vault Agent. After the restore job completes, you must configure Enterprise Vault to use the restored databases.

To configure Enterprise Vault to use the restored databases, see your Enterprise Vault documentation.

These items include the:

- Enterprise Vault 8.x/9.x/10.x Directory, Monitoring, Audit, FSAReporting, and Fingerprint databases
- Vault store databases, indexes, and partitions.
- Compliance and Discovery Accelerator Configuration and Customer databases.
- Discovery Accelerator Custodian database
- Symantec recommends that you use the Enterprise Vault service account or an account with rights to access the restore selections as the default logon account. Otherwise, you may have to enter proper credentials for each Enterprise Vault resource that you select for restore.
- After you restore Enterprise Vault, a message appears that says you need to run Enterprise Vault recovery tools. The recovery tools are used to re-synchronize Enterprise Vault with the newly restored databases after you complete the restore.

For information on running the Enterprise Vault recovery tools, see your Enterprise Vault documentation.

Before you restore Enterprise Vault sites, servers or other components, you should have the following items installed on the destination computer:

- Enterprise Vault

- The Backup Exec Agent for Windows

Note: You must install the Agent for Windows on remote Enterprise Vault computers where you want to restore Enterprise Vault components.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

Enterprise Vault restore options

Use the following table to select the restore option you want to use when you restore the Enterprise Vault databases.

Table I-6 Enterprise Vault restore options

Item	Description
Automatically take the Enterprise Vault databases offline when restoring selected databases	<p>Takes the shared Enterprise Vault Directory, Monitoring, Audit, FSA Reporting, and Fingerprint databases offline so Backup Exec can replace them during a restore job.</p> <p>Note: If you don't use this option, you must stop the Directory and Admin services on all Enterprise Vault servers before you restore the previously mentioned databases. In addition, you must also stop the Accelerator Manager server on all of the Compliance Accelerator servers and the Discovery servers. Only after you stop the Accelerator Manager can you restore the Customer, Configuration, and Custodian databases.</p> <p>This option causes the Enterprise Vault Admin and Directory services on all related Enterprise Vault servers to terminate the connection to the Directory database that you restore.</p> <p>It also terminates connections to the following:</p> <ul style="list-style-type: none">■ Monitoring database■ Audit, Fingerprint, and FSA Reporting databases (Enterprise Vault 8.x, 9.x, 10.x only)■ Configuration, Customer, and Custodian databases <p>When the restore job completes, you must manually restart the Enterprise Vault Admin and Directory services on your Enterprise Vault server. After you restart the services, the services reconnect to the restored databases and Enterprise Vault begins archival operations again.</p> <p>Note: This option causes the Enterprise Vault Admin and Directory services on all Enterprise Vault servers to terminate their connections to the Directory database that you restore. It also terminates the connections to the Enterprise Vault Accelerator Manager database.</p>

Table I-6 Enterprise Vault restore options (*continued*)

Item	Description
Do not take the Enterprise Vault databases offline	<p>Leaves all Enterprise Vault databases online.</p> <p>If you use this option, you must stop the Directory and Admin services on all Enterprise Vault servers before you restore the previously mentioned databases. In addition, you must also stop the Accelerator Manager server on all of the Compliance Accelerator servers and the Discovery servers. Only after you stop the Accelerator Manager can you restore the Customer, Configuration, and Custodian databases.</p>
Leave the database ready to use; additional transaction logs or differential backups cannot be restored	<p>Rolls back all uncompleted transactions when you restore the last database, differential, or log backup. After the recovery operation, the database is ready for use. If you do not select this option, the database is left in an intermediate state and is not usable.</p> <p>If you select this option, you cannot continue to restore backups. You must restart the restore operation from the beginning.</p>
Leave the database nonoperational; additional transaction logs or differential backups can be restored	<p>Creates and maintains a standby database.</p> <p>By using this option, you can continue restoring other backups sets for non-operational databases.</p> <p>See your SQL documentation for information on standby databases.</p>

Note: Symantec recommends that you select all required backup sets when you run a single restore job for a vault store database. All required backup sets can include full, differential, and incremental backup sets. The vault store database should also be restored in a ready-to-use state after the restore job completes.

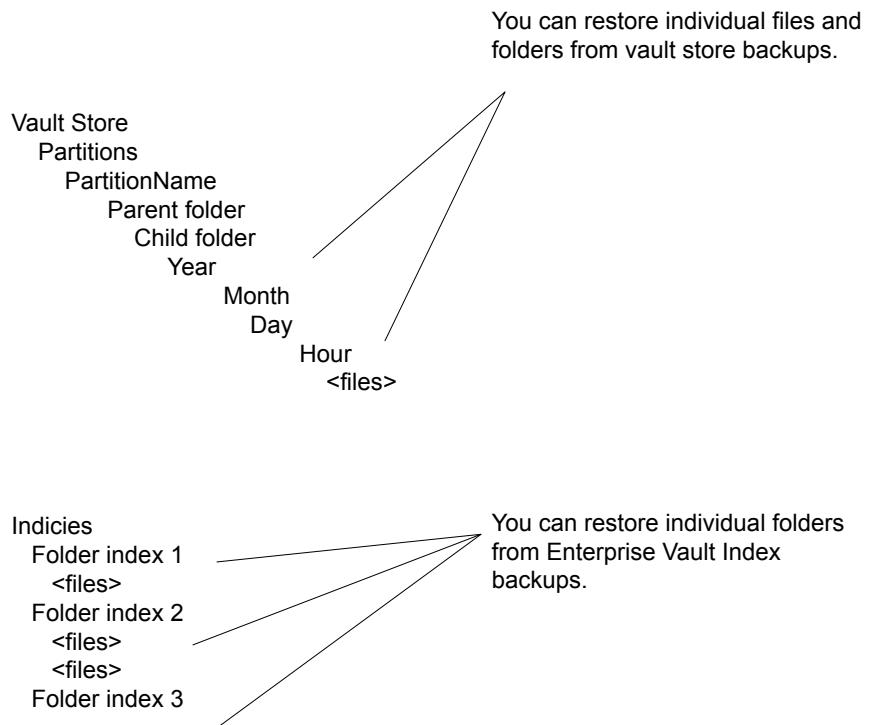
See [“Methods for restoring data in Backup Exec”](#) on page 232.

About restoring individual files and folders with the Enterprise Vault Agent

The Agent for Enterprise Vault (Enterprise Vault Agent) supports individual file and folder restores from vault store partition backups. You can also restore complete index locations or individual folders from Enterprise Vault index backups.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

Figure I-1 Restoring individual files from vault store partitions and complete folders from an Enterprise Vault index



About automatic redirection of Enterprise Vault components under an Enterprise Vault server

You can change the location of the vault store databases, Fingerprint databases, or partitions to a location that differs from where they were backed up. During restores of the vault store database, Fingerprint databases, or partitions, the Agent for Enterprise Vault (Enterprise Vault Agent) detects the location change. It then automatically redirects the component restores to the new location.

Note: Automatic redirected restores of the vault databases, partitions, or Fingerprint databases occur when you change only the location of these Enterprise Vault components. The names of the partitions, vault stores, and vault store groups must not change from the time the partition was originally backed up.

See [“About restoring Enterprise Vault”](#) on page 1031.

About redirecting a restore for an Enterprise Vault component

You can redirect the restore of the Enterprise Vault components.

The following table describes the requirements for redirecting a restore for an Enterprise Vault component:

Table I-7 Requirements for redirecting a restore for an Enterprise Vault component

Component	Requirement
All Enterprise Vault components	<p>The following are requirements for redirecting the restore of all Enterprise Vault components:</p> <ul style="list-style-type: none">■ All Enterprise Vault components must already exist on the server to which you redirect the restore. If they do not exist, you must create them. See your Enterprise Vault documentation.■ The Backup Exec logon account that you use must have the same credentials as the Vault Store service account.
Enterprise Vault databases	<p>The following are requirements for redirecting the restore of the databases:</p> <ul style="list-style-type: none">■ You must create a separate job for each database that you want to redirect.■ You must redirect all databases to the same SQL server.

Table I-7 Requirements for redirecting a restore for an Enterprise Vault component (*continued*)

Component	Requirement
Vault store databases	<p>The following are additional requirements for redirecting the restore of a vault store database:</p> <ul style="list-style-type: none">■ Ensure that the Directory database already exists on the server to which you redirect the restore.■ Ensure that the Directory database contains an entry for the vault store that uses the new SQL Server name.
Vault store partition	<p>The following are additional requirements for redirecting the restore of a vault store partition:</p> <ul style="list-style-type: none">■ The vault store must already contain a vault partition with the same name. If a vault partition does not exist, you must create it.
Index location	<p>The Directory database must already be configured with the new index location.</p> <p>See your Enterprise Vault documentation.</p>

Redirection options for Enterprise Vault

You can redirect a restore job for Enterprise Vault components.

See [“About redirecting a restore for an Enterprise Vault component”](#) on page 1037.

Table I-8 Redirection options for Enterprise Vault

Item	Description
To a new Microsoft SQL server	<p>Redirects the restore jobs of Enterprise Vault databases and Accelerator databases to a different SQL Server.</p> <p>Displays the name of the server to which you want to redirect the restore job for a vault store.</p> <p>Note: Vault store databases are restored for Enterprise Vault 8.x, 9.x and 10.x only.</p>

Table I-8 Redirection options for Enterprise Vault (continued)

Item	Description
Instance	Displays the name of the instance of the SQL Server to which you want to redirect the restore job for a vault store.
Restore index root to a new location	Redirects the restore job for the index root to a new location. If you redirect the restore of the Enterprise Vault server, you can specify an alternate path on the destination server. You can also redirect the index root location to an alternate path on the original server.
Path	Displays the path name to which you want to redirect the restore job for an index root.
Restore partition root to a new location	Redirects the restore job for a vault store partition to a new location. Partitions are restored for Enterprise Vault 8.x, 9.x , and 10.x only.
Path	Displays the path name to which you want to redirect the restore job for a vault store partition.
Enterprise Vault logon account	Specifies the logon account to use.

Configuring Enterprise Vault to use the name of the new SQL Server that holds the Directory database

Use the following steps to configure Enterprise Vault to use the name of the new SQL Server that holds the Directory database.

To configure Enterprise Vault to use the name of the new SQL Server that holds the Directory database

- 1 At each Enterprise Vault server, use Enterprise Vault to change the name of the previous SQL Server computer. Change the name to the name of the SQL Server computer that now holds the Directory database.

See your Enterprise Vault documentation.

- 2 Restart the Enterprise Vault Admin service on all Enterprise Vault servers that use the Directory database.

Two directory names appear in the backup selections view after you restart the Enterprise Vault Admin service on the Enterprise Vault server.

For example, **Directory on <OldSQL_computer_name>** and **Directory on <NewSQL_computer_name>**).

- 3 On the Backup and Restore tab, right-click the Enterprise Vault server that you want to back up.
- 4 On the Backup menu, select the backup definition that you want to use.
- 5 In the Selections box, click **Edit**.
- 6 Expand **Directory on <SQL server computer where you moved the Directory database>**.
- 7 Expand all items under **Directory on <SQL server computer where you moved the Directory database>**.

The **Directory** and **Monitoring** databases, Enterprise Vault 8.x, 9.x, or 10.x **FSA Reporting** and **Audit** databases, and the Enterprise Vault sites should appear. In addition, the Directory database should display the new SQL Server name and instance where it was redirected.

When you configure a new Directory database backup job, you must select the Directory database from the current Directory server. Backup Exec automatically removes the previous Directory server name 13 days after you complete the Directory database move.

- 8 To manually remove the previous server name, right-click **Directory on <OldSQL_computer_name>**.
- 9 Click **Remove**.

See your Enterprise Vault documentation.

Best practices for the Enterprise Vault Agent

Symantec recommends the following best practices when you use the Agent for Enterprise Vault (Enterprise Vault Agent).

- Back up the Enterprise Vault Directory database after you make any configuration changes in Enterprise Vault.
- Restore the Enterprise Vault Directory database in a separate Backup Exec restore job.
- Restore all Full, Differential, and Incremental backup sets of the vault store database in a single restore job.
- Do not allow the backup window and archive window to overlap.
- Do not allow the backup window and the migration window to overlap.
- Make sure Enterprise Vault components are not in Backup mode before you back up the Enterprise Vault Directory database.
- If you install both the Symantec Backup Exec NDMP Option and the Enterprise Vault Agent, pick only one product to protect an Enterprise Vault partition that resides on NDMP filers.
- Do not change the recovery model of any database that is created by Enterprise Vault. Enterprise Vault configures each database in full recovery mode when it creates them.

See [“About the Agent for Enterprise Vault”](#) on page 1019.

About the Backup Exec Migrator for Enterprise Vault

The Backup Exec Migrator for Enterprise Vault (Backup Exec Migrator) lets you automatically migrate archived Enterprise Vault data to the storage devices that Backup Exec manages. By migrating the archived Enterprise Vault data from a partition, you can reclaim disk space on the Enterprise Vault server without incurring the cost of additional hardware.

By migrating Enterprise Vault archive data to the Backup Exec server storage devices, you also ensure an added level of storage redundancy using an off-host environment.

See [“How the Backup Exec Migrator works”](#) on page 1042.

See [“Configuring the Backup Exec Migrator”](#) on page 1049.

Backup Exec Migrator for Enterprise Vault requirements

Before you configure the Backup Exec Migrator, ensure that your Enterprise Vault server meets the following requirements:

- Backup Exec Agent for Enterprise Vault must be installed on the Enterprise Vault server.
- Enterprise Vault migration and collections must be enabled for the Enterprise Vault partition from which you want to migrate data.
- Enterprise Vault 8.0 SP3 or higher must be installed on the Enterprise Vault server.

How the Backup Exec Migrator works

Enterprise Vault automatically initiates all data migration operations from the Enterprise Vault server after you configure the Backup Exec Migrator. Enterprise Vault makes decisions on what should be migrated based on the archival policies and the data retention policies that you configure in the Enterprise Vault Administration Console. The Backup Exec Migrator then migrates the archived data to a Backup Exec server after Enterprise Vault collects the eligible data from the vault store partitions. When you configure migration options for a partition, you can set the migration period. All migration options are configured at the Enterprise Vault server.

Table I-9 Enterprise Vault data migration process

Action	Notes
Enterprise Vault archives eligible partition data that is based on the file size or the file creation date.	All data that is eligible for archive is determined in the partition where you want to migrate data. See your Enterprise Vault documentation.

Table I-9 Enterprise Vault data migration process (continued)

Action	Notes
After Enterprise Vault completes the archival process, an Enterprise Vault collection process collects the archived data.	<p>The collection process places the archived data into Windows .cab files. The .cab files are stored in the partition where the migration occurs.</p> <p>Eligible data can include Enterprise Vault files with the following extensions:</p> <ul style="list-style-type: none">■ .dvf■ .dvssp■ .dvsc■ .dvs <p>Note: Some eligible data cannot be compressed into .cab files due to file size restrictions. However, the Backup Exec Migrator still migrates the data during the migration operation.</p> <p>See your Enterprise Vault documentation.</p>

Table I-9 Enterprise Vault data migration process (*continued*)

Action	Notes
The Backup Exec Migrator initiates the migration of the archived data files to a Backup Exec server.	

Table I-9 Enterprise Vault data migration process (continued)

Action	Notes
	<p>Migration period schedules are determined when you configure migration for a partition and when you configure a collection schedule for the partition.</p> <p>See “Configuring Enterprise Vault collections” on page 1050.</p> <p>See “Configuring the Backup Exec Migrator to communicate with Enterprise Vault” on page 1054.</p> <p>If you follow the Symantec configuration recommendations for the Backup Exec Migrator and Enterprise Vault partitions, one migration job for each partition runs during a migration period. However, the Backup Exec Migrator may create separate migration jobs for each partition folder if you do not follow the configuration recommendations. If separate jobs are created, the resulting overhead that is required to run the jobs results in degraded migration and retrieval performance.</p> <p>Note: If you schedule a file retrieval request from the Enterprise Vault server between migration periods, separate jobs are created even though you followed the configuration recommendations. In this case, the Backup Exec Migrator automatically creates separate jobs to facilitate retrieval of the requested file. During a migration operation, the restore job can be scheduled to run between migration jobs.</p> <p>If you do not follow the configuration recommendations, file retrieval performance can be affected.</p> <p>To ensure the most efficient migration and retrieval performance possible, follow the Symantec recommendations when you configure the Backup Exec Migrator and the Enterprise Vault partitions.</p> <p>See “Configuring the Backup Exec Migrator”</p>

Table I-9 Enterprise Vault data migration process (*continued*)

Action	Notes
	on page 1049.
Backup Exec completes the migration process by moving all of the migrated files to storage devices.	Symantec recommends configuring two storage devices for staged migration operations. See “About using staged migrations with Backup Exec and the Backup Exec Migrator” on page 1046. See “Configuring the Backup Exec Migrator” on page 1049.

After Backup Exec migrates the .cab files to the storage devices, you can review migration details by looking at the job history details for each Enterprise Vault server where the migration occurs.

See [“About the Job History”](#) on page 264.

About using staged migrations with Backup Exec and the Backup Exec Migrator

When you configure Backup Exec to work with the Backup Exec Migrator, Symantec recommends that you configure two storage devices for staged migration operations. When you consider the devices to use, consider selecting a high performance backup-to-disk folder and a slower performance tape device. By using two devices, archived data can be migrated in two stages.

During the first stage, Backup Exec migrates the data it receives from the Backup Exec Migrator to a backup-to-disk folder on a high performance hard drive. By using a backup-to-disk folder, you can minimize the amount of time it takes to perform the initial migration. During the second migration stage, Backup Exec creates a duplicate job to migrate the archived data from the backup-to-disk folder to a tape device. You can schedule the duplicate job to move the archived data to a tape device at times when Backup Exec server activity is low.

See [“Configuring the Backup Exec Migrator to work with a Backup Exec server”](#) on page 1052.

See [“Configuring the Backup Exec Migrator to communicate with Enterprise Vault”](#) on page 1054.

About Backup Exec Migrator events

The Backup Exec Migrator generates events that specify the status of the tasks that it runs. The events also provide useful information for troubleshooting purposes. You can view the events on the computer where you installed the Enterprise Vault Storage Service by viewing the Windows Event Viewer. From the Event Viewer, you can see the events under **Enterprise Vault**. You can also view the events in the Enterprise Vault Dtrace Utility.

For more information on the Enterprise Vault Dtrace Utility, see your Enterprise Vault documentation.

See [“About Backup Exec Migrator logs”](#) on page 1047.

About Backup Exec Migrator logs

The Backup Exec Migrator can create log files that log all migration activity. The log files reside on both the Enterprise Vault server and the Backup Exec server. Backup Exec Migrator log files can help you troubleshoot migration issues.

Before you can view the log files, you must enable Backup Exec Migrator logging on the Enterprise Vault server and on the Backup Exec server. To enable Backup Exec Migrator logs on the Enterprise Vault server, edit the Windows registry.

For information on enabling Backup Exec Migrator logging on the Enterprise Vault server, see the following:

<http://entsupport.symantec.com/umi/V-269-27>

You must also enable Backup Exec Migrator logging on the Backup Exec server.

See [“Using the Backup Exec Debug Monitor for troubleshooting”](#) on page 707.

Note: Partition Recovery Utility log files are enabled by default.

After you enable logging on the Enterprise Vault server and on the Backup Exec server, the following types of log files are created:

- VxBSA log files
For example, <computer_name>-vxbsa<00>.log
- Partition Recovery Utility log files
For example, partitionrecovery<00>.log
- Backup Exec server log files
For example, <computer_name>-bengine<00>.log

Each time the Backup Exec Migrator is started, separate VxBSA log files are created. As a result, each new log file's sequential number increments by one.

For example, <computer_name>vxbsa00.log, <computer_name>vxbsa01.log.

Similarly, a new log file is created each time the Partition Recovery Utility is started. As a result, each new Partition Recovery Utility log file's sequential number increments by one.

For example, partitionrecovery00.log, partitionrecovery01.log

Backup Exec server log file numbers also increment by one as multiple log files are created.

For example, <computer_name>-bengine00.log, <computer_name>-bengine01.log

You can find the log files in the following locations.

Table I-10 Backup Exec Migrator and Partition Recovery Utility log file locations

Log file	Computer	Directory location
VxBSA log files Partition Recovery Utility log files	Enterprise Vault server	C:\Program Files\Symantec\BACKUP EXEC\RAWS\logs
Backup Exec server log files	Backup Exec server	C:\Program Files\Symantec\Backup Exec\Logs

See [“About Backup Exec Migrator events”](#) on page 1047.

About deleting files migrated by Backup Exec Migrator

Enterprise Vault automatically deletes archived items when the item's Enterprise Vault retention periods expire. An Enterprise Vault retention period indicates how long Enterprise Vault retains archived items before it deletes them.

The Backup Exec Migrator maintains existing Enterprise Vault retention periods for archived items when it migrates the archived items to tape. As a result, when an item's data retention period expires, Enterprise Vault issues a command to delete the item from the storage tape that Backup Exec manages. To delete the expired archive item, the .cab file it resides in must be deleted from tape.

Note: Although the Backup Exec Migrator maintains existing Enterprise Vault retention periods, it does not initiate the deletion of expired archived items or archived partitions from tape. Only Enterprise Vault can initiate the deletion of expired items and partitions.

For more information on deleting expired items, see your Enterprise Vault documentation.

Because the .cab files may contain archived items with different retention periods, an expired item may be marked as deleted in the Backup Exec catalogs. However, it may not be immediately deleted from tape. All archived items in a .cab file must have expired retention periods before Enterprise Vault issues a command to delete the .cab file from tape.

Enterprise Vault can also delete entire archived vault store partitions from tape. After you delete an active Enterprise Vault vault store partition by using the Enterprise Vault Administration Console, Enterprise Vault deletes the associated archived partition from tape.

Backup Exec automatically recycles the tapes when all of the items on the tape are marked as deleted in the catalogs. Backup Exec checks for expired Enterprise Vault Migrator media once every 24 hours. If Backup Exec detects such media, it logically moves the media to the Scratch Media set and then generates an information alert informing you of the move.

Note: Expired Enterprise Vault Migrator media is defined as media that contains only migrated Enterprise Vault data that is marked as deleted in the Backup Exec catalogs.

See [“Managing tapes”](#) on page 378.

Note: You should ensure that migrated Enterprise Vault data remains accessible on the tapes that are used for migration purposes until the Enterprise Vault data retention periods expire. Therefore, Symantec recommends that you configure a retention period of 999 years for all tapes that are used for migration purposes.

See [“Overwrite protection periods and append periods in media sets”](#) on page 382.

Configuring the Backup Exec Migrator

All of the program files that are required to run the Backup Exec Migrator are installed when you install the Agent for Enterprise Vault (Enterprise Vault Agent) on the Enterprise Vault server. However, before you can use the Backup Exec Migrator,

you must configure it to work with both a destination Backup Exec server and the Enterprise Vault server.

Table I-11 Enterprise Vault configuration process

Step	Description
Step 1	Configure Enterprise Vault collections. See “Vault store partition properties - Collections” on page 1051.
Step 2	Configure the Backup Exec Migrator to work with a Backup Exec server. See “Configuring the Backup Exec Migrator to work with a Backup Exec server” on page 1052.
Step 3	Configure the Backup Exec Migrator to work with Enterprise Vault. See “Configuring the Backup Exec Migrator to communicate with Enterprise Vault” on page 1054.

Use the following configuration recommendations for both the Backup Exec Migrator and the Enterprise Vault partitions:

- Configure the Enterprise Vault partitions to save migrated data locally.
Do not configure Enterprise Vault partitions to delete files immediately after a migration operation finishes.
See your Enterprise Vault documentation for details on configuring a partition for migration.
- Configure the Backup Exec server template to run staged migrations.
See [“About using staged migrations with Backup Exec and the Backup Exec Migrator”](#) on page 1046.

Failure to follow the configuration recommendations results in degraded migration and retrieval performance.

Configuring Enterprise Vault collections

Before you can use the Backup Exec Migrator to migrate Enterprise Vault archived data from a partition, Enterprise Vault first needs to collect the data.

To configure Enterprise Vault collections

- 1 From the Enterprise Vault Console, navigate to a vault store partition from which you want to migrate data.
- 2 Right-click the partiition, and then click **Properties**.
- 3 On the **Collections** tab, check **Use collection files**.
- 4 Set collection options as appropriate.
See [“Vault store partition properties - Collections”](#) on page 1051.
- 5 Click **OK**.

Vault store partition properties - Collections

Before you can use the Backup Exec Migrator to migrate Enterprise Vault archived data from a partition, Enterprise Vault needs to collect the data to be migrated.

See [“Configuring Enterprise Vault collections”](#) on page 1050.

Table I-12 Vault store partition properties - Collection options

Item	Description
Use collection details	Lets you set Enterprise Vault as the collector.
Start at	Indicates the local time at which you want collection to start.
End at	Indicates the local time at which you want collection to finish. Enterprise Vault stops collecting at this time or when it has no more files to collect, whichever comes first.
Limit collection files to <number> megabytes	Indicates the maximum size for collection files. The default size is 10 MB, although you can specify a file size range from 1 MB to 99 MB. You may want to change this value to optimize the use of your backup media
Collect files older than	Indicates the amount of time that must elapse since items were archived before they are eligible for collection.

Configuring the Backup Exec Migrator to work with a Backup Exec server

Use the following steps to configure the Backup Exec Migrator to work with a destination Backup Exec server.

Note: Symantec recommends that you configure two server storages devices when you configure the Backup Exec Migrator to work with Backup Exec. Configuring two storage devices lets you create a staged migration for your archived Enterprise Vault data.

See [“About using staged migrations with Backup Exec and the Backup Exec Migrator”](#) on page 1046.

See [“Configuring the Backup Exec Migrator to communicate with Enterprise Vault”](#) on page 1054.

To configure the Backup Exec Migrator to work with a Backup Exec server

- 1 At the Backup Exec server, start Backup Exec.
- 2 Create a logon account that uses the Enterprise Vault server Vault Service account credentials.

Vault Service account credentials are used so that Backup Exec and the Backup Exec Migrator can complete the migration operation

See [“Backup Exec logon accounts”](#) on page 610.
- 3 Click the Backup Exec button and then select **Configuration and Settings**.
- 4 Click **Backup Exec Settings**, and then click **DBA-initiated Job Settings**.
- 5 Select the **DEFAULT** template, and then click **Edit**.

You can also use an existing template, or you can create a new template specifically for Enterprise Vault migrations.
- 6 Under **Storage**, select **Any disk storage** as the primary storage location for migrated data, and then set the options you want to use with the device.
- 7 Under **Migrator for Enterprise Vault**, click the down arrow next the field for **Vault Service account credentials**.
- 8 Select the logon account that you created in step 2.

See [“Migrator for Enterprise Vault options”](#) on page 1053.
- 9 In the **DBA-initiated Job Settings** dialog box, set other options as appropriate.

See [“DBA-initiated job templates”](#) on page 597.
- 10 Do one of the following:

If you want to configure staged migrations Do the following in the order listed.

See [“About using staged migrations with Backup Exec and the Backup Exec Migrator”](#) on page 1046.

- Under **Duplicate Job Settings**, check **Enable settings to duplicate backup sets for this job**.
- In the **Storage** list, select a type of storage.
- Set other options as appropriate. See [“Duplicate job settings for DBA-initiated jobs”](#) on page 606.
- Click **OK**.

If you do not want to configure staged migrations

Continue with step 12.

11 Click **OK**.

12 Configure Backup Exec Migrator to work with Enterprise Vault.

See [“Configuring the Backup Exec Migrator to communicate with Enterprise Vault”](#) on page 1054.

Migrator for Enterprise Vault options

The Backup Exec Migrator uses the Enterprise Vault server Vault Service account during the Backup Exec Migrator to Backup Exec server authentication process.

Table I-13 Migrator for Enterprise Vault options

Item	Description
Vault Service account credentials	<p>Specifies the Enterprise Vault server Vault Service account credentials to use so that Backup Exec and the Backup Exec Migrator can complete the migration operation.</p> <p>The Vault Service account must be included in either the Administrators group or the Backup Operators group at the Backup Exec server.</p> <p>Note: If the Enterprise Vault server and the Backup Exec server are in different domains, a trust relationship must be established between the domains. The Vault Service account user must be a trusted user at the Backup Exec server. Trust relationships are required so that the Microsoft Security Support Provider Interface (SSPI) can authenticate the Vault Service account user.</p> <p>For more information on domain trust relationships, see your Microsoft documentation.</p>
New	<p>Lets you create a new logon account or edit an existing account.</p> <p>See “Backup Exec logon accounts” on page 610.</p>

See [“Configuring the Backup Exec Migrator to work with a Backup Exec server”](#) on page 1052.

Configuring the Backup Exec Migrator to communicate with Enterprise Vault

Use the following steps to configure the Backup Exec Migrator to communicate with Enterprise Vault.

See [“Configuring the Backup Exec Migrator”](#) on page 1049.

To configure the Backup Exec Migrator to communicate with Enterprise Vault

- 1 At the Enterprise Vault server, navigate to a vault store partition from which to migrate data.
- 2 Right-click the vault store partition, and then click **Properties**.

- 3 On the **Migration** tab, check **Migrate files**.
- 4 In **Remove collection files from primary storage**, set the time period for this option to something longer than zero days.

Do not set it to zero days. Setting the time period to zero days causes Enterprise Vault to immediately delete the migrated data from the partition. More importantly, it causes the Backup Exec Migrator to create separate migration jobs for each partition folder being migrated during a migration period. If separate jobs are created, the resulting overhead that is required to run the jobs results in degraded migration and retrieval performance.

See [“Configuring the Backup Exec Migrator”](#) on page 1049.
- 5 Set other migration options as appropriate.

See [“Vault store partition properties - Migration options”](#) on page 1056.
- 6 On the **Advanced** tab, ensure that **Symantec Backup Exec** appears in the **List setting from** field.
- 7 In the window below the **List setting from** field, select **Backup Exec server**.
- 8 Click **Modify**.
- 9 Type the name or the IP address of the destination Backup Exec server.
- 10 Click **OK**.
- 11 Select **Backup Exec DBA-initiated template**.
- 12 Click **Modify**.
- 13 Enter the name of an existing template that uses the Enterprise Vault server Vault Service account credentials.

The template that you select must be configured to use the Enterprise Vault server Vault Service account. The template you use must also match the template name that you used when you configured the Backup Exec Migrator to work with a Backup Exec server.

See [“Configuring the Backup Exec Migrator to work with a Backup Exec server”](#) on page 1052.
- 14 Click **OK**.
- 15 Ensure that the name of the template that contains the Enterprise Vault server Vault Service account credentials appears in the **Setting** pane.

See [“Configuring the Backup Exec Migrator to work with a Backup Exec server”](#) on page 1052.
- 16 To test the communications between the Enterprise Vault server and the Backup Exec server, click **Test Configuration**.

- 17 If the test fails, ensure that you used the correct credentials for the Vault Service account , and then click **Test Configuration** again.
- 18 Click **OK** after the test successfully completes.
- 19 Click **OK**.

Vault store partition properties - Migration options

Select the Enterprise Vault migration property options that you want to use.

Table I-14 Vault store partition properties - Migration options

Item	Description
Migrate files	<p>Lets you migrate archived Enterprise Vault data to a Backup Exec storage device.</p> <p>Migration can help reduce storage costs by moving collection files to tertiary storage devices. However, retrieval times can increase.</p> <p>See your Enterprise Vault documentation.</p>
Migrator	<p>Indicates the name of the migration application.</p> <p>Symantec Backup Exec must appear in this field.</p>
Migrate files older than	<p>Indicates the amount of time that must elapse since files were last modified before they are eligible for migration.</p> <p>See your Enterprise Vault documentation.</p>

Table I-14 Vault store partition properties - Migration options (*continued*)

Item	Description
Remove collection files from primary storage	<p>Indicates the age at which migrated collection files are removed from the primary storage location.</p> <p>Files that have been migrated to Backup Exec storage media can remain in their primary location for the period of time you specify.</p> <p>Note: Symantec recommends that you set the time period for this option to something longer than zero days, with a longer time period being best. Do not set it to zero days. Setting the time period to zero days causes the Backup Exec Migrator to create separate migration jobs in a migration period for each partition being migrated. If separate jobs are created, the resulting overhead that is required to run the jobs results in degraded migration and retrieval performance.</p> <p>See “Configuring the Backup Exec Migrator” on page 1049.</p>

See [“Configuring the Backup Exec Migrator to communicate with Enterprise Vault”](#) on page 1054.

About viewing migrated Enterprise Vault data

The Backup Exec **Backup Sets** view shows the migrated items for the Enterprise Vault partition. In the **Backup Sets** view, backup sets containing the migrated .cab files appear under a partition name that reflects the Enterprise Vault partition from which the data was migrated. Because the **Backup Sets** view displays the archived data in a read-only mode, you cannot select the data for restore. However, you can retrieve the data in the application where the data resides.

See [“Backup sets ”](#) on page 342.

Note: You can completely retrieve of all archived items that appear in the **Backup Sets** view by using the Partition Recovery Utility.

See [“About the Partition Recovery Utility”](#) on page 1059.

See [“About retrieving migrated Enterprise Vault data”](#) on page 1058.

About retrieving migrated Enterprise Vault data

All file retrieve operations start from the Enterprise Vault server console. You cannot restore archived Enterprise Vault data from Backup Exec.

When files are migrated from a partition, Enterprise Vault creates a shortcut in the partition that replaces the migrated file. The shortcut also links to the storage location of the migrated file. You retrieve files by double-clicking their shortcuts in the Enterprise Vault partition itself. If a partition retains a local copy of the migrated files, Enterprise Vault retrieves the files from the local copies. If Enterprise Vault deletes the migrated files because the partition's file retention period passes, the requested files must be retrieved from Backup Exec storage media.

Note: Backup Exec does not support restore from storage device media and from backup sets for Backup Exec Migrator for Enterprise Vault because restore of data is initiated at the Enterprise Vault server or where the data resides. Backup Sets view displays the archived data in a read-only mode; therefore, you cannot select the data for restore.

Table I-15 How migrated data is retrieved

Action	Notes
Enterprise Vault works with the Backup Exec Migrator to begin the process.	The Backup Exec Migrator identifies the Backup Exec server where the files are stored.
The Backup Exec Migrator schedules a Backup Exec restore job at the server.	Backup Exec restores the requested files.
The Backup Exec Migrator migrates the restored files to the Enterprise Vault server partition from the Backup Exec server.	The Backup Exec Migrator moves the restored files to a location specified by Enterprise Vault, using the name provided by Enterprise Vault.

The retrieval process is automatic after you start the operation at the Enterprise Vault server. It requires no user intervention other than perhaps placing a tape in the tape device if you removed the storage media.

See [“Retrieving migrated Enterprise Vault data”](#) on page 1058.

Retrieving migrated Enterprise Vault data

Use the following steps to restore migrated Enterprise Vault files.

Note: To successfully retrieve the files you want, you may need to place a tape in a tape drive at the Backup Exec server.

To retrieve migrated Enterprise Vault data

- 1 At the Enterprise Vault server, navigate to the partition where you want to retrieve the data.
- 2 Double-click the file that you want to retrieve.

About the Partition Recovery Utility

The Partition Recovery Utility is a command-line application that is automatically installed when you install the Backup Exec Agent for Windows. The utility lets you restore all of a partition's archived files from the Backup Exec storage media in a single operation. You can also use it to recover the archived partition data for each of the Enterprise Vault partitions in a disaster recovery situation.

After you use the Partition Recovery Utility, you can review recovery details by looking at the Backup Exec job history for each Enterprise Vault server where the recovery occurs.

See [“Partition Recovery Utility requirements”](#) on page 1059.

See [“Finding an archive ID”](#) on page 1060.

See [“Starting the Partition Recovery Utility”](#) on page 1060.

Partition Recovery Utility requirements

You must know the following when you use the Partition Recovery Utility:

- The vault store partition name for the data that you want to recover.
- The Archive ID of the partition data that you want to recover.
- An Enterprise Vault server user account with Vault Service Account privileges.

Note: If you run the Partition Recovery Utility on a Windows Server 2008/2008 R2 computer, administrator privileges are required.

In addition, the Partition Recovery Utility must run at the Enterprise Vault server that originally migrated the data you want to restore.

See [“Finding an archive ID”](#) on page 1060.

See [“Starting the Partition Recovery Utility”](#) on page 1060.

Finding an archive ID

You use the archive ID of the data you want to restore along with the vault store partition name when you run the Partition Recovery Utility. The archive ID is an alpha-numeric number of considerable length.

For example, 1D69957C6D917714FB12FEA54C9A8299A1110000ev8archive.EVMBE

You can find the Archive ID listed among the properties of an archived file set.

To find an archive ID

- 1 In the left view of the the Enterprise Vault Administration Console, expand **Archives**.
- 2 Navigate the folder structure and select the folder of the type for data you want to restore.
- 3 In the right view, right-click an archive, and then select **Properties**.
- 4 On the **Advanced** tab, note the archive ID at the bottom.

See [“Starting the Partition Recovery Utility”](#) on page 1060.

Starting the Partition Recovery Utility

Use the following steps to start the Partition Recovery Utility.

To start the Partition Recovery Utility

- 1 From the Enterprise Vault server, open a Windows command prompt.
- 2 Navigate to the Enterprise Vault Agent installation directory.
For example, C:\Program Files\Symantec\Backup Exec\RAWS
- 3 Do the following:

If you start the Partition Recovery Utility on a Windows Server 2008/2008 R2 computer

Type the following command:

```
runas  
/user:<domain\administrator>  
partitionrecovery.exe -vs  
<vault_store_name> -ap  
<archive_ID>
```

If you start the Partition Recovery Utility on all other supported Windows operating system versions

Type the following command:

```
partitionrecovery.exe -vs  
<vault_store_name> -ap  
<archive_ID>
```

- 4 Press **Enter**.

See [“About the Partition Recovery Utility”](#) on page 1059.

Best practices for using the Backup Exec Migrator

Consider the following best practices when you use the Backup Exec Migrator:

- Symantec recommends that you regularly back up the Backup Exec catalogs. In the event the catalogs become corrupt, you can restore them from backups. After you restore the catalogs, you must re-catalog the storage media on which Backup Exec Migrator data is stored. Re-cataloging the storage media ensures that the latest catalog entries are available.
- For best performance, configure the Backup Exec Migrator to migrate data to a backup-to-disk folder and then to a tape device by using a duplicate job. See [“About using staged migrations with Backup Exec and the Backup Exec Migrator”](#) on page 1046. See [“Duplicating backup sets or a job history manually”](#) on page 221.
- In the Enterprise Vault **Migration** options tab, set the time period for **Remove collection files from primary storage** to something longer than zero days. Setting the time period to zero days causes Enterprise Vault to immediately delete the migrated data from the partition.

If you set the time period to zero days, Symantec recommends the following:

- Increase the number of concurrent jobs that are allowed for the backup-to-disk folder you use for migration purposes.
Increase the number of concurrent jobs based on the following formula:
$$\text{<number of recommended concurrent jobs>} = \text{<number of installed tape drives plus two>}$$

For example, if you have two installed tape drives, you should configure the backup-to-disk folder to allow four concurrent jobs.

Concurrent jobs let the Backup Exec Migrator continue to migrate data to disk storage while tape drives process duplicate jobs in a staged migration environment.

Note: You can increase the number of concurrent jobs that run by increasing the total concurrency level of the backup-to-disk devices.

- Symantec recommends that you first collect all of the archived files in one collection and migration operation and then migrate them in the next collection and migration operation. This process helps ensure that the Backup Exec Migrator creates a single job for each migration operation, which improves the migration performance.

See [“About the Backup Exec Migrator for Enterprise Vault”](#) on page 1041.

Troubleshooting Backup Exec Migrator and Partition Recovery Utility issues

Review the following error messages for possible solutions to errors that you may encounter:

- The Backup Exec Migrator logs migration activity in the Windows Event Viewer and in the Enterprise Vault Dtrace Utility on the Enterprise Vault server. It also logs migration activity on the Backup Exec server.
The details that are provided in the log files can help you troubleshoot issues with the Backup Exec Migrator.
See [“About Backup Exec Migrator events”](#) on page 1047.
See [“About Backup Exec Migrator logs”](#) on page 1047.
- The Partition Recovery Utility cannot find any files to be recalled.
There are no file to be recalled from the vault store database using the Archive ID that you provided.
- The Partition Recovery Utility operation will be terminated due to a user request.
You may have stopped the Partition Recovery Utility operation by pressing **Ctrl + C** or **Ctrl + Break**.
- The migrated file name `<file_name>` with ID `<migrated_file_id>` was not found in the Backup Exec backup sets. The recall is skipped for this file.
The Partition Recovery Utility skips collection files if they already exist in the vault store database. To restore the files, delete them from the vault store database, and then run the Partition Recovery Utility again.
- The Partition Recovery Utility cannot find any partitions. Ensure that the name of the vault store is valid, and that there are partitions in the vault store.
The vault store name that you provided may be invalid.

See [“About the Backup Exec Migrator for Enterprise Vault”](#) on page 1041.

See [“About the Partition Recovery Utility”](#) on page 1059.

Symantec Backup Exec Agent for Lotus Domino

This appendix includes the following topics:

- [About the Agent for Lotus Domino Server](#)
- [Lotus Domino Agent requirements](#)
- [About installing the Lotus Domino Agent on the Backup Exec server](#)
- [About the Lotus Domino Agent and the Domino Attachment and Object Service \(DAOS\)](#)
- [Viewing Lotus Domino databases](#)
- [About backing up Lotus Domino databases](#)
- [About restoring Lotus Domino databases](#)
- [How to prepare for disaster recovery on a Lotus Domino server](#)
- [Recovering a Lotus Domino server from a disaster](#)
- [About disaster recovery of a Lotus Domino server using archive logging](#)
- [Disabling the monitor change journal](#)
- [Recovering the Lotus Domino server, databases, and transaction logs when archive logging is enabled](#)
- [Re-enabling the monitor change journal](#)
- [Recovering a Lotus Domino server that uses circular logging](#)

About the Agent for Lotus Domino Server

The Symantec Backup Exec Agent for Lotus Domino Server (Lotus Domino Agent) lets you back up and restore Lotus Domino on local Backup Exec servers and on remote computers. The Lotus Domino Agent backs up Lotus Domino databases, Domino Attachment and Object Service (DAOS) - related NLO files, and transaction logs. You can integrate Lotus Domino database backups with regular server backups without separately administering them or using dedicated hardware.

The Lotus Domino Agent provides support for the following:

- Dynamic inclusion
- Full and incremental online backups of Lotus Domino databases, DAOS-related NLO files, and transaction logs using Lotus Domino APIs.
- Restores the Lotus Domino databases, .nlo files, archived transaction logs, and point-in-time restores.
- Reuse the archived Lotus Domino transaction logs after a successful backup.
- Flexible scheduling capabilities.
- Backup of DAOS-related NLO files when DAOS is configured on a remote share.
- Backup and restore of partitioned and clustered Lotus Domino servers.
- Lotus Domino databases in a Microsoft Cluster Server (MSCS) cluster in both Active-Active and Active-Passive configurations.

See [“Backing up data”](#) on page 158.

See [“Adding a stage to a backup definition”](#) on page 219.

See [“Editing backup definitions”](#) on page 206.

See [“Changing default backup job settings”](#) on page 515.

Lotus Domino Agent requirements

The Agent for Lotus Domino Server (Lotus Domino Agent) supports the backup and restore of Lotus Domino versions 8.x and 9.x.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Following are the requirements for backing up Lotus Domino database files residing on the Backup Exec server, or for remote Windows computers and workstations.

Note: Backup Exec does not support two versions of Lotus Domino on the same computer.

If the Lotus Domino files you want to back up are on the local Backup Exec server, the server must have the following:

- Backup Exec
- An Intel-compatible processor
- The Lotus Domino data directory on the Lotus Domino server

If the Lotus Domino files you want to back up are on a remote computer, the remote computer must have the following:

- Windows operating system
- Backup Exec Agent for Windows
- An Intel-compatible processor
- Corresponding Windows Administrative Share for each volume that contains Lotus Domino databases
- The Lotus Domino data directory on the Lotus Domino server

The following are required to back up Lotus Domino transaction logs:

- The Lotus Domino archive-style transaction logging must be enabled to perform incremental backups and to perform point in time recovery.

The following are required to back up Lotus Domino DAOS-related NLO files:

- The DAOS state must be in read-only mode or enabled.
- The DAOS catalog should be synchronized.

If the Lotus Domino databases are running in a Microsoft Cluster Server cluster, you must have the following:

- Lotus Domino must be running on a Microsoft Cluster Server cluster. For more information, see your Lotus Domino documentation for instructions for setting up Lotus Domino in a Microsoft Cluster Server cluster.
- The Backup Exec Lotus Domino Agent must be installed on all nodes in the Microsoft Cluster Server cluster.

See [“About installing the Lotus Domino Agent on the Backup Exec server”](#) on page 1066.

About installing the Lotus Domino Agent on the Backup Exec server

The Symantec Backup Exec Agent for Lotus Domino Server (Lotus Domino Agent) is installed as part of the Agent for Applications and Databases. It can back up local or remote Lotus Domino databases.

See [“Using a command prompt to install the Agent for Windows on a remote computer”](#) on page 98.

Note: If you install Lotus Domino on the same server on which Backup Exec is already installed, you must restart Backup Exec services to display Lotus Domino database selections.

See [“About backing up Lotus Domino databases”](#) on page 1068.

See [“Starting and stopping Backup Exec services”](#) on page 621.

See [“Editing backup definitions”](#) on page 206.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

About the Lotus Domino Agent and the Domino Attachment and Object Service (DAOS)

Lotus Domino 8.5 incorporates the Domino Attachment and Object Service (DAOS). DAOS-enabled databases (DAOS databases) save significant hard drive space by sharing data between applications on a server. DAOS databases do not save separate copies of every document attachment. Instead, DAOS databases save a single copy of a file attachment to an internal repository. The databases then create and save reference pointers to the stored file attachments.

File attachments are saved to the internal repository with the .nlo file extension. During a full backup of the entire Lotus Domino server, Backup Exec backs up all .nlo files, along with the Domino <server>.id file.

Backup Exec adds one container per partition under **Lotus Domino Databases** called **Domino Attachment and Object Service** in the restore selections view. All backed up DAOS NLO files reside in backup sets under **Domino Attachment and Object Service**. In addition, all backed up <server>.id files reside in the **Databases** container under **Lotus Domino Databases**.

Note: Domino uses `<server>.id` for NLO encryption purposes. If you enable NLO file encryption at the Domino server, the `<server>.id` file must be backed up.

When you select individual DAOS-enabled databases for backup, the referenced .nlo files for each database are included in the backup job. However, the `<server>.id` file is excluded.

With incremental backups, only the databases and the .nlo files that are created since the last full backup of the server are backed up.

Note: In cases where incremental backup jobs fully back up DAOS-enabled databases, all of the .nlo files that each database references are backed up. This scenario occurs when DAOS-enabled databases use circular logging, or when DAOS-enabled databases are in archive log mode and their DBIID changes.

During a full DAOS-enabled Domino database restore, all database data, .nlo files, and the `<server>.id` file are restored. When you restore individual DAOS-enabled databases, Backup Exec restores all database data, including the .nlo files. However, Backup Exec does not restore any .nlo files that match .nlo files in the internal repository. After the DAOS-enabled databases are restored, Backup Exec resynchronizes the Domino DAOS catalog.

During a point-in-time restore of a DAOS-enabled database, some required .nlo files may not be generated when the archive transaction logs are replayed. When this condition occurs, Backup Exec reports the names of the missing .nlo files. You can individually restore the missing .nlo files and then start a Domino DAOS catalog resynchronization operation at the Domino server.

Note: Only the .nlo files that change are backed up when DAOS-enabled databases are in archive log mode and their DBIIDs do not change.

For information on Domino DAOS catalog resynchronization operations, see your Lotus Domino documentation.

See [“About the Agent for Lotus Domino Server”](#) on page 1064.

Best practices for restoring the missing .nlo files

If you decide to individually restore the missing .nlo files, Symantec recommends the following best practices:

- Always restore the .nlo files to the internal repository of the current DAOS-enabled Domino server.

- Instead of making random .nlo files selections, select all of them when you make your selections in the restore selections view. Then use the restore option, **Skip if file exists**. By using the **Skip if file exists** option, Backup Exec restores only the missing .nlo files.
See [“About redirecting the restore of DAOS NLO files”](#) on page 1079.

Viewing Lotus Domino databases

Use the following steps to view Domino databases that are on both local and remote servers.

Lotus Domino transaction logs do not appear under **Lotus Domino Databases**; however, when you select the database for backup, the transaction logs are automatically included.

The same process applies to DAOS NLO files. They do not appear under **Lotus Domino Databases**; however, when you select the database for backup, the .nlo files are automatically included.

See [“About backing up Lotus Domino databases”](#) on page 1068.

To view Lotus Domino databases on both local and remote servers

- 1 On the **Backup and Restore** tab, right-click a Lotus Domino server.
- 2 On the **Backup** menu, select a backup option that you want to use.
- 3 In the selections box, click **Edit**.
- 4 Expand **Lotus Domino Databases**.
- 5 Expand each folder under **Lotus Domino Databases** until a database folder named appears.
- 6 Click each database folder, and then view the Lotus Domino databases in the right pane.

About backing up Lotus Domino databases

When a Lotus Domino backup job is submitted, Backup Exec uses Lotus Domino APIs to obtain the backup of the database. When you back up a DAOS-enabled Domino database, the DAOS NLO files are automatically included. In addition, the transaction logs associated with the Lotus Domino databases are included in the backup only if archive logging is turned on at the server. If they are backed up, the archive logs and the DAOS NLO files are stored in a separate backup set that is stored along with the Lotus Domino database backup set.

The Agent for Lotus Domino Server (Lotus Domino Agent) supports the backup of the following types of files:

- .ntf - Lotus Notes Template Files
- .nsf - Lotus Notes Database Files
- .box - Lotus Mailbox Files
- .dsk - Cache Files
- .txn - Transaction log files
- .nlo - DAOS attachment files

Transaction log files and DAOS attachment files do not appear in the Backup Exec backup selections view; however, they do appear in restore selections view.

Note: If your backup of Lotus Domino data includes a duplicate stage, make sure that you select all related Lotus Domino transaction logs and Lotus Domino DAOS. Otherwise, you cannot successfully restore the Lotus Domino data from the duplicate backup set.

You must back up .nsf, .ntf, and .box files to properly recover Lotus Domino databases. If you want to back up .njf, .ncf, .id, .dic, or notes.ini files, you must select them for backup from the volume in which the Lotus Domino Program directory is located.

The Domino agent uses the full and incremental backup methods to back up Lotus Domino databases. An incremental backup is smaller and faster than a full backup because only archived transaction logs, DAOS NLO files, unlogged databases, and logged databases with DBIIDs that have changed are backed up.

Note: The full backup method should be used when the DBIID for the database has changed since previous transactions cannot be applied to the new database.

Although DAOS and non-DAOS Domino servers use additional Domino-related databases and support files, Backup Exec does not back them up. Domino automatically recreates the items after you restart the Domino servers.

Backup Exec excludes the following support files from backup jobs:

- daos.cfg
- daoscat.nsf
- dbdirman.nsf

Note: You cannot back up databases to devices that are attached to a computer on which the Remote Media Agent for Linux Servers is installed.

See [“Lotus Domino Agent backup options”](#) on page 1070.

See [“Backing up data”](#) on page 158.

See [“Editing backup definitions”](#) on page 206.

See [“Editing a stage”](#) on page 221.

About selecting backup options for Lotus Domino databases

You should back up Lotus Domino databases during off-peak hours and disable third-party Lotus Domino agents before you run the backup. The archived transaction logs are included automatically.

All Lotus Domino databases and transaction logs that reside on single or multiple volumes must be backed up by the same Backup Exec server. In addition, you should not back up a Lotus Domino server simultaneously from multiple Backup Exec servers.

See [“Backing up data”](#) on page 158.

Lotus Domino Agent backup options

The following option is available for Agent for Lotus Domino Server (Lotus Domino Agent) backup jobs. This option appears when you select **Lotus Domino** on the **Backup Options** dialog box, and when the backup definition for the Domino backup job includes the incremental backup method.

Table J-1 Lotus Domino Agent incremental backup method option

Item	Description
Mark archive logs for reuse	<p>Reuses the transaction log after it has been backed up.</p> <p>Backup Exec will not delete the transaction log. Selecting this option only indicates that the transaction log is ready to be reused after it has been backed up successfully; the Lotus Domino server actually deletes the transaction logs.</p> <p>This option is selected automatically when you select the full backup method. You cannot clear this option when you are using the full backup method.</p>

About automatic exclusion of Lotus Domino files during volume-level backups

If you select a volume that contains Lotus Domino data for backup, the Agent for Lotus Domino Server (Lotus Domino Agent) determines which Domino data should not be included in a volume level backup. For example, .ntf and .nsf files, nlo files, as well as any active log files, should not be part of the backup because they are opened for exclusive use by the Lotus Domino system. These files will be automatically excluded for backup by a feature called Active File Exclusion. If this exclusion did not happen during a non-snapshot backup, these files would appear as in use - skipped. If this exclusion did not happen during a snapshot backup, the files would be backed up in a possible inconsistent state, which could create restore issues.

See [“Backing up data”](#) on page 158.

See [“Lotus Domino Agent backup options”](#) on page 1070.

About supported Lotus Domino database configurations

You can back up the following types of Lotus Domino database configurations using the Agent for Lotus Domino Server (Lotus Domino Agent):

- **Domino Server Databases.**
Domino Server databases can be Logged or Unlogged, with their DAOS states being not-enabled, read-only, or enabled. DAOS cannot be enabled on Domino databases that do not use logging. Domino databases are located in a folder in

the Domino data directory, typically Lotus\Domino\Data, but may also be linked to the Domino data directory using Lotus Linked Databases.

The following types of Lotus Domino databases are supported:

- **Logged Domino Server Databases.**
A logged Domino Server database logs transactions for one or more Lotus databases. If transaction logging is enabled on the server, all database transactions go into a single transaction log.
- **Unlogged Domino Server Databases.**
An unlogged Domino Server database does not have transaction logging enabled, or the transaction logging has been disabled for specific server databases. Unlogged Domino Server databases will be backed up in their entirety when a full or incremental backup is performed, but the database can only be restored to the point of the latest database backup.
- **Local Databases.**
Lotus databases are considered Local when they cannot be found in the Domino data directory, cannot be shared, and cannot be logged. This type of database requires a backup of the database itself when using any of the Lotus Domino backup methods. The database can be restored only to the point of the latest database backup.

See [“About Lotus Domino transaction logs”](#) on page 1072.

See [“About selecting backup options for Lotus Domino databases”](#) on page 1070.

About Lotus Domino transaction logs

Lotus Domino has the ability to log transactions for one or more Lotus Domino databases. Lotus Domino databases are logged by default when transaction logging is enabled on the Lotus Domino server and the database is in the Domino data directory.

When transaction logging is enabled on the server, each Lotus Domino database is assigned a database instance ID (DBIID). Each transaction recorded in the log includes the DBIID, which is used to match transactions to the database during a restore.

A new DBIID may be assigned to the database when some Lotus Domino operations are performed. When the new DBIID is assigned, all new transactions recorded in the log use the new DBIID; however, previous transactions have the old DBIID and will not match the new DBIID for the database. To prevent data loss, it is recommended that a full backup be performed when a database receives a new DBIID since transactions with the old DBIID cannot be restored to the database. A

full backup includes all current transactions on the database and ensures that only the transactions with the new DBIID are needed to restore the database.

You can select only one logging style when transaction logging is enabled on the server.

Following are the two styles of logging for Lotus Domino databases:

- **Archive logging.**
This logging style produces a transaction log that is limited only by the capacity of your mass storage. Archive logging is the recommended logging style to be used with the Lotus Domino Agent since all the transaction logs can be backed up and marked for recycling. When the transaction logs are recycled the Lotus Domino server reuses the existing transaction logs after they are backed up to create space for new transaction logs.
- **Circular logging.**
This logging style reuses the log file after a specific log file size is reached. By reusing the log file you are saving resources; however, you are also limiting your recovery options because the database can only be recovered to the point of the last full backup. If the incremental backup method is selected for a backup job, a full backup of the changed databases is performed since transaction logs cannot be backed up.

Caution: When circular logging is enabled, the circular transaction log cannot be backed up, which could result in the loss of changes made to the database since the last backup was performed.

See [“About supported Lotus Domino database configurations”](#) on page 1071.

See [“About selecting backup options for Lotus Domino databases”](#) on page 1070.

About restoring Lotus Domino databases

To restore data, selections should be made in the Restore Wizard from the backup set that contains the Lotus Domino databases; the required transaction logs and the DAOS NLO files are automatically restored with the selected database.

Lotus Domino data is usually contained in the most recent backup set. However, some subsequent incremental backup jobs run after a full backup job may not contain data in the backup set because only the transaction log was backed up. If the data you want to restore is not located in the most recent backup set, check the previous backup sets until you find the data.

Note: If a new DBIID has been assigned to databases and you run an incremental backup, the data will be contained in the most recent backup set since transactions with the new DBIID will not match the old DBIID.

Note: When restoring Lotus Domino databases to Microsoft Cluster Server cluster, the virtual computer name or the virtual IP address of the Domino server should be used when browsing or making Domino database selections in the **Resource View** tab in the Restore Wizard.

Restore options for Lotus Domino databases are set using the Restore wizard. The wizard provides definitions for Domino-specific restore options.

When you select a Lotus Domino backup set to restore, all database files and necessary transaction logs are automatically restored. You can also choose to restore specific database files.

See [“Lotus Domino Agent restore options”](#) on page 1076.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

Restoring Lotus Domino databases

Restoring a Lotus Domino database is a three-part process.

Table J-2 Restoring a Lotus Domino database

Step	Description
Step 1	<p>Restore database files to the Domino server.</p> <p>During a restore of the Lotus Domino database, the existing database is taken offline and deleted, the database is restored, and changed records contained in the backup job are applied to the database.</p> <p>Note: Domino servers include databases with names such as <code>admin4.nsf</code>, <code>names.nsf</code>, and <code>busytime.nsf</code>. The Notes client computers include databases with names such as <code>bookmark.nsf</code>, <code>cache.dsk</code>, and <code>homepage.nsf</code>. These databases are critical and cannot be taken offline when the Domino server is running. In addition, you should only restore these databases in disaster recovery situation.</p> <p>If the database is unlogged or local, the database is brought back online. If the database is logged and multiple databases are being restored, the database name is added to a list for recovery. During the restore process, Backup Exec assigns a unique name to databases and then before databases are brought online, reassigns the original name. Changing the name during the restore process has no effect on restored databases.</p>
Step 2	<p>Restore DAOS-related NLO files that are missing.</p>

Table J-2 Restoring a Lotus Domino database (continued)

Step	Description
Step 3	<p>Run transaction logs to bring the database up-to-date.</p> <p>The internal Domino recovery process automatically begins after the DAOS NLO files are restored to the server. The database is restored to a point in time using transactions from the required transaction logs. Required transaction logs that were backed up and recycled are also included in the recovery process. After the recovery process completes, the Lotus Domino database is brought online.</p> <p>If you back up your Lotus Domino databases regularly, then restoring the most recent backup set containing the Lotus Domino data is all that is required to restore the most recent backups of your Lotus Domino databases.</p> <p>Note: If circular logging is enabled and both the databases and the Domino transaction logs are lost, the database can only be recovered to the point of the last full backup.</p>

Use the same procedures to restore a server in a Microsoft Cluster Server cluster that you use to restore a server in a non-clustered environment.

When restoring a Lotus Domino database to a MCSC cluster and a failover occurs during the restore operation, active restore jobs are paused for 15 minutes as they wait for existing connections to resolve themselves. If the restore job does not restart before the failover time-out period expires, the job fails. If this occurs, the restore job must be resubmitted.

See [“About restoring Lotus Domino databases”](#) on page 1073.

Lotus Domino Agent restore options

The following options are available for Agent for Lotus Domino Server (Lotus Domino Agent) restore jobs. These options appear when you run the **Restore Wizard**.

See [“Lotus Domino Agent backup options”](#) on page 1070.

Table J-3 Lotus Domino restore options

Item	Description
To the latest available backup set time	Uses the latest available backup set time when multiple backup sets are available for restore.
To a point in time in the transaction log up to and including the specified time	<p>Specifies the date and time to restore the database. The option is only available for logged databases when the archive logging style is set. Backup Exec will restore the Lotus Domino database you selected in the Restore selections dialog box and then automatically restore the necessary transaction logs required to bring the databases up to the date and time specified.</p> <p>If a point in time is not specified, the databases will be restored up to the last committed transactions in the log file.</p> <p>This option may require additional time since the archived transaction logs are also restored.</p>
Specify the number of seconds to wait for a database to go offline	Specifies the number of seconds for the restore process to wait for a database that is in use. When a Lotus database is restored it must first be taken offline. This will ensure that the database is not being accessed, closed, or deleted while the restore operation is being processed. If the database is still in use and cannot be taken offline after the specified wait time, the restore will fail.
Restore over existing database	Overwrites and restores the original database in the destination folder.

Table J-3 Lotus Domino restore options (*continued*)

Item	Description
Create a copy of the database	<p>Creates a copy of the original database during a restore. After creating a copy of the database, you can use the copied database to extract individual items such as emails at a later time. In creating the copy, Backup Exec assigns the copied database a new database ID and a replica ID. It does not take the original production database offline or alter it in any way.</p> <p>After the restore finishes, the copied database appears in the destination folder in the following format:</p> <p><OriginalDatabaseName>_<Date_Time in UTC format>.nsf. For example, when Backup Exec creates a copy of the User.nsf database during a restore, the copy of User.nsf database appears in the destination folder as User_2012Apr05_220547.nsf.</p>
Retain original database ID	Restores the original database ID.
Assign new database ID	Assigns new IDs to the database.
Assign new database ID and replica ID	Assigns new IDs to the database. A replica ID is used to synchronize two or more databases that are being replicated in the Lotus Domino environment. You can assign a new replica ID during a restore to prevent other databases under replication from overwriting the restored database files.

About redirecting restore jobs for Lotus Domino databases

Point in time restores cannot be redirected. Redirecting the restore of a DAOS-enabled Domino database does not restore the nlo files.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

See [“Backup Exec logon accounts”](#) on page 610.

About redirecting the restore of DAOS NLO files

You can restore DAOS NLO files without restoring the entire DAOS-enabled Domino database. When you restore the DAOS NLO files you must specify a redirection destination path. In most cases, the path points to the DAOS internal repository that you set when you configured Lotus Domino.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

See [“About the Lotus Domino Agent and the Domino Attachment and Object Service \(DAOS\)”](#) on page 1066.

How to prepare for disaster recovery on a Lotus Domino server

A Disaster Preparation Plan is necessary for restoring Lotus Domino databases efficiently and effectively in the event of a catastrophic failure. The goal is to minimize the time to recover. Developing a backup strategy for your Windows computers and Lotus Domino databases is the critical part of this plan.

When developing a strategy for backing up your Lotus Domino databases, consider the following recommendations:

- Keep linked databases on one volume. This allows Backup Exec to synchronize the databases before they are backed up.
- Back up active databases often. This reduces the amount of effort required to update the databases to the point following the most recent backup.
- Ensure that the notes.ini, cert.id, and <server>.id files are protected and available if a disaster occurs.
- Configure the DAOS prune period as recommended in the Lotus Domino documentation. However, Symantec recommends that you do not set the DAOS prune period for a period less than the time between two Domino backups.

See [“Recovering a Lotus Domino server from a disaster”](#) on page 1079.

See [“About disaster recovery of a Lotus Domino server using archive logging”](#) on page 1081.

See [“Recovering a Lotus Domino server that uses circular logging”](#) on page 1084.

Recovering a Lotus Domino server from a disaster

Recovering a Lotus Domino system can be performed in the following ways:

- Manually

For more information about doing a manual disaster recovery, click the following URL:

<http://entsupport.symantec.com/umi/V-290-706>

- By using Backup Exec's Simplified Disaster Recovery
See "[About Simplified Disaster Recovery](#)" on page 736.

When you recover a DAOS-enabled Domino server from a disaster, all of the .nlo files that each recovered Domino database references are automatically restored.

Note: Disaster recovery of a Lotus Domino server in a Microsoft Cluster Server cluster uses the same steps as recovering a Domino server in a non-clustered environment.

Use the following steps as a guide when you want to do a disaster recovery operation on a Lotus Domino server.

Table J-4 Steps to take to recover a Lotus Domino server from a disaster

Step	Description
Step 1	Recover the Windows computer.
Step 2	Disable the monitor change journal. See " Disabling the monitor change journal " on page 1081.
Step 3	Recover or re-install Lotus Domino to the same location as before the disaster occurred. All Lotus Domino system data must be recovered. System data includes log.nsf, names.nsf, template files, notes.ini, mail.box, and ID files. Note: If transaction logging is enabled, you must run a disaster recovery operation that is based on the style of logging selected on the Lotus Domino server. After rebuilding the server, you can restore the databases from your most recent backup.
Step 4	Re-enable the monitor change journal. See " Re-enabling the monitor change journal " on page 1083.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

See [“Recovering a Lotus Domino server that uses circular logging”](#) on page 1084.

About disaster recovery of a Lotus Domino server using archive logging

If the active transaction log is lost, you can recover the database only up to the transactions contained in the last transaction log.

However, if all of the transaction logs are lost, you must have the following to recover the database:

- An up-to-date Notes.ini file from the Lotus Domino server.
- The backups of the database.
- All archived log extents.

In addition, if the monitor change journal is enabled, you must disable it in the registry before beginning the Lotus Domino server recovery.

See [“Disabling the monitor change journal”](#) on page 1081.

Disabling the monitor change journal

Use the following steps to disable and enable the monitor change journal in the registry. Then you can recover the Lotus Domino server, databases, and transaction logs.

See [“About disaster recovery of a Lotus Domino server using archive logging”](#) on page 1081.

To disable the monitor change journal

- 1 Open the registry and browse to the following key:
HKEY_LOCAL_MACHINE\SOFTWARE\SYMANTEC\BACKUP EXEC FOR WINDOWS\BACKUP EXEC\ENGINE\DOMINO
- 2 Disable monitor change journal by setting the value of the key Enable Change Journal to 0.
- 3 Restart the Symantec Backup Exec Agent for Windows systems.

Recovering the Lotus Domino server, databases, and transaction logs when archive logging is enabled

Use the following steps to recover the Lotus Domino server, databases, and transaction logs when archive logging is enabled.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

See [“About disaster recovery of a Lotus Domino server using archive logging”](#) on page 1081.

To recover the Lotus Domino server, databases, and transaction logs when archive logging is enabled

- 1 On the **Backup and Restore** tab, in the **Restores** group, click **Restore**.
- 2 Select **Files, folders, or volumes**, and then click **Next**.
- 3 Select **File and folder backups from a backup set**, and then click **Next**.
- 4 Select a backup set that contains the non-database Domino server files (*.id and notes.ini) that you want to restore. Do not launch the Domino server after reinstalling it.

If necessary, reinstall the Domino server but do not configure it. Restore the non-database Domino files, which include the notes.ini file and the *.id files. Use the same directory structure, directory location, and logdir path as was created during the original installation.

- 5 Click **Next**.
- 6 On the **Where do you want to restore the data** wizard panel, select **To the original location**, and then click **Next**.
- 7 On the **How do you want to maintain file integrity and hierarchy for the restored data?** wizard panel, select **Restore over existing files**.

Make additional selections on the wizard panel as desired.

- 8 Click **Next**.
- 9 On the **How do you want to restore operating system features?** wizard panel, accept the default selections or make selections as desired and then click **Next**.
- 10 On the **What additional tasks do you want to perform before and/or after a restore?** wizard panel, make selections as desired, and then click **Next**.
- 11 On the **What job name and schedule do you want to use?** wizard panel, name the job and schedule the job to run, or click **Run now** to run the job immediately.

- 12 Click **Next**.
- 13 On the **Restore Summary** panel, click **Finish**.
- 14 After the restore job finishes, navigate to the location of the notes.ini file.
- 15 Using a text editor, change the TRANSLOG_Status setting in the notes.ini file on the Domino server to 0.

For example, TRANSLOG_Status=0
- 16 Using the Backup Exec Restore Wizard, run a redirected restore of the last transaction log backed up before the loss of the active transaction log.
- 17 Verify the transaction log restore was successful.
- 18 Restart the Symantec Backup Exec Agent for Windows Systems.
- 19 Delete all transaction logs except the transaction log restored in step 16, from the Domino transaction log directory.
- 20 Using a text editor, change the notes.ini file for the Domino server to match the following:

TRANSLOG_Recreate_Logctrl=1

TRANSLOG_Status=1
- 21 Run a full restore of the Domino databases or a point-in-time state within the archived log extents.

Backup Exec automatically restores all DAOS NLO files along with the DAOS-enabled databases from full backups. In addition, Domino automatically recreates both the `daos.cfg` file and the `daoscat.nsf` when you restart the Domino server.

After the full restore finishes, the TRANSLOG_Logctrl parameter in the notes.ini file is reset to 0.
- 22 Restore DAOS NLO files from incremental backups.
- 23 Start the Domino server.

Disaster recovery is complete.

If monitor change journal was disabled before beginning the disaster recovery process, you must re-enable it.

See ["Re-enabling the monitor change journal"](#) on page 1083.

Re-enabling the monitor change journal

Use the following steps to re-enable the monitor change journal.

See [“About disaster recovery of a Lotus Domino server using archive logging”](#) on page 1081.

See [“Lotus Domino Agent restore options”](#) on page 1076.

To re-enable monitor change journal

- 1 Open the registry and browse to the following key:
HKEY_LOCAL_MACHINE\SOFTWARE\SYMANTEC\BACKUP EXEC FOR WINDOWS\BACKUP EXEC\ENGINE\DOMINO
- 2 Enable monitor change journal by setting the value of the key Enable Change Journal to 1.
- 3 Restart the Symantec Backup Exec Agent for Windows Systems.

Recovering a Lotus Domino server that uses circular logging

If circular logging is enabled and the transaction log is lost, the Domino database can only be recovered to the point of the last backup.

To recover a Lotus Domino server that uses circular logging

- 1 If the monitor change journal is enabled, you must disable it in the registry before beginning the Lotus Domino server recovery.
See [“Disabling the monitor change journal”](#) on page 1081.
- 2 Restore or reinstall the Lotus Domino server program directory (**excluding** the notes.ini, cert.id, and <server>.id files) to the same location as before the disaster occurred.
- 3 Check that the log directory (logdir) is created and does not contain old files.

If the log directory was not created, recreate the directory to the same location as before the disaster occurred.

Do not start the Lotus Domino server after performing the previous steps.

- 4 Restore the notes.ini, cert.id, and <server>.id files from the last full backup of the Lotus Domino server program directory to the same location as before the disaster occurred.

For this restore, select **Restore over existing files** option on the **How do you want to maintain file integrity and hierarchy for the restored data?** page of the **Restore Wizard**.

- 5 To have Lotus Domino create the circular log file in the log directory when the server starts, set the following parameter in the notes.ini file:

```
translog_path=logdir
```

- 6 Restart the Symantec Backup Exec Agent for Windows Systems.
- 7 Use the Lotus Domino Agent to restore the databases to the Domino data directory.

Backup Exec automatically restores all DAOS NLO files along with the DAOS-enabled databases from full backups. In addition, Domino automatically recreates both the daos.cfg file and the daoscat.nsf file when you restart the Domino server.

- 8 Restore databases and DAOS NLO files from incremental backups.
- 9 Start the Domino server.

Disaster recovery is complete.

- 10 If monitor change journal was disabled before beginning the disaster recovery process, you must enable it again.

See [“Re-enabling the monitor change journal”](#) on page 1083.

Symantec Backup Exec Agent for Microsoft Active Directory

This appendix includes the following topics:

- [About the Agent for Microsoft Active Directory](#)
- [Requirements for the Agent for Microsoft Active Directory](#)
- [About backing up Active Directory and ADAM/AD LDS](#)
- [About restoring individual Active Directory and ADAM/AD LDS objects](#)

About the Agent for Microsoft Active Directory

The Symantec Backup Exec Agent for Microsoft Active Directory uses full backups for which Granular Recovery Technology (GRT) is enabled to let you restore individual Active Directory objects and attributes without performing an authoritative or non-authoritative full restore. You can also restore individual Active Directory Application Mode (ADAM) and Active Directory Lightweight Directory Services (AD LDS) objects and attributes.

The Agent for Microsoft Active Directory works with backups of the Windows System State, where Active Directory is installed, and with ADAM/AD LDS. When you back up the Windows System State, the Active Directory is included in the backup job because Active Directory is a component of Windows System State. You can also use the Agent for Microsoft Active Directory to restore individual ADAM/AD LDS objects and attributes. If multiple ADAM/AD LDS instances are backed up, each instance appears under the Active Directory Application Mode node.

Installing the Agent for Microsoft Active Directory

The Agent for Microsoft Active Directory is installed as part of the Agent for Applications and Databases.

See “[Installing additional Backup Exec options to the local Backup Exec server](#)” on page 76.

See “[About restoring individual Active Directory and ADAM/AD LDS objects](#)” on page 1092.

Requirements for the Agent for Microsoft Active Directory

Review the following requirements for the Agent for Microsoft Active Directory:

- The Agent for Windows must be installed on the computer on which the Active Directory is installed.
- For specific operating system requirements for the Agent for Microsoft Active Directory, see the Backup Exec Software Compatibility List at the following URL: <http://entsupport.symantec.com/umi/v-269-1>
- Ensure that the following option is selected on the backup job properties: **Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual objects from Active Directory backups (not supported for Read-Only Domain Controllers)**. This option is selected by default. This option must be selected so that you can restore individual attributes and properties from full Active Directory and ADAM/AD LDS backups.

Note: You cannot back up databases to devices that are attached to a computer on which the Remote Media Agent for Linux is installed.

- To perform a GRT-enabled backup of a Windows Server 2012 R2 Active Directory Application Server, you must use a Backup Exec server that runs Windows Server 2012 R2.

Following are the requirements for the restore of individual objects and attributes from backup jobs that used the Agent for Microsoft Active Directory:

- You must have a full backup of the ADAM/AD LDS or the Windows System State (where Active Directory is installed).
- You must use a version of the Windows operating system that supports minifilter drivers on the Backup Exec server that runs the restore job.

- You must designate a location on the Backup Exec server where Backup Exec can temporarily place the objects and attributes that are being restored when you restore from tape.

See [“About restoring individual Active Directory and ADAM/AD LDS objects”](#) on page 1092.

About backing up Active Directory and ADAM/AD LDS

Symantec recommends that you run Active Directory and ADAM/AD LDS backups to disk storage and then copy the backups to tape. This strategy provides you with shorter backup windows and disk storage provides the most efficient method of storage for GRT-enabled backups, and the most efficient method of restore. This method also lets you administer Active Directory or ADAM/AD LDS without requiring the individual cataloging of the backed up objects and properties.

For example, if you back up to tape, you must create a temporary hard disk staging location on a local NTFS volume to restore individual items from GRT-enabled backups on tape. The data is first copied from tape to the temporary staging location before it can be restored. As such, a restore from tape takes more time. For best results, you should specifically select disk storage when you configure your GRT-enabled backup jobs.

Note: You cannot back up databases to devices that are attached to a computer on which the Remote Media Agent for Linux is installed.

Granular Recovery Technology (GRT) lets you restore individual objects and attributes from Active Directory and ADAM/AD LDS backups without performing an authoritative or non-authoritative full restore. To restore individual items, you must enable the Granular Recovery Technology feature when you create a backup job. You should review the requirements for a GRT-enabled backup before you configure it. To perform a GRT-enabled backup of a Windows Server 2012 R2 Active Directory Application Server, you must use a Backup Exec server that runs Windows Server 2012 R2.

When you back up any Windows Active Directory or ADAM/AD LDS application database directly to tape, objects and properties that are added or deleted during the backup will not match the individual objects and properties that are available for restore from the backup set. The backup of the database is a snapshot backup of the live Active Directory or ADAM/AD LDS database and the cataloging of the individual Active Directory or ADAM/AD LDS objects occurs after the snapshot is performed. Since the catalog operation catalogs objects and properties from the live Active Directory or ADAM/AD LDS database, object and property changes can occur after the snapshot was taken.

See [“Backing up data”](#) on page 158.

See [“Editing backup definitions”](#) on page 206.

See [“Adding a stage to a backup definition”](#) on page 219.

See [“Microsoft Active Directory backup job options”](#) on page 1091.

See [“Granular Recovery Technology”](#) on page 590.

Editing options for Active Directory and ADAM/AD LDS backup jobs

You can edit the default settings for all Active Directory and ADAM/AD LDS backup jobs.

See [“About restoring individual Active Directory and ADAM/AD LDS objects”](#) on page 1092.

To edit options for Active Directory and ADAM/AD LDS backup jobs

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Job Defaults**.
- 2 Select a backup option.
- 3 On the left, click **Microsoft Active Directory**.

- 4 Edit the following backup options for the Agent for Microsoft Active Directory as appropriate:

Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual objects from Active Directory backups (not supported for Read-Only Domain Controllers)

Enables the restore of individual items from full backups of the Active Directory or ADAM/ AD LDS.

This option is selected by default. This option must be selected so that you can restore individual attributes and properties from full Active Directory and ADAM/AD LDS backups.

Note: You cannot restore individual objects and attributes from Active Directory backups for a read-only domain controller (RODC). You should do GRT backups and restores of the Active Directory to a writable, centralized data center domain controller.

Ensure that you meet the requirements for Granular Recovery Technology.

To perform a GRT-enabled backup of a Windows Server 2012 R2 Active Directory Application Server, you must use a Backup Exec server that runs Windows Server 2012 R2.

See [“Granular Recovery Technology”](#) on page 590.

Perform consistency check before backup when using Microsoft Volume Shadow Copy Service (VSS) snapshot provider (Windows Server 2008 and above)

Checks snapshots for data corruption. This option applies only to snapshots that are performed by the Microsoft Volume Shadow Copy Services (VSS).

Continue with backup if consistency check fails

Enables the backup job to continue even if the consistency check fails. You may want the job to continue if a backup of the database in its current state is better than no backup at all. Or you may want the job to continue if you back up a large database that may have only a small problem.

- 5 Click **OK**.

Microsoft Active Directory backup job options

You can edit the default settings for Active Directory and ADAM/AD LDS backup jobs.

Note: Only full backups of Active Directory are allowed.

See [“Editing options for Active Directory and ADAM/AD LDS backup jobs”](#) on page 1089.

Table K-1 Microsoft Active Directory backup default options

Item	Description
Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual objects from Active Directory backups (not supported for Read-Only Domain Controllers)	<p>Enables the restore of individual items from full backups of the Active Directory or ADAM/AD LDS.</p> <p>This option is selected by default. This option must be selected so that you can restore individual attributes and properties from full Active Directory and ADAM/AD LDS backups.</p> <p>Note: You cannot restore individual objects and attributes from Active Directory backups for a read-only domain controller (RODC). You should do GRT backups and restores of the Active Directory to a writable, centralized data center domain controller.</p> <p>Ensure that you meet the requirements for Granular Recovery Technology.</p> <p>To perform a GRT-enabled backup of a Windows Server 2012 R2 Active Directory Application Server, you must use a Backup Exec server that runs Windows Server 2012 R2.</p> <p>See “Granular Recovery Technology” on page 590.</p>
Perform consistency check before backup when using Microsoft Volume Shadow Copy Service (VSS) snapshot provider (Windows Server 2008 and above)	<p>Checks snapshots for data corruption. This option applies only to snapshots that are performed by the Microsoft Volume Shadow Copy Services (VSS).</p>

Table K-1 Microsoft Active Directory backup default options (*continued*)

Item	Description
Continue with backup if consistency check fails	Enables the backup job to continue even if the consistency check fails. You may want the job to continue if a backup of the database in its current state is better than no backup at all. Or you may want the job to continue if you back up a large database that may have only a small problem.

See [“About backing up Active Directory and ADAM/AD LDS”](#) on page 1088.

About restoring individual Active Directory and ADAM/AD LDS objects

When you restore Active Directory and ADAM/AD LDS objects from tape, you must specify an on-disk staging location where the objects will be placed prior to being restored. The staging location must be a path on a local NTFS volume on the Backup Exec server running the restore job and the Backup Exec service account must also have access to it.

By default, the Agent for Microsoft Active Directory restores deleted Active Directory or ADAM/AD LDS objects from the Active Directory Deleted Objects container if their tombstone lifetimes have not passed.

When objects in Active Directory are deleted, they are removed from their current Active Directory or ADAM/AD LDS container, converted into tombstones, and then placed in the Active Directory Deleted Objects container where their tombstone lifetime is monitored. After their tombstone lifetime passes, the tombstones are purged from the Active Directory Deleted Objects container, which permanently deletes the objects from the Active Directory and ADAM/AD LDS databases.

The Agent for Microsoft Active Directory lets you restore tombstoned objects from the Active Directory Deleted Objects container in the following situations:

- Their tombstone lifetimes have not passed.
- They have not been purged from the Deleted Objects container.
- You are restoring to a Windows Server 2003/2008/2008 R2/XP Professional x64 Edition system.

When you restore data with the ADAM Writer, Backup Exec stops the service for the ADAM instance you want to restore before the restore job starts. However, Backup Exec does not restart the ADAM service when the restore job completes

because post-processing jobs, such as authoritative restores using Adamutil.exe, may be needed. You must restart the ADAM service. If Backup Exec cannot stop the ADAM service or if Backup Exec cannot restore all of the ADAM files, the ADAM restore fails.

When you restore Active Directory user objects, you must reset the object's user password and then re-enable the object's user account. For ADAM/AD LDS user objects, you must reset the object's user password and then re-enable the object's user account. For Active Directory user objects, use the Microsoft Active Directory Users and Computers application. For ADAM/AD LDS user objects, use ADSI Edit.

For Active Directory computer objects, you must reset the object's account.

Note: Some objects in the Active Directory Configuration Partition node cannot be reanimated from the Active Directory Deleted Objects container. However, recreated objects may not be recognized by some applications.

For more information, see your Microsoft Active Directory documentation.

Resetting the Active Directory computer object and the computer object account

In Active Directory, computer objects are derived from user objects. Some attributes that are associated with a computer object cannot be restored when you restore a deleted computer object. The attributes can only be restored if the attributes were saved through schema changes before the computer object was originally deleted. Because computer object credentials change every 30 days, the credentials from the backup may not match the credentials that are stored on the actual computer.

Note: To reset a computer object, you must use the Microsoft Active Directory Users and Computers application.

For more information on resetting a computer object, see your Microsoft Active Directory Users and Computers application documentation.

If a computer object's **userAccountControl** attribute was not preserved before the object was deleted, you must reset the object's account after you restore the object.

To reset the Active Directory computer object account

- 1 Remove the computer from the domain.
- 2 Re-join the computer to the domain. The SID for the computer remains the same since it is preserved when you delete a computer object. However, if the object's tombstone expires and a new computer object is recreated, the SID is different.

Recreating purged Active Directory and ADAM/AD LDS objects

You can attempt to recreate deleted Active Directory objects and ADAM/LDS objects after they have been purged from the **Active Directory Deleted Objects** container by restoring the object from a previous Active Directory backup.

You can attempt to recreate the deleted objects if their tombstone lifetimes have passed and the objects have been purged from the Active Directory Deleted Objects container.

However, you should be aware of the following:

- Most applications will not recognize a recreated object since recreated objects are not identical to the original deleted object. Recreated objects are assigned new global unique identifiers (GUIDs) and security identifiers (SIDs) that cannot be identified by the applications that created the original object.
- Attributes created by the Windows operating system cannot be recreated when a purged object is recreated. Hence, objects that rely on attributes set by the operating system will not be recognized by Windows when the objects are recreated.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

See [“About restoring individual Active Directory and ADAM/AD LDS objects”](#) on page 1092.

Symantec Backup Exec Central Admin Server Option

This appendix includes the following topics:

- [About the Central Admin Server Option](#)
- [Requirements for installing CASO](#)
- [How to choose the location for CASO storage and media data](#)
- [About installing the Central Admin Server Option](#)
- [Upgrading an existing CASO installation](#)
- [Changing a Backup Exec server to a central administration server](#)
- [Changing a Backup Exec server to a managed Backup Exec server](#)
- [Deleting a managed Backup Exec server from a CASO environment](#)
- [Renaming a central administration server](#)
- [Renaming a managed Backup Exec server](#)
- [How to reduce network traffic in CASO](#)
- [CASO distributed, centralized, and replicated catalog locations](#)
- [Changing the settings for a managed Backup Exec server](#)
- [What happens when CASO communication thresholds are reached](#)
- [Alerts and notifications in CASO](#)

- [Enabling managed Backup Exec servers to use any available network interface card](#)
- [About job delegation in CASO](#)
- [About adding storage devices in a CASO environment](#)
- [How data lifecycle management \(DLM\) works in a CASO environment](#)
- [Obtaining media audit information for a managed Backup Exec server](#)
- [How to use Backup Exec server pools in CASO](#)
- [How centralized restore works in CASO](#)
- [About recovering failed jobs in CASO](#)
- [Pausing or resuming a managed Backup Exec server](#)
- [Stopping or starting Backup Exec services for a managed Backup Exec server](#)
- [Viewing managed Backup Exec server properties](#)
- [Viewing the settings for a central administration server](#)
- [Disaster recovery in CASO](#)
- [Troubleshooting CASO](#)
- [Running the Backup Exec Utility for CASO operations](#)
- [Uninstalling Backup Exec from the central administration server](#)
- [Uninstalling Backup Exec from a managed Backup Exec server](#)

About the Central Admin Server Option


The Symantec Backup Exec Central Admin Server Option (CASO) enables a central administration server to delegate jobs to managed Backup Exec servers across the network. Job delegation is the automatic load balancing of jobs across available managed Backup Exec servers in the CASO environment. If your organization includes more than one Backup Exec server, you can benefit from using CASO.

All backup information in the CASO environment can be centralized on the central administration server. The managed Backup Exec servers perform the actual processing of backup and restore jobs. You create jobs on the central administration server and then delegate the jobs to run on a managed Backup Exec server. The jobs are delegated, or load-balanced, across the available storage devices on the managed Backup Exec server. Multiple Backup Exec servers can share a storage

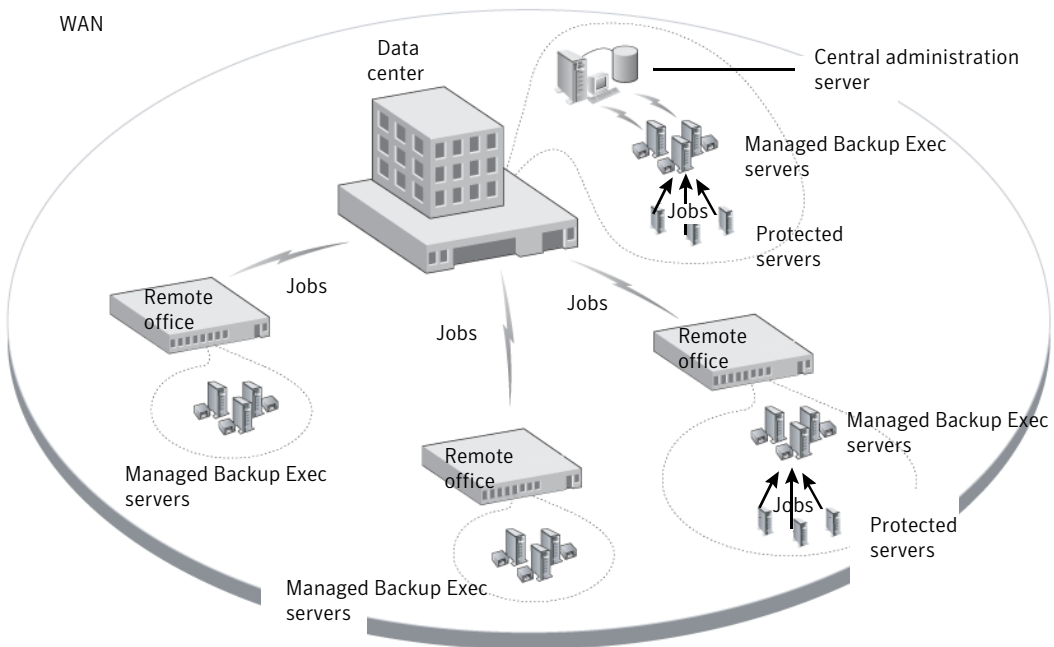
device when the sharing option is enabled. Centralized restore jobs can also be delegated to managed Backup Exec servers. Additionally, the central administration server can function as a managed Backup Exec server and process delegated jobs. A managed Backup Exec server can also run the jobs that are created locally at its local administration console.

The following graphic shows a local area network (LAN) environment with a central administration server and several managed Backup Exec servers.

Figure L-1 CASO-configured Backup Exec environment - LAN



The same communications that occur over a LAN between the central administration server and the managed Backup Exec servers take place over a WAN.

Figure L-2 CASO-configured Backup Exec environment - WAN

See [“Sharing storage devices”](#) on page 448.

See [“How to choose the location for CASO storage and media data”](#) on page 1099.

See [“About installing the Central Admin Server Option”](#) on page 1101.

See [“Upgrading an existing CASO installation”](#) on page 1110.

Requirements for installing CASO

The system requirements for the Central Admin Server Option (CASO) are the same as the minimum requirements for Backup Exec, with the exception of RAM. However, processor speed, memory, and disk space requirements may increase based on the number of managed Backup Exec servers, the number of servers being backed up, and the amount of catalog storage space that is required.

On the computer on which you install the central administration server, 1 GB RAM is required, although 2 GB RAM or more is recommended for better performance. Other applications on the Backup Exec server also require a certain amount of physical RAM to function properly. The requirements for RAM may also increase

when the central administration server manages more Backup Exec servers or tape hardware.

Before you install CASO, do the following:

- Ensure that you have administrative rights on computers on which you want to install CASO.
- Ensure that when you install CASO on Backup Exec servers in multiple domains, the Backup Exec service account is in the trusted domain, and has administrative rights on all of the Backup Exec servers that you want to use as managed Backup Exec servers.

If the Backup Exec Database for the central administration server is installed on a SQL Server instance on a different computer, the account must be a domain account with local administrative privileges on that computer as well.

- Ensure that the central administration server and the managed Backup Exec servers are part of a domain or domains. CASO is not supported in a workgroup.
- Use only NetBIOS computer names for managed Backup Exec servers and central administration servers. You cannot enter fully qualified domain names or IP addresses as server names.
- Ensure that you have the appropriate licenses for Backup Exec. A license for Backup Exec is required in addition to a license for CASO.
- Ensure that when you install a managed Backup Exec server, both the logged-in user and the Backup Exec service account for the managed Backup Exec server have administrative rights on the central administration server.

See [“System requirements for Backup Exec”](#) on page 63.

See [“About installing the Central Admin Server Option”](#) on page 1101.

How to choose the location for CASO storage and media data

During the installation of the managed Backup Exec server feature, you can choose the location of the managed Backup Exec server's storage and media data.

The following table compares how CASO tasks are performed depending on the location of the managed Backup Exec server's storage and media data:

Table L-1 Comparison of CASO tasks

Task	Storage and media data on the central administration server	Storage and media data on the managed Backup Exec server
Delegate jobs from the central administration server to the managed Backup Exec server	Yes	No. Instead, you can create jobs on the central administration server, and then copy them to the managed Backup Exec server.
Manage storage devices and media on the managed Backup Exec server from the central administration server	Yes	No
Hold, delete, run, cancel, and change the priority of copied jobs from the central administration server if the option to monitor jobs is enabled on the managed Backup Exec server	Yes	Yes
Monitor the jobs that are created on the local managed Backup Exec server if the option to monitor jobs is enabled on the managed Backup Exec server	Yes	Yes
Send job status updates, job logs, and job histories to the central administration server if the option to monitor jobs is enabled on the managed Backup Exec server	Yes	Yes
Centralize, distribute, or replicate the catalog	Yes	No Only a distributed catalog location can be selected.

Table L-1 Comparison of CASO tasks *(continued)*

Task	Storage and media data on the central administration server	Storage and media data on the managed Backup Exec server
Run centralized restore	Yes	Yes You can browse the backup sets and run restore operations for the managed Backup Exec server from the central administration server.

Note: In a CASO environment, you can add an NDMP Server only to a central administration server or a managed Backup Exec server on which the storage and media database is located.

- See [“Upgrading an existing CASO installation”](#) on page 1110.
- See [“About the Central Admin Server Option”](#) on page 1096.
- See [“Running the Backup Exec Utility for CASO operations”](#) on page 1153.

About installing the Central Admin Server Option

The Central Admin Server Option is installed as part of the Enterprise Server Option. After you enter license information for the Enterprise Server Option, on the **Configure Options** panel, you must expand the **Backup Exec Options** item, and then expand the **Enterprise Server Option** item to select the Central Admin Server Option for installation. When you select the Central Admin Server Option for installation, then the central administration server is installed. After the central administration server is installed, you can install managed Backup Exec servers.

Note: You must use the custom installation option in the installation wizard to install CASO. The typical installation option does not support the installation of CASO.

- See [“Push-installing a managed Backup Exec server from the central administration server ”](#) on page 1102.
- Before you start the installation, review the information about the location of storage and media data.
- See [“How to choose the location for CASO storage and media data ”](#) on page 1099.

A managed Backup Exec server may be installed outside the firewall that the central administration server is installed in or in a different firewall. However, you must set up a static port for SQL Server and create an alias for the managed Backup Exec server.

See [“Installing a managed Backup Exec server across a firewall”](#) on page 1107.

Push-installing a managed Backup Exec server from the central administration server

After you install the central administration server, you can push-install the managed Backup Exec server feature to a standalone server.

Before you install a managed Backup Exec server, decide where to locate the storage and media database for it. During the installation of the managed Backup Exec server, you can choose the location of the managed Backup Exec server's storage and media data. Your choice affects how you can manage jobs in the CASO environment.

See [“How to choose the location for CASO storage and media data”](#) on page 1099.

To push-install a managed Backup Exec server from the central administration server

- 1 From the central administration server, click the Backup Exec button, and then select **Installation and Licensing**.
- 2 Select **Install Agents and Backup Exec Servers on Other Servers**.
- 3 In the installation wizard, click **Add**, and then select either **Add a Single Computer** or **Add Multiple Computers with the Same Settings**.
- 4 Select **Symantec Backup Exec**, and then click **Next**.
- 5 In the **Remote computer** field, type the name of the managed Backup Exec server that you want to add, or click **Browse Remote Computers** to locate the server.
- 6 Click **Add to List**.

This option is not necessary if you selected **Add a Single Computer** in step 3.

- 7 Under **Remote computer credentials** , complete the fields as follows:

User Name	Type the user name for an account that has administrative rights on the remote computer.
Password	Type the password for an account that has administrative rights on the remote computer.
Domain	Select the domain in which the remote computer is located.

- 8 Click **Next**.

- 9 Select one of the following methods to enter license keys:

To enter serial numbers manually	In the Enter a serial number field, type a serial number from your sales certificate, and then click Add .
To import a Symantec license file	Click Import Symantec License File , and then navigate to the location of your .slf file.
To install a trial version	Do not enter serial numbers or import license files. Proceed to the next step.

- 10 Click **Next**.

- 11 After your serial numbers are validated, click **Next**.

- 12 On the list of features to install, expand **Backup Exec 2014**, and then select **Managed Backup Exec server**.

- 13 Do one of the following:

To change the directory where the Backup Exec files are installed	In the Destination Folder field, type the name of the directory.
To accept the default directory (recommended)	Proceed to the next step.

Symantec recommends that you do not select a mount point as the destination directory because if you delete the mount point, Backup Exec is uninstalled.

- 14 Click **Next**.

- 15 Provide a user name, password, and domain for an Administrator account that the Backup Exec system services can use, and then click **Next**.

- 16 On the **Choose SQL Server** panel, choose the location to store the Backup Exec Database, and then click **Next**.
- 17 In the **Central Administration Server** field, type the name of the central administration server that will manage this managed Backup Exec server.

Use only NetBIOS computer names for managed Backup Exec servers and central administration servers. You cannot enter fully qualified domain names or IP addresses as server names.
- 18 Select from the following options to determine how storage devices and data are managed:

See [“Managed Backup Exec Server Configuration options”](#) on page 1106.

Centrally managed Backup Exec server Select this option to enable the central administration server to manage this Backup Exec server, its storage devices, media, and job delegation. This option also enables this Backup Exec server to share storage devices with other managed Backup Exec servers.

Unrestricted access to catalogs and backup sets for restore Select this option to enable this managed Backup Exec server to have unrestricted access to all centrally stored catalogs. This option also enables this managed Backup Exec server to restore data from any backup set on any storage devices that it shares.

This option can be selected only if the **Centrally managed Backup Exec server** option is selected. Selecting both of these options enables the central administration server to have the greatest amount of control over this managed Backup Exec server.

Locally managed Backup Exec server Select this option to enable the central administration server to monitor this managed Backup Exec server and create restore jobs for it. However, the server and its devices, media, and backup jobs are controlled locally.

- 19 Click **Next**.
- 20 Select the Symantec device drivers that you want to use, and then click **Next**.

21 After Backup Exec validates the remote computers, you can change the list in any of the following ways:

To manually add one remote computer	Click Add , and then click Add a Single Server .
To manually add multiple remote computers	Click Add , and then click Add Multiple Servers with the Same Settings .
To add multiple remote computers by importing an existing list of computers	<p>Click Import and Export, and then select one of the following options:</p> <ul style="list-style-type: none">■ Select Import from File to enable Backup Exec to add the names of the remote computers from a selected list.■ Select Import Servers Published to this Backup Exec Server to enable Backup Exec to add the names of all the remote computers that are set up to publish to this Backup Exec server. <p>You must enter remote computer logon credentials for the list of remote computers.</p>
To change the product that you selected to install or to change other properties you selected for this installation	Select the remote computer that you want to change, and then click Edit .
To delete a remote computer from the list	Select the remote computer that you want to delete, and then click Delete .
To save this list of remote computers and the associated remote computer logon credentials	<p>Verify that Save the server list for future remote install sessions is checked.</p> <p>This option enables the names of all of the remote computers and their credentials to be added automatically the next time you want to install Backup Exec or options to these remote computers.</p>
To save this list of remote computers to an XML file	<p>Click Import and Export, and then click Export to File.</p> <p>You can select the location to save the XML file. This option is useful if you want to use the same list for multiple Backup Exec servers. When you import the list, you must re-enter the remote computer logon credentials.</p>

To fix the errors that were located during the validation Right-click the name of the computer, and then click **Fix Error**.

To enable Backup Exec to attempt to re-validate an invalid remote computer Right-click the name of the computer, and then click **Retry Validation**.

22 After all of the computers are validated, click **Next**.

23 Read the Backup Exec installation review, and then click **Install**.

24 Click **Next**, and then click **Finish**.

If you did not restart the remote computer, you may need to do it now in order for the configuration to take effect.

Managed Backup Exec Server Configuration options

The following information can help you to determine which configuration options to choose when you install a managed Backup Exec server.

The following information applies if you select the option **Centrally managed Backup Exec server** and also select the option **Unrestricted access to catalogs and backup sets for restore**:

- This Backup Exec server becomes a managed Backup Exec server.
- A persistent network connection is required between the managed Backup Exec server and the central administration server.
- The catalogs are centralized and are stored on the central administration server. Note that this combination of options may not be suitable if you have a low-bandwidth network connection to the central administration server.
- This managed Backup Exec server can access and restore backup sets for all storage devices that it shares with other Backup Exec servers.
- The backup jobs that are created on the central administration server can be load-balanced and delegated to this managed Backup Exec server.
- A rolling upgrade cannot be performed with this configuration. This managed Backup Exec server must be upgraded at the same time as the central administration server.

The following information applies if you select the option **Centrally managed Backup Exec server**, but do not select the option **Unrestricted access to catalogs and backup sets for restore**:

- This Backup Exec server becomes a managed Backup Exec server.
- A persistent network connection is required between the managed Backup Exec server and the central administration server.

- The catalogs are in distributed mode by default, but can be changed. The catalogs for the jobs that run on this managed Backup Exec server are stored locally.
- This managed Backup Exec server can access and restore any backup set that is stored on the storage devices that it hosts, regardless of which Backup Exec server ran the backup job. However, for the shared storage devices that other Backup Exec servers host, this managed Backup Exec server can access and restore only the backup sets that were created from backup jobs it ran.
- The backup jobs that are created on the central administration server can be load-balanced and delegated to this managed Backup Exec server.
- This option is recommended for use with the private cloud configuration.

The following information applies if you select the option **Locally managed Backup Exec server**:

- This Backup Exec server becomes a managed Backup Exec server.
- A persistent network connection is not required between the managed Backup Exec server and the central administration server. Therefore, this option may be useful when a very low-bandwidth connection exists between the managed Backup Exec server and the central administration server. It may also be useful when the managed Backup Exec server cannot always be connected to the central administration server.
- The catalogs are in distributed mode by default. The catalogs for the jobs that run on this managed Backup Exec server are stored locally.
- The central administration server does not delegate jobs to this managed Backup Exec server.
- This managed Backup Exec server cannot be used in a private cloud configuration.

See [“Push-installing a managed Backup Exec server from the central administration server”](#) on page 1102.

Installing a managed Backup Exec server across a firewall

A managed Backup Exec server may be installed outside the firewall that the central administration server is installed in or in a different firewall.

The following rules apply to the managed Backup Exec servers that are installed across a firewall:

- Port 3527 must be open in both directions to enable communication for the Backup Exec Server service.

- Port 10000 must be open for the Agent for Windows, which allows browsing for remote selections.
- A SQL port must be open in both directions to the central administration server's database to enable database connections.
- A static port must be used.

The Backup Exec SQL instance is configured by default to use a dynamic port. Each time SQL Server is started, the port number can change. You must change the dynamic port to a static port. After you change the configuration of the port from dynamic to static, you must add the static port to the Windows Firewall Exceptions list.

See your Windows operating system documentation.

See ["To change the dynamic port for a SQL Express instance to a static port and create an alias for the managed Backup Exec server"](#) on page 1108.

See ["To open a SQL port for a SQL 2008 instance and create an alias for a managed Backup Exec server"](#) on page 1109.

To change the dynamic port for a SQL Express instance to a static port and create an alias for the managed Backup Exec server

- 1 On the central administration server, click **Start > All Programs > Microsoft SQL Server 2008 R2 > Configuration Tools > SQL Server Configuration Manager**.
- 2 Expand **SQL Server Network Configuration**.
- 3 Click **Protocols for BKUPEXEC**, and then in the right pane, double-click **TPC/IP**.
- 4 On the **TCP/IP Properties** dialog box, click the **IP Addresses** tab.
- 5 Under **IPAll**, in **TCP Dynamic Ports**, remove the value and leave the field blank.
- 6 Under **IPAll**, type in a port number in **TCP Port**.

The port number can be between 1025 and 65535 and must not be in use by another application.
- 7 Under the heading for the specific network interface card that is being used, such as IP1 or IP2, change **Enabled** from **No** to **Yes**.
- 8 Under that same heading, in **TCP Dynamic Ports**, remove the value of 0, and type the same port number you entered for **TCP Port**.
- 9 Click **Apply**.
- 10 Restart the Backup Exec and SQL services.

- 11 On the managed Backup Exec server, click **Start > All Programs > Microsoft SQL Server 2008 R2 > Configuration Tools > SQL Server Configuration Manager**.
- 12 Expand **SQL Native Client Configuration**.
- 13 Click **Aliases**, and then double-click the alias name that contains the central administration server name and the Backup Exec SQL instance name.
- 14 On the alias properties dialog box, enter the appropriate information as described in the following table:

Alias Name	Type the name of the central administration server and the Backup Exec SQL instance name using the format server name\instance name.
Port No	Type the port number of the remote Backup Exec SQL Server instance that you noted in the previous procedure.
Protocol	Select TCP/IP .
Server	Type the name of the central administration server and the Backup Exec SQL instance name using the format server name\instance name.

- 15 Click **Apply**, and then click **OK**.
- 16 Close the SQL Server Configuration Manager utility.

To open a SQL port for a SQL 2008 instance and create an alias for a managed Backup Exec server

- 1 On the central administration server, click **Start > All Programs > Microsoft SQL Server 2008 R2 > Configuration Tools > SQL Server Configuration Manager**.
- 2 Expand **SQL Server Network Configuration**, and then click **Protocols** for the SQL Server instance that is in use for the central administration server.
- 3 Under **Enabled Protocols**, select **TCP/IP**, and then click **Properties**.
- 4 Double-click **TCP/IP**, and then click the **IP Addresses** tab.
- 5 Write down the TCP Dynamic port number.
- 6 On the managed Backup Exec server, to create an alias for the managed Backup Exec server, go to `\Windows\System32` and double-click **cliconfg.exe**.
- 7 On the **Alias** tab, click **Add**.
- 8 In the **Server alias** field, type: *server name\instance name*

- 9 Under **Network libraries**, select **TCP/IP**.
- 10 In the **Server name** field, type: *server name\instance name*
- 11 Uncheck **Dynamically determine port**.
- 12 In the **Port number** field, type the port number of the remote Backup Exec SQL Server instance.

Upgrading an existing CASO installation

In an existing CASO environment, upgrade the central administration server, and then upgrade the managed Backup Exec servers. You can upgrade from Backup Exec 2010 or Backup Exec 2012.

If necessary, you can perform rolling upgrades in the CASO environment. A rolling upgrade lets you upgrade the central administration server from the previous version to the current version first, and then upgrade the managed Backup Exec servers from the previous version to the current version over a period of time. Note that a rolling upgrade can be performed only from Backup Exec 2012. Rolling upgrades are not supported from any versions before Backup Exec 2012. You must have the most recent Backup Exec service pack to perform rolling upgrades.

Note: Forward compatibility is not supported in rolling upgrades. Therefore, any system that runs the previous version cannot protect a system that runs the current version.

Symantec recommends that you do not keep a mix of versions in the CASO installation for an extended time. Key functionality for administering managed Backup Exec servers is missing in a mixed-version environment, which decreases your ability to properly administer the CASO environment.

Note: Do not rename any managed Backup Exec servers or the central administration server during the upgrade process. Although you can rename the servers before you upgrade, Symantec recommends that you rename them after the upgrade process has completed.

After you upgrade the central administration server to the current version, the following operations are supported on managed Backup Exec servers that run the previous version:

- Backup
- Restore

- Inventory
- Catalog

To upgrade an existing CASO installation

- 1 Verify that the latest service pack for Backup Exec is installed.
- 2 Place all scheduled jobs on hold on the central administration server and the managed Backup Exec servers.
See [“Holding jobs”](#) on page 258.
- 3 Allow all active jobs to complete.
- 4 Run a database maintenance job to delete job histories and catalogs that you no longer need in order to shorten the upgrade window.
- 5 Stop all Backup Exec services on each managed Backup Exec server
- 6 From the installation media browser, select the option to install Symantec Backup Exec.
- 7 On the **Welcome** panel, click **Next**.
- 8 Select **I accept the terms of the license agreement**, and then click **Next**.
- 9 Check **Local Install**, and then click **Install Backup Exec software and options**.
- 10 Click **Next**.
- 11 Follow the prompts in the wizard.
- 12 On the **Back Up Existing Catalog and Data** page, enter or browse to a directory to which all existing catalogs and data will be backed up. The default location is:

`C:\Program Files\Symantec\Backup Exec\Data`

If you do not want to keep previous catalogs and data, click **Do not back up previous data and catalogs**.

- 13 Click **Next** to continue.

An upgrade summary is displayed. When the upgrade is complete, communication with the managed Backup Exec servers is automatically enabled.

- 14 Release the hold on all the jobs.

See [“Removing the hold on jobs”](#) on page 259.

- 15 Upgrade some or all of the managed Backup Exec servers.

Note: If you are upgrading the central administration server from Backup Exec 2010 R3 to Backup Exec 2014 and some of the managed Backup Exec servers were not backed up before the upgrade, those managed Backup Exec servers do not appear on the **Backup and Restore** tab until you upgrade the managed Backup Exec servers to Backup Exec 2014 or manually add them using the **Add Server** wizard.

See [“CASO distributed, centralized, and replicated catalog locations”](#) on page 1118.

See [“Changing the settings for a managed Backup Exec server”](#) on page 1120.

Changing a Backup Exec server to a central administration server

You can change a standalone Backup Exec server to a central administration server.

To change a Backup Exec server to a central administration server

- 1 On the Backup Exec server that you want to be the central administration server, start Backup Exec.
- 2 Click the Backup Exec button, select **Installation and Licensing**, and then select **Install Options and Licenses on this Backup Exec Server**.
- 3 Select one of the following methods to enter licenses:

To enter licenses manually	Do the following in the order listed:
----------------------------	---------------------------------------

- Type a serial number into the **Enter a serial number** field.
- Click **Add to list**.
- Repeat for each serial number for each option or agent that you want to add.

To import licenses from a file	Do the following in the order listed:
--------------------------------	---------------------------------------

- Click **Import Symantec License File**.
- Select the Symantec license file.

To install a trial version

Proceed to the next step.

A license key is not required for a fully functional Trial version.

- 4 Click **Next**.
- 5 On the features list, expand **Backup Exec Options**, expand **Enterprise Server Option**, and then select **Central Admin Server Option**.
- 6 Click **Next**.
- 7 Read the Backup Exec installation review, and then click **Install**.
- 8 Click **Finish**.

See [“Changing a Backup Exec server to a managed Backup Exec server”](#) on page 1113.

Changing a Backup Exec server to a managed Backup Exec server

To change a Backup Exec server to a managed Backup Exec server, you set the central administration server that will manage the Backup Exec server.

If the managed Backup Exec server does not appear on the **Storage** tab after you follow these instructions, and if your network contains firewalls, you may need to open some ports between the central administration server and the managed Backup Exec server.

To change a Backup Exec server to a managed Backup Exec server

- 1 At the standalone Backup Exec server, click the Backup Exec button, select **Installation and Licensing**, and then select **Install Options and Licenses on this Backup Exec Server**.
- 2 On the **Add Licenses** panel, click **Next**.
- 3 On the **Configure Options** panel, under **Backup Exec 2014**, select **Managed Backup Exec server**, and then click **Next**.
- 4 When you are prompted to select an additional language, click **Next**.
- 5 Enter the name of the central administration server that you want to manage this Backup Exec server.

- 6 Under **Managed Backup Exec Server Configuration**, select the appropriate option, and then click **Next**.

Centrally managed Backup Exec server Select this option to enable the central administration server to manage this Backup Exec server, its storage devices, media, and job delegation. This option also enables this Backup Exec server to share storage devices with other managed Backup Exec servers.

Unrestricted access to catalogs and backup sets for restore Select this option to enable this managed Backup Exec server to have unrestricted access to all centrally stored catalogs. This option also enables this managed Backup Exec server to restore data from any backup set on any storage devices that it shares.

This option can be selected only if the **Centrally managed Backup Exec server** option is selected. Selecting both of these options enables the central administration server to have the greatest amount of control over this managed Backup Exec server.

Locally managed Backup Exec server Select this option to enable the central administration server to monitor this managed Backup Exec server and create restore jobs for it. However, the server and its devices, media, and backup jobs are controlled locally.

- 7 Click **Install**.
- 8 After the installation completes, click **Finish**.

See [“Changing a Backup Exec server to a central administration server”](#) on page 1112.

Deleting a managed Backup Exec server from a CASO environment

If communications are still active between the managed Backup Exec server and the central administration server, you can remove a managed Backup Exec server by deleting it from the **Storage** tab. When you delete a managed Backup Exec

server from the **Storage** tab, the server is changed to a standalone Backup Exec server. After the managed Backup Exec server is removed from the CASO environment, it still remains on the list of servers on the **Backup and Restore** tab, so it can still be backed up and restored as a standalone server.

Note: If the central administration server experiences a disaster, communications are lost between the managed Backup Exec server and the central administration server and the managed Backup Exec server cannot be deleted from the **Storage** tab. However, you can change the managed Backup Exec server to a standalone Backup Exec server by using the Windows Change Program feature.

See [“Disaster recovery in CASO”](#) on page 1150.

To delete a managed Backup Exec server from a CASO environment

- 1 If the storage devices on the managed Backup Exec server that you want to delete are shared with other managed Backup Exec servers, on the **Storage** tab, right-click the device, and then click **Share**. Uncheck the servers that share the device.

You must remove the sharing reference to avoid overwriting the media.

Note: If any devices are shared by Fibre Channel SAN, disconnect the devices that are on the managed Backup Exec server before you delete it.

- 2 On the central administration server's **Storage** tab, right-click the managed Backup Exec server that you want to delete from the CASO environment and convert to a standalone Backup Exec server.
- 3 Select **Delete**.
- 4 Click **Yes** to confirm that you want to convert the server to a standalone Backup Exec server.
- 5 After you receive an alert on the central administration server that confirms the server was deleted, restart the server that was changed to a standalone Backup Exec server.

Note: If Backup Exec is unable to complete the request, you can choose to remove the reference to the managed Backup Exec server from the central administration server's database. Then, when the managed Backup Exec server is brought online again, Backup Exec automatically adds the managed Backup Exec server back to the central administration server's database. At that time, you can perform this procedure again.

See [“Changing a Backup Exec server to a managed Backup Exec server”](#) on page 1113.

Renaming a central administration server

If you back up and then rename a server, the new server name and the old server name both appear on the **Backup and Restore** tab. You should select the icon with the old server name to restore any data that you backed up before you changed the server name. You should select the icon with the new server name to restore any data that you backed up after you changed the server name.

After you rename a central administration server, that catalogs folder uses the new central administration server's name. An automatic process called catalog self-healing modifies the catalog files on the server and the catalog metadata in the Backup Exec Database.

Note: If you want to rename a central administration server that does not have any managed Backup Exec servers associated with it, you do not need to follow this procedure. Instead, you can use the standard Windows renaming process.

To rename a central administration server

- 1 On the central administration server's **Storage** tab, right-click a managed Backup Exec server, and then click **Delete**.

The managed Backup Exec server is not deleted from Backup Exec. It is removed only from the CASO environment, so it becomes a standalone Backup Exec server.
- 2 Repeat step 1 for every managed Backup Exec server that the central administration server manages.
- 3 Rename the server by using the standard Windows renaming process.
- 4 To rejoin the managed Backup Exec servers to the central administration server, install the managed Backup Exec server option on the servers that were deleted from the **Storage** tab in steps 1 and 2.

Note: Restart the server when the installation is complete.

- 5 Verify that the managed Backup Exec servers are rejoined to the renamed central administration server.

Renaming a managed Backup Exec server

Before you can rename a managed Backup Exec server, you must delete it from the CASO environment so that it becomes a standalone Backup Exec server. After you rename the server, you must reinstall the managed Backup Exec option to the server, and then redelegate all of the jobs to the new managed Backup Exec server name.

If you back up and then rename a server, the new server name and the old server name both appear on the **Backup and Restore** tab. You should select the icon with the old server name to restore any data that you backed up before you changed the server name. You should select the icon with the new server name to restore any data that you backed up after you changed the server name.

After you rename a managed Backup Exec server, the catalogs folder uses the new managed Backup Exec server's name. An automatic process called catalog self-healing modifies the catalog files on the server and the catalog metadata in the Backup Exec Database.

To rename a managed Backup Exec server

- 1 On the central administration server's **Storage** tab, right-click the managed Backup Exec server that you want to rename.
- 2 Click **Delete** to remove the managed Backup Exec server from the CASO environment.

The managed Backup Exec server is not deleted from Backup Exec. It is removed only from the CASO environment, so it becomes a standalone Backup Exec server. The server must be removed from the CASO environment before it can be renamed.

- 3 Restart the services on the server that you want to rename.
- 4 Rename the server by using the standard Windows renaming process.
- 5 Restart the server.

After the server is restarted, Backup Exec may prompt you to restart the Backup Exec services and the Backup Exec deduplication services. You must choose the new server name on the **Backup Exec Services Manager** dialog box.

- 6 To rejoin the server to the central administration server, install the managed Backup Exec server option on the renamed server.

Note: Restart the server when the installation is complete.

- 7 Verify that the renamed managed Backup Exec server is rejoined to the central administration server.
- 8 Delegate all jobs that were associated with the managed Backup Exec server prior to the renaming to the renamed managed Backup Exec server.

How to reduce network traffic in CASO

To accommodate a low-bandwidth network connection or to reduce network traffic, you can do the following:

- Reduce the frequency of job status updates that are sent from the managed Backup Exec servers to the central administration server.
- Prevent the central administration server from monitoring the jobs that are created on the local managed Backup Exec servers.
- Reduce the frequency that job logs and job histories are sent from the managed Backup Exec servers to the central administration server.
- Increase the amount of time that Backup Exec waits before changing the Backup Exec server's status if the Backup Exec server becomes unresponsive.
- Keep the catalogs on the managed Backup Exec server (distributed). If there is a persistent network connection between the central administration server and the managed Backup Exec server, then you can browse the catalog and perform restore operations from both servers, regardless of the catalog location.

See ["Changing the settings for a managed Backup Exec server"](#) on page 1120.

CASO distributed, centralized, and replicated catalog locations

In the CASO environment, you can choose the catalog location. Regardless of the catalog location, if a persistent network connection is available between the central administration server and the managed Backup Exec server, then you can browse the backup sets in the catalog and perform restore operations from both servers.

The following catalog locations are available:

Table L-2

Item	Description
Distributed	<p>Image files, which are small files that contain information about the backup set, are distributed to the central administration server from every managed Backup Exec server. History files, which contain detailed information about the backup set, remain on the managed Backup Exec server.</p> <p>Note: It is important that you back up the catalog files on the managed Backup Exec server since most catalog information is kept here when the distributed catalog location is used.</p> <p>When the catalog is distributed, the view of the restore selections on the central administration server displays only the backup set at the volume level. Backup set details are not displayed if the managed Backup Exec server that created this backup set is not available, but the whole volume can be restored from the central administration server.</p> <p>A distributed catalog provides increased performance, default centralized restore capability, and decreased network traffic. If a managed Backup Exec server does not have a persistent connection to the central administration server, then whenever the managed Backup Exec server does connect, the image files in the catalog are automatically distributed to the central administration server. The temporary increase in network traffic that is caused by the catalog distribution is not significant.</p>
Centralized	<p>All catalog files and information for the managed Backup Exec server are kept on the central administration server.</p>

Table L-2 (continued)

Item	Description
Replicated	<p>All catalog files are replicated from the managed Backup Exec server to the central administration server. Both the managed Backup Exec server and the central administration server store the catalogs that are produced by the managed Backup Exec server.</p> <p>Deletions of catalog files are replicated between the managed Backup Exec server and the central administration server only when the catalog files are deleted by Backup Exec according to the catalog settings. If catalog files on the managed Backup Exec server are deleted as a result of a backup job or a manual deletion, the deletions are replicated the next time that the catalogs are synchronized.</p>

When choosing the catalog location, consider the following:

- If there is enough available disk space on the managed Backup Exec server to keep a distributed or replicated catalog.
- If there is enough network bandwidth to handle the traffic that is generated by a centralized or replicated catalog. Centralized and replicated catalogs require a high-bandwidth network connection.
- If it is important for your data recovery needs to keep catalog information in one location. For example, when the catalog location is centralized or replicated, all catalog information is kept in one location, which makes it easier to back up. When the catalog location is distributed, most catalog information is kept on the managed Backup Exec server.

See [“Changing the settings for a managed Backup Exec server”](#) on page 1120.

Changing the settings for a managed Backup Exec server

The settings for a managed Backup Exec server determine how the managed Backup Exec server communicates and interacts with the central administration server. For example, you can change the connection type, the catalog location, and

the job reporting and monitoring functionality. You can change the settings for a managed Backup Exec server at any time.

Note: You may need to restart the services on the managed Backup Exec server after you change the settings. For example, if you change the catalog location, you must restart the services to enable the change to take effect.

To change the settings for a managed Backup Exec server

- 1 On the central administration server, on the **Storage** tab, double-click the managed Backup Exec server.
- 2 In the left pane, select **Settings**.
- 3 Select the appropriate options.

Connection settings

Select one of the following types of connections with the central administration server:

- **Fast connection**
Configures frequent communications between the central administration server and the managed Backup Exec server. By default, when you choose this setting, job status updates are sent every 10 seconds to the central administration server. Job logs and job histories are sent whenever a job on the managed Backup Exec server completes.
- **Slow connection**
Configures less frequent communications between the central administration server and the managed Backup Exec server. By default, when you choose this setting, job status updates are sent every 120 seconds to the central administration server. Job logs and job histories are sent only when a job on the managed Backup Exec server fails.
- **Custom**
Lets you change the thresholds that trigger the communication statuses when managed Backup Exec servers become unresponsive. You can also set how often the managed Backup Exec server sends active job status updates to the central administration server. The frequency affects the network traffic.

Communication stalled

Note: This option appears only if **Custom** is selected in the **Connection settings** field.

Indicate the amount of time to allow before the managed Backup Exec server's status changes to Communication Stalled if the managed Backup Exec server is unresponsive.

The central administration server does not delegate jobs to the managed Backup Exec server when it has a status of Communication Stalled. Job delegation resumes if the managed Backup Exec server returns to an Enabled status before the threshold is exceeded.

The default threshold is five minutes.

No communication

Note: This option appears only if **Custom** is selected in the **Connection settings** field.

Indicate the amount of time to allow before the managed Backup Exec server's status changes from Communication Stalled to No Communication.

When the status of the managed Backup Exec server changes from Communication Stalled to No Communication, the central administration server marks the active jobs on the managed Backup Exec server as Failed. The custom error-handling rule Recovered Jobs is applied to any job that is active when the No Communication status appears.

The default threshold is 15 minutes.

Send active job status updates to the central administration server

Note: This option appears only if **Custom** is selected in the **Connection settings** field.

Select this option to enable the managed Backup Exec server to send a job status update to the central administration server. You can adjust the number of seconds that a managed Backup Exec server waits between sending job status updates to the central administration server. To preserve network bandwidth when many jobs are running, increase the amount of time between job update statuses. Decrease the amount of time if you want to send more updates.

The default is 10 seconds, which provides near real-time monitoring. This setting is recommended only for fast network connections.

Send status updates to the central administration server every

Note: This option appears only if **Custom** is selected in the **Connection settings** field and **Yes** is selected in the **Send active job status updates to the central administration server** field.

Set the amount of time that a managed Backup Exec server waits between sending job status updates to the central administration server. To preserve network bandwidth when many jobs are running, increase the amount of time between job update statuses. Decrease the amount of time if you want to send more updates.

The default is 10 seconds, which provides near real-time monitoring. This setting is recommended only for fast network connections.

For low-bandwidth network connections, consider a setting of 120 seconds. This frequency allows updates to be displayed for a medium-sized job while still significantly decreasing the network traffic caused by job status updates.

If you uncheck the check box, job status updates are not sent. Job progress is not displayed on the central administration server. When the job is complete, the **Job History** on the central administration server is updated.

Send job log details to the central administration server

Choose when the job log for the managed Backup Exec server is sent to the central administration server. You can choose to send the job log one time per day, after a job completes, or never.

The following options are available:

- **Never**
If you select this option, job logs are stored locally at the managed Backup Exec server.
- **Once a day**
If you select this option, the **Send job logs at** field appears. You must select the time to send the job log to the central administration server.
- **On job completion**
If you select this option, the **Send job log only if the job fails** field appears. Select **Yes** to send the job log only for failed jobs. Select **No** to send the job log regardless of the job disposition.

Send job logs at

Choose the time when Backup Exec sends the job logs for the managed Backup Exec server to the central administration server. This option appears only if **Once a day** is selected in the option **Send job log details to the central administration server**

Send job log only if the job fails

Choose whether to send the job log for failed jobs only or for all jobs. Select **Yes** to send the job log only for failed jobs. Select **No** to send the job log regardless of the job disposition. This option appears only if **On job completion** is selected in the option **Send job log details to the central administration server**

Send job history details to central administration server

Choose when the job history for the managed Backup Exec server is sent to the central administration server.

The following options are available:

- **Never**
If you select this option, job histories are stored locally at the managed Backup Exec server.
- **Once a day**
If you select this option, the **Send job history logs at** field appears. You must select the time to send the job history to the central administration server.
- **On job completion**
If you select this option, the **Send job history only if the job fails** field appears. Select **Yes** to send the job history only for failed jobs. Select **No** to send the job history regardless of the job disposition.

Send job history at

Choose the time when Backup Exec sends the job history for the managed Backup Exec server to the central administration server. This option appears only if **Once a day** is selected in the option **Send job history details to the central administration server**.

Send job history details only if the job fails

Choose whether to send job history details for failed jobs only or for all jobs. Select **Yes** to send the job history details only for failed jobs. Select **No** to send the job history details regardless of the job disposition. This option appears only if **On job completion** is selected in the option **Send job history details to the central administration server**.

Monitor jobs that are created locally on the managed Backup Exec server	<p>Select this option if you want to view delegated jobs and the jobs that are created on the local managed Backup Exec server.</p> <p>You can also hold, delete, run, cancel, and change the priority order of the jobs that are created on or copied to the local managed Backup Exec server.</p>
Display an alert when the time is not synchronized between the managed Backup Exec server and the central administration server	<p>Select this option to enable Backup Exec to create an alert if the clock on the managed Backup Exec server differs from the clock on the central administration server. An alert is generated when the number of seconds indicated is exceeded.</p> <p>CASO monitors the internal computer clocks on both the managed Backup Exec servers and the central administration server. If time differences develop between the central administration server and the managed Backup Exec servers, jobs could run at unexpected times. To prevent problems, the time that is reported on managed Backup Exec servers should match the time that is reported on the central administration server. If you receive time difference alerts, reset the clock on the managed Backup Exec server to match the system clock on the central administration server.</p> <p>If you change the system time on either the managed Backup Exec server or the central administration server, you must restart the Backup Exec services on that server.</p>

Send the alert after the servers are not synchronized for	<p>Indicate the number of seconds that the clocks on the managed Backup Exec server and the central administration server must differ before Backup Exec sends an alert.</p> <p>Note: This option appears only if Enabled is selected in the Display an alert when the time is not synchronized between the managed Backup Exec server and the central administration server field.</p>
Storage and media database location	<p>Indicate whether the storage and media database is located on the central administration server or a managed Backup Exec server.</p>

Keep the catalogs on

Set the location of the catalog to one of the following locations:

- **Managed Backup Exec server (distributed)**
Distributes the catalog files between the central administration server and the managed Backup Exec server. If storage and media data is kept in a local database on the managed Backup Exec server, then the distributed location is the only available catalog location. Select this option if you have a low-bandwidth network connection.
- **Central administration server (centralized)**
Keeps all catalog files on the central administration server. A high-bandwidth network connection is required for this option.
- **Both servers (replicated)**
Replicates all catalog files from the managed Backup Exec server to the central administration server. If a managed Backup Exec server is unavailable, you can still browse the catalog from the central administration server. However, you cannot restore data because the managed Backup Exec server is unavailable. A high-bandwidth network connection is required for this option.

See [“CASO distributed, centralized, and replicated catalog locations”](#) on page 1118.

Private cloud server

Select this option to enable a managed service provider to locate a Backup Exec server in its data center, and then configure it for a CASO environment with other Backup Exec servers that are located across the WAN at the managed service provider's customer locations. As an alternative to shipping tapes off-site for storage, backups can be run and stored locally, and then copied to the cloud server's deduplication disk storage device. Additionally, this feature can be used by customers with widely distributed networks who want to use Backup Exec servers in remote offices for local backups, and then copy the backup sets to a Backup Exec server that is located in a central data center. This option is part of the Cloud Services for Backup Exec feature.

4 Click **Apply**.

What happens when CASO communication thresholds are reached

In a CASO environment, communications that occur between managed Backup Exec servers and the central administration server can sometimes be disrupted even if network communications are normal. If job-related communication disruptions occur between a managed Backup Exec server and the central administration server, the managed Backup Exec server's communication status changes from Enabled to Stalled or No Communication. The jobs waiting to be processed by the managed Backup Exec server are held in the managed Backup Exec server's job queue until the communications are restored.

You can set the amount of time that Backup Exec waits before changing the managed Backup Exec server's status if it becomes unresponsive. When a managed Backup Exec server's status changes to Stalled or No Communication, the central administration server changes how it handles current and future jobs that are delegated to the stalled managed Backup Exec server.

For example, if communications from a managed Backup Exec server are not received at the central administration server after the set amount of time, the central administration server changes the Backup Exec server's communication status to Stalled. Job delegation to the managed Backup Exec server is suspended as it

continues to wait for the managed Backup Exec server to return to an Enabled status. Jobs are delegated to other managed Backup Exec servers that are represented in the destination storage device or Backup Exec server pool.

CASO continues to monitor the amount of time during which no communications are received from the managed Backup Exec server. After a set amount of time passes after a Stalled status appears, CASO changes the status of the managed Backup Exec server to No Communication. CASO marks the jobs as Failed, and then begins job recovery by invoking the custom error-handling rule Recovered Jobs for any job that is active at the time the No Communication status appears.

See [“Changing the settings for a managed Backup Exec server”](#) on page 1120.

Enabling or disabling communications between the managed Backup Exec server and the central administration server

You can manually enable or disable communications between the managed Backup Exec server and the central administration server. When communications are disabled, jobs cannot be delegated to the managed Backup Exec server.

To enable communications between the managed Backup Exec server and the central administration server

- 1 From the central administration server's **Storage** tab, right-click the managed Backup Exec server for which you want to enable communications.
- 2 Select **Communication Enabled**.

To disable communications between the managed Backup Exec server and the central administration server

- 1 From the central administration server's **Storage** tab, right-click the managed Backup Exec server for which you want to disable communications.
- 2 Select **Communication Enabled** to remove the check mark.

Alerts and notifications in CASO

In a Central Admin Server Option (CASO) environment, the alerts that are generated on a managed Backup Exec server are automatically rolled up to the central administration server. To see those alerts on the central administration server, you must configure alert categories to enable or disable alerts on each managed Backup Exec server and on the central administration server itself.

After you respond to and clear the active alert on the central administration server, the alert is cleared on the managed Backup Exec server as well.

If you enable Backup Exec alerts on a managed Backup Exec server without enabling alerts on the central administration server, alerts appear only on the managed Backup Exec server where they are generated; they do not appear on the central administration server.

On the central administration server, you can view alerts for all managed Backup Exec servers, or you can filter the alerts to view only the alerts for a specific managed Backup Exec server or Backup Exec server pool.

You can configure a notification on either the central administration server or the managed Backup Exec server. Regardless of where you configure the notification, if it is for a delegated job, it is sent by the central administration server. You can choose to notify the local administrator of the managed Backup Exec server, or the administrator of the central administration server, or both.

Copying alert configurations to managed Backup Exec servers

You can enable and configure alerts at the central administration server, and then copy the alert configurations to a managed Backup Exec server. After the alert configurations are copied, the alerts that are generated on a managed Backup Exec server appear on both the managed Backup Exec server and the central administration server.

To copy alert configurations to managed Backup Exec servers

- 1 On the central administration server, click the Backup Exec button.
- 2 Select **Configuration and Settings**, and then select **Copy Settings to Other Servers**.
- 3 Under **Select settings to copy**, check **Alert configuration**.
- 4 Click **Add**.
- 5 Enter the name of a managed Backup Exec server to which you want to copy the alert configuration.
- 6 Click **OK**.
- 7 On the **Copy Settings** dialog box, click **OK**.

An alert on the central administration server confirms that the copy succeeded.

Enabling managed Backup Exec servers to use any available network interface card

By default, jobs that are delegated or copied to a managed Backup Exec server from the central administration server use the network and security settings that are set on the managed Backup Exec server.

However, you can select an option on the central administration server to let a job use any network interface to access Backup Exec agents if the selected network interface is unavailable. Enabling this option for a backup job lets the managed Backup Exec server use an alternative network interface to run important backup jobs that would otherwise fail to run.

To enable managed Backup Exec servers to use any available network interface card

- 1 On the central administration server, create a backup job.
- 2 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**.
- 3 In the left pane, select **Network**.
- 4 Check **Allow managed Backup Exec server to use any network interface to access Backup Exec agents**.
- 5 Set any additional options for the backup job
- 6 Click **OK**.

About job delegation in CASO

Job delegation is the automatic load balancing of jobs among the various storage devices that are attached to the managed Backup Exec servers. The job is created on the central administration server, but can be run on any managed Backup Exec server.

When the storage devices are logically grouped in Backup Exec server pools, and as the storage devices become available, they process jobs that are delegated from the central administration server. For example, if a storage pool contains two storage devices and one is busy processing a job, the central administration server automatically delegates another job to the idle storage device.

See [“How to use Backup Exec server pools in CASO ”](#) on page 1136.

About copying jobs instead of delegating jobs in CASO

If the managed Backup Exec server's storage and media data are kept on a local database on the managed Backup Exec server, the central administration server cannot delegate jobs to it. Instead, you can copy job options, default schedules, error-handling rules, and alert configurations from the central administration server to the managed Backup Exec server. A persistent network connection to the central administration server is not needed when the jobs are run locally on the managed Backup Exec server.

Use the same names for objects on the central administration server and all of the managed Backup Exec servers that you want to copy jobs to. For example, use the same name for a storage pool on the central administration server and on the managed Backup Exec server. Then, it is not necessary to customize settings or names for each managed Backup Exec server that you copy jobs to.

See [“Copying configuration settings to another Backup Exec server”](#) on page 626.

About adding storage devices in a CASO environment

From the central administration server, you can run the **Configure Storage Wizard** to set up devices for the central administration server or for any of the managed Backup Exec servers. After managed Backup Exec servers are installed, they appear on the **Storage** tab of the central administration server. When you start the **Configure Storage Wizard**, you are prompted to select the server for which you want to configure storage. You can choose the central administration server or any managed Backup Exec server that runs the same version of Backup Exec as the central administration server.

How data lifecycle management (DLM) works in a CASO environment

Backup Exec uses data lifecycle management (DLM) to delete the expired backup sets on disk storage, disk cartridge media, deduplication storage, storage arrays, and virtual disks. By default, Backup Exec keeps the most recent backup sets that are necessary to restore any backed-up component of a server, even if the backup sets expire. If backup sets are dependent of other backup sets, then Backup Exec does not delete the backup set until all expiration dates on the backup sets are reached. Even if the backup set is displayed as expired, the data is available until all dependent backup sets expire as well.

If you want Backup Exec to delete all expired backup sets, even if they are the last remaining backup sets that you need to restore a server, you can select the option **Allow Backup Exec to delete all expired backup sets** on the **Storage** settings dialog box. In a CASO environment, this option appears only on the central administration server. If you enable this option on the central administration server, DLM deletes all expired backup sets on the central administration server as well as on the managed Backup Exec servers. This option deletes all expired backup sets on both centrally managed and locally managed Backup Exec servers.

Warning: If you enable the option **Allow Backup Exec to delete all expired backup sets**, the data that you need to restore a server may not be available.

If you manually expire a backup set from the central administration server, DLM immediately runs on the server on which the backup set was created. The server can be either the central administration server or the managed Backup Exec server. DLM runs only on the storage device from which the backup set was manually expired. If you manually expire a backup set from a managed Backup Exec server, DLM runs immediately on the storage device from which the backup set was manually expired.

See [“How data lifecycle management \(DLM\) deletes expired backup sets on disk-based storage”](#) on page 335.

Obtaining media audit information for a managed Backup Exec server

The Media Audit report lists the recent configuration changes that were made to your media. In a CASO environment, if you run this report from the central administration server the report provides data only for the media for the central administration server. It does not provide any data for any of the managed Backup Exec servers. To obtain the media audit data for a managed Backup Exec server, you must either log on to the local managed Backup Exec server or access the managed Backup Exec server from the Remote Administration Console.

To obtain media audit information for a managed Backup Exec server

- 1 Do one of the following:
 - Log on locally to the managed Backup Exec server.
 - From a remote Windows server or workstation, click **Start**, point to **Symantec Backup Exec**, and then enter the name of the server you want to connect to and the credentials for that server.
- 2 On the **Reports** tab, under **Report Groups**, select **Configuration**.
- 3 Select **Media Audit** from the list of reports, and then click **Run Report Now** to run the report immediately or click **New Scheduled Report** to schedule the report to run later.

How to use Backup Exec server pools in CASO

In a CASO environment, you can group multiple managed Backup Exec servers together into Backup Exec server pools. If you create a pool of managed Backup Exec servers, all of the pools on those managed Backup Exec servers are available for job delegation. If there are multiple devices attached to each of the managed Backup Exec servers in the Backup Exec server pool, you can create multiple,


smaller pools that are made up of fewer storage devices. Use this method to send some jobs to a specific pool in the Backup Exec server pool, and send other jobs to a different pool in the same Backup Exec server pool.

Backup Exec server pools can contain multiple managed Backup Exec servers or just one managed Backup Exec server. A managed Backup Exec server can belong to more than one Backup Exec server pool. The central administration server can be used as a managed Backup Exec server and can be included in the Backup Exec server pool.

Any managed Backup Exec server or Backup Exec server in a pool must be able to access the destination device for the backup. If there is no intersection between the device and the managed Backup Exec server or Backup Exec server pools, the job does not run. The Jobs list displays the following status: Ready; No Backup Exec server available in Backup Exec server pool.

This graphic shows a Backup Exec server pool.

Figure L-3 An example of a CASO-configured Backup Exec server pool inside a corporate network



See [“Creating a Backup Exec server pool”](#) on page 1138.

See [“Selecting a Backup Exec server pool for backups”](#) on page 1137.

Selecting a Backup Exec server pool for backups

You can run a job on storage devices that are on a specific managed Backup Exec server or on storage devices that are in a group of managed Backup Exec servers. This filter lets you control where certain jobs are delegated. For example, to always run backups of Exchange databases only on the devices that are attached to

managed Backup Exec servers in a pool named Exchange Backups, you can select this option, and then select the Exchange Backups server pool.

To select a Backup Exec server pool for backups

- 1 Create a backup definition.
- 2 On the **Backup** box, click **Edit**.
- 3 On the **Backup Options** dialog box, in the left pane, select **Storage**.
- 4 In **Backup Exec server or Backup Exec server pool**, select the pool that you want to use for all of the backups in the backup definition.
- 5 Configure any additional options for the backup definition.

See [“How to use Backup Exec server pools in CASO ”](#) on page 1136.

Creating a Backup Exec server pool

You can group, or pool, Backup Exec servers.

See [“How to use Backup Exec server pools in CASO ”](#) on page 1136.

See [“Adding managed Backup Exec servers to a Backup Exec server pool ”](#) on page 1138.

To create a Backup Exec server pool

- 1 On the central administration server's **Storage** tab, in the **Configure** group, select **Configure Storage**.
- 2 Select the server on which you want to create the server pool, and then click **Next**.
- 3 Select **Storage pools**, and then click **Next**.
- 4 Select **Backup Exec server pool**, and then click **Next**.
- 5 Enter a name and description for the pool, and then click **Next**.
- 6 Check the check boxes next to the names of the servers that you want to include in the pool, and then click **Next**.
- 7 Click **Finish**.

Adding managed Backup Exec servers to a Backup Exec server pool

You can add managed Backup Exec servers to existing Backup Exec server pools.

See [“Creating a Backup Exec server pool”](#) on page 1138.

To add managed Backup Exec servers to a Backup Exec server pool

- 1 On the central administration server, select the **Storage** tab.
- 2 Expand **All Storage Pools**, and then double-click a Backup Exec server pool to which you want to add managed Backup Exec servers.
- 3 Under **Backup Exec servers that belong to the pool**, click **Add**.
- 4 Check the check boxes for the Backup Exec servers that you want to add to the pool, and then click **OK**.
- 5 Click **Apply**.

Deleting a Backup Exec server pool

You can delete a Backup Exec server pool at any time.

To delete a Backup Exec server pool

- 1 On the central administration server, select the **Storage** tab.
- 2 Expand **All Storage Pools**.
- 3 Right-click the Backup Exec server pool that you want to delete, and then click **Delete**.
- 4 Click **Yes** to confirm that you want to delete the pool.

See [“Removing a managed Backup Exec server from a Backup Exec server pool”](#) on page 1139.

Removing a managed Backup Exec server from a Backup Exec server pool

Removing a managed Backup Exec server deletes it from a Backup Exec server pool, but does not remove it from Backup Exec.

To remove a managed Backup Exec server from a Backup Exec server pool

- 1 On the central administration server, select the **Storage** tab.
- 2 Expand **All Storage Pools**, and then double-click the Backup Exec server pool that contains the server you want to delete.
- 3 Under **Backup Exec servers that belong to the pool**, select the Backup Exec servers that you want to remove from the pool, and then click **Remove**.
- 4 Click **Apply**.

See [“How to use Backup Exec server pools in CASO ”](#) on page 1136.

How centralized restore works in CASO

Depending on whether the required storage media resides in storage devices or is stored off-site, initiating restore operations from the central administration server can be an automated process with little user intervention necessary.

When you use centralized restore with online media, you run the Restore Wizard at the central administration server. During the data selection process, CASO determines which media are required to complete the restore operation, and then queries the Backup Exec storage and media database to determine the identity of the storage device where the media reside. After you run the Restore Wizard, CASO begins the restore operation by delegating the jobs to the central administration server or managed Backup Exec servers that control the selected storage devices. If the data that is being restored spans multiple storage media, you are prompted to load additional media as needed to successfully complete the restore operation.

When you use centralized restore with offline media, you run the Restore Wizards at the central administration server. During the data selection process, CASO determines what media is required to complete the restore operation, and then queries the Backup Exec storage and media database to determine the identity of the storage device where the primary media resides. If the media is not found in a storage device, the media is considered offline. CASO then presents you with a selection of drive pools and storage devices that are compatible with the type of media being used during the restore operation, thus giving you the flexibility of choosing a storage device in which to load your media.

After noting the identity and location of the storage device you have selected to run the job, you do the following:

- Submit the restore job on hold as a scheduled job.
- Retrieve the media, place it in the storage device.
- Remove the job from hold at the central administration server, at which time the restore job begins.

CASO then delegates the job to the managed Backup Exec server that controls the selected storage device. If the data being restored spans multiple storage media, you are prompted to load additional media as needed to successfully complete the restore operation.

Before restore operations from the central administration server can be initiated, the following requirements must be met:

- Managed Backup Exec server communication status must be Enabled.
- Managed Backup Exec servers must be online with all Backup Exec server statuses showing Online.

Best practices for centralized restore in CASO

Symantec recommends the following best practices for centralized restore:

- Select only one resource to restore for each job.
- Select the same restore device or Backup Exec server for all of the selections that are in the same restore job.
- Select a Backup Exec server that has compatible devices for all media that is required for the restore job.

How CASO restores data that resides on multiple storage devices

If the data selected for restore is located on a single device that is attached to a managed Backup Exec server, then a single restore job is created at, and then delegated from, the central administration server. However, if the data being selected for restore is located on multiple devices in the CASO environment, then the single restore job is split into separate restore jobs, depending on the number of devices involved.


All split restore jobs have the same name as the original job, but are differentiated and linked with a subscript numeral that is appended to the job name.

For example, if you create a restore job and the data you select for restore resides in one device on a managed Backup Exec server, CASO creates one restore job. However, if you create one restore job and the data you select resides on two or more devices that are attached to a managed Backup Exec server, CASO creates two or more restore jobs.

The following graphic shows how CASO restores data that is stored on a single device.

Figure L-4 For data stored on a single storage device

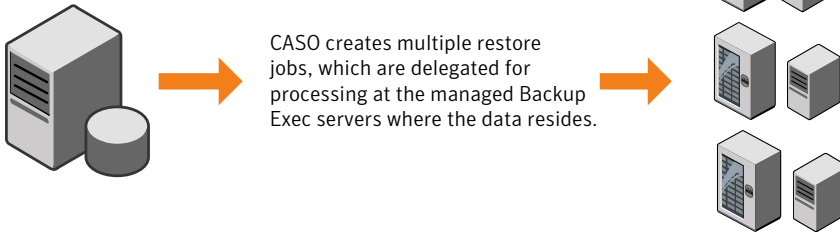
You select the data to restore from the administration console of the central administration server.



The following graphic shows how CASO restores data that is stored on multiple devices.

Figure L-5 For data stored on multiple storage devices

You select the data to restore from the administration console of the central administration server.



About recovering failed jobs in CASO

The Backup Exec error-handling rule named Recovered Jobs is a custom error-handling rule that is used by CASO to recover the jobs that failed because of issues with internal job communications. This rule is created when Backup Exec is installed and is enabled by default.

The retry options for this rule are to retry the job twice, with an interval of five minutes between the retry attempts. During the first retry attempt, CASO attempts to re-delegate the jobs to another available managed Backup Exec server.

If this attempt fails, CASO makes a second attempt at finding another available managed Backup Exec server to process the jobs. If no managed Backup Exec servers are available, the final job disposition is to place the job on hold until you have fixed the error condition.

Note: If the checkpoint restart error-handling rule is enabled, then recovered jobs are not resubmitted to a Backup Exec server pool to be run on a different server. The checkpoint restart error-handling rule reschedules the job to run on the original server when that server comes online. To enable a recovered job to be resubmitted to a Backup Exec server pool, you must disable the checkpoint restart error-handling rule.

Note: If you target a job to a Backup Exec server pool that contains multiple managed Backup Exec servers and a job failure occurs, the recovery process uses only the managed Backup Exec servers in the Backup Exec server pool. Managed Backup Exec servers that are not in the Backup Exec server pool are not used for job recovery.

When you open the job history entry for a recovered job, the reason for the failure is listed as Job Errors, with an explanation of the type of internal communication error that occurred. The job history entry also indicates that the job was recovered.

Note: Job logs are not created for the jobs that are recovered.

The following table describes the CASO error codes that are selected by default for the Recovered Jobs custom error-handling rule:

Table L-3 Error codes for Recovered Jobs custom error-handling rule

Error code	Description
0xE000881B JOBDISPATCH	The displayed message is: Job failed while being dispatched. The job will be recovered.
0xE000881D JOB_CASO_QUEUE FAILURE	The displayed message is: The job could not be delegated to the destination managed Backup Exec server. The managed Backup Exec server may not be online, or there may be a communications failure. The job will be recovered.
0xE000881E JOB_CASO_REMOTEMMS_STARTFAILURE	The displayed message is: The job failed to start on the destination managed Backup Exec server, possibly because a database error occurred. The job will be recovered.

See [“Error-handling rules for failed or canceled jobs”](#) on page 272.

See [“Custom error-handling rule for recovered jobs”](#) on page 275.

See [“Changing the settings for a managed Backup Exec server”](#) on page 1120.

Pausing or resuming a managed Backup Exec server

You can pause and resume a managed Backup Exec server from the central administration server.

Pausing a managed Backup Exec server prevents the central administration server from delegating jobs to it. When the managed Backup Exec server is paused, its status changes from Online to Paused.

Caution: When you install Backup Exec options at a managed Backup Exec server, the managed Backup Exec server must be paused so that no further jobs are delegated to it from the central administration server while the installation process occurs. If jobs are running, let them finish or cancel them before beginning the installation.

When you resume a paused managed Backup Exec server, the following changes occur:

- Jobs can be delegated from the central administration server to the managed Backup Exec server.
- The managed Backup Exec server's status changes from Paused to Online in the **State** column on the **Storage** tab.

To pause a managed Backup Exec server

- 1 From the central administration server's **Storage** tab, right-click the managed Backup Exec server that you want to pause.
- 2 Click **Pause**.

To resume a paused managed Backup Exec server

- 1 From the central administration server's **Storage** tab, right-click the managed Backup Exec server that you want to resume.
- 2 Click **Pause** to remove the check mark next to **Pause**.

Stopping or starting Backup Exec services for a managed Backup Exec server

You can stop or start the Backup Exec services on a managed Backup Exec server from the central administration server.

To stop Backup Exec services for a managed Backup Exec server

- 1 On the central administration server's **Storage** tab, right-click the managed Backup Exec server for which you want to stop services.
- 2 Select **Backup Exec Services**.
- 3 On the **Backup Exec Services Manager** dialog box, click **Stop all services**.
- 4 Click **Close**.

To start Backup Exec services for a managed Backup Exec server

- 1 On the central administration server's **Storage** tab, right-click the managed Backup Exec server for which you want to start services.
- 2 Select **Backup Exec Services**.
- 3 On the **Backup Exec Services Manager** dialog box, click **Start all services**.
- 4 Click **Close**.

Viewing managed Backup Exec server properties

From the central administration server, you can view properties for managed Backup Exec servers.

To view managed Backup Exec server properties

- 1 On the central administration server's **Storage** tab, double-click the managed Backup Exec server for which you want to view properties.

2 In the left pane, select **Properties**.

Name	Displays the name of the managed Backup Exec server or the central administration server.
Description	Indicates whether the server is a managed Backup Exec server or a central administration server. You can change this description.
Server status	Indicates the current status of the server, such as online, paused, unavailable, or offline.
Version	Indicates the version of Backup Exec that is installed.
License	Provides information about the Backup Exec license that is installed on the server.
Time zone	Indicates the time zone that is set for the server.
Start date and time	Indicates when the server was started.
Current date and time	Indicates the current date and time on the server.
Operating system type	Indicates the type of operating system that is installed on the server.
Operating system version	Indicates the version of the operating system that is installed on the server.
Operating system build	Indicates the build number of the operating system that is installed on the server.
Processor type	Indicates the type of processor that the server has.
Number of processors	Indicates the number of processors that the server has.
Total physical memory	Indicates the total amount of physical memory that the server has.
Available physical memory	Indicates the amount of physical memory that is available on the server.

Total virtual memory	Indicates the total amount of virtual memory that the server has.
Available virtual memory	Indicates the amount of virtual memory that is available on the server.
Total pagefile size	Indicate the total amount of memory that is available in the server's page file.

Viewing the settings for a central administration server

If you have the Central Admin Server Option (CASO), you can view information about the location of the databases for Backup Exec. The databases include the Backup Exec Database, the Advanced Device and Media Database (ADAMM), and the catalog database.

During Backup Exec installation, if you chose the default option to create a local Backup Exec SQL Express instance on which to store the Backup Exec Database, the databases are all located on the local Backup Exec server. If you chose another instance on the network on which to store the Backup Exec Database, then the databases are all located on the Microsoft SQL Server that contains that instance.

To view the settings for a central administration server

- 1 Do one of the following:
 - Click the Backup Exec button, select **Configuration and Settings**, and then click **Local server properties**.

- On the **Storage** tab, double-click the central administration server.
- 2 In the left pane, click **Settings**.

Server	Shows the name of the Microsoft SQL Server that contains the Backup Exec Database.
Instance	Shows the name of the instance that the Backup Exec Database is installed on.
Name	Shows the Microsoft SQL Server name of the Backup Exec Database.
Path	Shows the path of the Backup Exec Database.
Server	Shows the name of the Microsoft SQL Server that contains the ADAMM database.
Instance	Shows the name of the instance that the ADAMM database is installed on.
Name	Shows the Microsoft SQL Server name for the ADAMM database.
Path	Shows the path of the ADAMM database.
Server	Shows the name of the Microsoft SQL Server that contains the Backup Exec catalog database.
Instance	Shows the database instance that contains the catalog database.
Name	Shows the Microsoft SQL Server name for the Backup Exec catalog database.
Path	Shows the path of the Backup Exec catalog database.
Private cloud server	Indicates if the private cloud server option is enabled or disabled.

Disaster recovery in CASO

Use the Symantec Backup Exec Simplified Disaster Recovery (SDR) option to protect both managed Backup Exec servers and the central administration server in a CASO environment.

See [“About Simplified Disaster Recovery”](#) on page 736.

Before implementing the SDR option in a CASO environment, review the following:

- To create recovery media for any managed Backup Exec server or central administration server, the **Create Simplified Disaster Recovery Disk Wizard** must be run at the central administration server. If you use a remote administration environment, connect to the central administration server.
- If you want managed Backup Exec servers to be protected using a bootable disk image, you must run the **Create Simplified Disaster Recovery Disk Wizard** at each of the managed Backup Exec servers where a bootable disk device is installed.
- You must locally back up and restore a central administration server.

Disaster recovery of a managed Backup Exec server that is managed by a failed central administration server

If communications are still active between the managed Backup Exec server and the central administration server, you can change a managed Backup Exec server to a standalone Backup Exec server by deleting it from the **Storage** tab. However, if the central administration server experiences a disaster, communications are lost between the managed Backup Exec server and the central administration server and the managed Backup Exec server cannot be deleted from the **Storage** tab. However, you can change the managed Backup Exec server to a standalone Backup Exec server by using the Windows Change Program feature.

To change a managed Backup Exec server to a standalone Backup Exec server when communications between the two servers have been lost

- 1 On the managed Backup Exec server, open the Windows control panel.
- 2 Select **Add/Remove Programs** or **Programs and Features**, depending on the version of Windows you are using.
- 3 Select **Backup Exec** from the list, and then click **Change**.
- 4 On the **Additional Options** panel, click **Next**.
- 5 On the **Add Licenses** panel, click **Next**.
- 6 On the **Configure Options** panel, click **Next**.
- 7 On the **Choose Languages** panel, click **Next**.

- 8 On the **Central Admin Server** panel, select **Locally managed Backup Exec server**, and then click **Next**.
- 9 When the error message about the managed Backup Exec server being unable to contact the central administration server appears, click **OK**.

Note: If an error message does not appear, then Backup Exec detected a server with the same name as the central administration server on the network. If that server is the central administration server and it is running correctly, you should be able to change it to a managed Backup Exec server from the **Storage** tab on the central administration server.

- 10 Click **Next**.
- 11 On the **Installation Review** panel, click **Install**.
- 12 When the installation is complete, on the Backup Exec server, add any storage devices that may be needed.

Note: You may also need to inventory and catalog the storage device, depending on which catalog method was in use when the managed Backup Exec server was joined to the central administration server.

When the central administration server is online again, you can change this standalone Backup Exec server back to a managed Backup Exec server.

See [“Changing a Backup Exec server to a managed Backup Exec server”](#) on page 1113.

Troubleshooting CASO

If you encounter issues with CASO, review the following questions and answers.

Table L-4 Troubleshooting CASO

Question	Answer
I received error 1065 that says "Database specified does not exist". What causes this error?	<p>This error can occur for the following reasons:</p> <ul style="list-style-type: none">■ UDP traffic is blocked on the network between the central administration server and the managed Backup Exec server.■ The SQL configuration on the central administration server is not set correctly.■ When the central administration server is installed to a named SQL instance and the SQL browser service is not running.■ The Named Pipes or TCP/IP protocols are not enabled or are not set for remote connections.
I changed the system time, but the change hasn't gone into effect on my managed Backup Exec servers or central administration server. Why?	<p>If you change the system time on either the managed Backup Exec server or the central administration server, you must restart the Backup Exec services. Backup Exec processes the time change when the services restart.</p>
I received the error "Ready, Job storage does not contain any eligible devices". What causes this error?	<p>This error can occur if you remove managed Backup Exec servers from the wrong location. To remove managed Backup Exec servers from a central administration server, use the Delete option on the Storage tab. When you remove a managed Backup Exec server from the Storage tab, the server becomes a standalone server, so it can still be backed up and restored as a standalone server, but it can no longer be controlled by the central administration server. Do not remove the managed Backup Exec server from the Backup and Restore tab.</p> <p>See "Deleting a managed Backup Exec server from a CASO environment" on page 1114.</p>

See ["About the Central Admin Server Option"](#) on page 1096.

Running the Backup Exec Utility for CASO operations

A separate application called Backup Exec Utility is available to help you perform the following CASO operations:

- Move a managed Backup Exec server.
- Disable or enable communication with a managed Backup Exec server.

Use the Backup Exec Utility only with the guidance of Symantec Technical Support. Improper use of this utility can result in configuration changes that may prevent Backup Exec from running.

To run the Backup Exec Utility

- 1 From the Backup Exec installation directory, in `\Program Files\Symantec\Backup Exec`, double-click **BEUtility**.
- 2 On the **Backup Exec Utility** menu, click **Help** to learn about how to use BE Utility.

Uninstalling Backup Exec from the central administration server

Before you uninstall Backup Exec from the central administration server, you must delete all managed Backup Exec servers from the **Storage** tab on the central administration server.

Caution: Failure to uninstall in the following sequence may result in long delays when shutting down Backup Exec services during the uninstall of Backup Exec on the managed Backup Exec servers.

To uninstall Backup Exec from the central administration server

- 1 On the central administration server's **Storage** tab, right-click a managed Backup Exec server.
- 2 Select **Delete**.
- 3 Click **Yes** to confirm that you want to delete.
- 4 Repeat steps 1 through 3 for each managed Backup Exec server that the central administration server manages.
- 5 Uninstall Backup Exec from the central administration server.

See [“Uninstalling Backup Exec”](#) on page 130.

Uninstalling Backup Exec from a managed Backup Exec server

You must delete the managed Backup Exec server from the **Storage** tab on the central administration server before you uninstall Backup Exec.

To uninstall Backup Exec from a managed Backup Exec server

- 1 On the central administration server's **Storage** tab, right-click a managed Backup Exec server.
- 2 Select **Delete**.
- 3 Click **Yes** to confirm that you want to delete.
- 4 Uninstall Backup Exec from the managed Backup Exec server.

See [“Uninstalling Backup Exec”](#) on page 130.

Symantec Backup Exec Advanced Disk-based Backup Option

This appendix includes the following topics:

- [About the Advanced Disk-based Backup Option](#)
- [How to use synthetic backups in place of recurring full backups](#)
- [Configuring off-host backup options for a backup job](#)
- [Setting default backup options for the Advanced Disk-based Backup Option](#)
- [About true image restore for synthetic backups](#)
- [How to use off-host backup to process remote computer backups on the Backup Exec server](#)
- [Troubleshooting off-host backups](#)

About the Advanced Disk-based Backup Option

The Advanced Disk-based Backup Option provides the following features:

- Synthetic backup

This feature enables a full backup to be assembled, or synthesized, from a baseline full backup and subsequent incremental backups.

The benefits of using a synthetic backup include the following:

- A reduced backup window, since the synthetic backup can be scheduled outside of the time-critical backup window.

- Reduced network traffic, since the synthetic backup does not need to access the network.
- True image restore, which enables Backup Exec to restore the contents of directories to what they were at the time of any synthetic full or incremental backup.
- Off-host backup
This feature enables the backup operation to be processed on a Backup Exec server instead of on the remote computer or host computer. Moving the backup from the remote computer to the Backup Exec server enables better backup performance and frees the remote computer as well.

The Advanced Disk-based Backup Option (ADBO) is installed as part of the Enterprise Server Option, for which you must enter a license on the Backup Exec server.

See [“How to use synthetic backups in place of recurring full backups”](#) on page 1156.

See [“How to use off-host backup to process remote computer backups on the Backup Exec server”](#) on page 1163.

See [“About the Backup Exec installation process and licensing options”](#) on page 54.

See [“Setting default backup options for the Advanced Disk-based Backup Option”](#) on page 1159.

How to use synthetic backups in place of recurring full backups

The synthetic backup feature eliminates the need to perform recurring full backups for supported remote resources. The synthetic backup is assembled from a full backup (called a baseline) and subsequent incremental backups.

The resulting synthetic backup then becomes the new baseline. Only incremental backups are required until the next synthetic backup is created. The synthetic backup is as current as the last incremental backup that it contains.

The components of a synthetic backup are as follows:

- Baseline backup.
The baseline backup is the first full backup to run that is associated with the synthetic backup. The full baseline backup runs one time only, and backs up all of the files on the selected computer when it runs.
- Recurring incremental backups.
Incremental backup jobs back up the files that change after the baseline backup.

- **Recurring synthetic backups.**

The synthetic backup process combines the data from the baseline backup and the incremental backups to form a synthesized full backup of the selected computer. This synthesized full backup becomes a new baseline backup, which is combined with subsequent incremental backup sets to form a new synthesized full backup.

For any of the backups in a synthetic backup, you can add a stage to duplicate the backup data to tape.

True image restore is automatically enabled for synthetic backups. True image restore lets you restore directories as they existed at the time of the synthetic backup. Files that were deleted before the time of the synthetic backup are not restored. In true image restore, only the correct versions of files are restored from the appropriate synthetic full or incremental backups that contain them.

Only file system data is supported for synthetic backup. Supported data includes common file system objects, such as volumes, drives, and folders.

Requirements for synthetic backup

Before you create a synthetic backup, review the following information:

- If you use an encryption key, all of the associated backups must use the same encryption key. Do not change the encryption key after the backups are created. The encryption key that is selected in the associated backups is automatically applied to the synthetic backup.
- You must configure disk storage before you can create a synthetic backup. For synthetic backups, incremental backups must use disk storage. The baseline full backup and the synthetic full backup can use tape or disk storage. See [“Configuring disk storage”](#) on page 318.
- If you send the baseline backup job to tape storage, and you want to use tape storage for the synthetic backup job, you must have two tape drives. You must use a tape drive to mount the baseline backup, and use a tape drive to mount the synthetic backup.
- You can select only file system data for synthetic backup.

See [“Configuring off-host backup options for a backup job”](#) on page 1157.

See [“About true image restore for synthetic backups”](#) on page 1161.

Configuring off-host backup options for a backup job

Off-host backup enables Backup Exec to move backup processing from the host computer, which is the remote computer that contains the volumes selected for

backup, to the Backup Exec media server. The off-host backup creates a snapshot of the volume or volumes that are selected for backup on the remote computer. The snapshots are then imported to the media server, where they are backed up.

To configure off-host backup options for a backup job

- 1 On the **Backup and Restore** tab, right-click the server, servers, or server group that you want to back up.
- 2 Select **Backup**, and then select the type of backup that you want to perform.
- 3 On the **Backup Definition Properties** dialog box, in the **Selections** box, click **Edit**.
- 4 On the **Backup Selections** dialog box, check the boxes for the resources that you want to back up and uncheck the check boxes for the resources that you do not want to back up.
- 5 Click **OK**.
- 6 On the **Backup Definition Properties** dialog box, in the **Backup** box, click **Edit**.
- 7 On the **Backup Options** dialog box, select the schedule for this job.
- 8 On the **Backup Options** dialog box, in the left pane, select **Advanced Disk-based Backup**.
- 9 Set any of the following options for this job:

Use offhost backup to move backup processing from the remote computer to the Backup Exec server

Select this option to enable off-host backup.

Off-host backup enables Backup Exec to move backup processing from the host computer to the Backup Exec server. The off-host backup creates a snapshot of the volume or volumes that are selected for backup on the remote computer. The snapshots are then imported to the Backup Exec server, where they are backed up.

Continue the backup job (offhost backup is not used)

Select this option to let the backup job complete without using the off-host feature if either of the following conditions occur:

- The selected volumes do not support off-host backup.
- An error occurs that is related to the snapshot import or the volume import.

Fail the backup job (further selections are not backed up after failure occurs)

Select this option to fail the backup job if either of the following conditions occur:

- The selected volumes do not support off-host backup.
- An error occurs that is related to the snapshot import or the volume import.

Process logical volumes for offhost backup one at a time

Select this option to enable the backup of multiple volumes in one job, while a snapshot is created of only one logical volume at a time. To ensure database integrity, or if a volume contains mount points, multiple volumes may need to be snapped at one time.

After the logical volume is snapped and backed up, the snapshot is deleted before the next logical volume is snapped. This option increases the ability to meet the minimum quiet time that is needed to complete a snapshot.

A logical volume can comprise multiple physical volumes. A single logical volume can encompass all of the volumes on which databases reside.

10 On the **Backup Options** dialog box, click any of the optional settings in the left pane that you want to set for this job.

11 Click **OK**.

12 On the **Backup Definition Properties** dialog box, click **OK**.

See [“About the Advanced Disk-based Backup Option”](#) on page 1155.

See [“How to use off-host backup to process remote computer backups on the Backup Exec server”](#) on page 1163.

Setting default backup options for the Advanced Disk-based Backup Option

You can use the defaults that Backup Exec sets during installation for all Advanced Disk-based Backup Option backup jobs, or you can choose your own defaults. You can override the default settings when you create individual jobs.

To set default backup options for Advanced Disk-based Backup Option

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Job Defaults**, and then select a backup option.

For example, if you want to set up the default options for Advanced Disk-based Backup Option backups to disk, select **Back Up to Disk**. The options that appear vary depending on what types of storage devices you have configured. Different default options can be configured for backup jobs to different types of storage.

- 3 In the left pane, select **Advanced Disk-based Backup**.
- 4 Select the appropriate options:

Use offhost backup to move backup processing from the remote computer to the Backup Exec server

Select this option to enable off-host backup.

Off-host backup enables Backup Exec to move backup processing from the host computer to the Backup Exec server. The off-host backup creates a snapshot of the volume or volumes that are selected for backup on the remote computer. The snapshots are then imported to the Backup Exec server, where they are backed up.

Continue the backup job (offhost backup is not used)

Select this option to let the backup job complete without using the off-host feature if either of the following conditions occur:

- The selected volumes do not support off-host backup.
- An error occurs that is related to the snapshot import or the volume import.

Fail the backup job (further selections are not backed up after failure occurs)

Select this option to fail the backup job if either of the following conditions occur:

- The selected volumes do not support off-host backup.
- An error occurs that is related to the snapshot import or the volume import.

Process logical volumes for offhost backup one at a time

Select this option to enable the backup of multiple volumes in one job, while a snapshot is created of only one logical volume at a time. To ensure database integrity, or if a volume contains mount points, multiple volumes may need to be snapped at one time.

After the logical volume is snapped and backed up, the snapshot is deleted before the next logical volume is snapped. This option increases the ability to meet the minimum quiet time that is needed to complete a snapshot.

A logical volume can comprise multiple physical volumes. A single logical volume can encompass all of the volumes on which databases reside.

See [“About the Advanced Disk-based Backup Option”](#) on page 1155.

See [“Backing up data”](#) on page 158.

About true image restore for synthetic backups

True image restore is automatically enabled for synthetic backups. True image restore enables Backup Exec to restore the contents of directories to what they were at the time of any full backup or incremental backup. Restore selections in backup sets are made from a view of the directories as they existed at the time of the synthetic backup. Files that were deleted before the time of the backup are not restored. In true image restore, only the correct versions of files are restored from the appropriate full or incremental backups that contain them. Previous versions are not restored and then overwritten.

Note: The **Use storage-based catalogs** option must be enabled in the **Catalogs** settings to use true image restore.

See [“Configuring default options for catalogs”](#) on page 249.

Backup Exec collects the information that is required to detect the files and directories that have been moved, renamed, or newly installed from a tape archive (tar) or a compressed archive. Depending on how the files were packaged and how they were installed, some newly installed files are not backed up by normal incremental backups. With true image restore enabled, Backup Exec compares path names with path names from the previous full or incremental backup. If a name is new or changed, the file or directory is backed up.

The following are examples where using true image restore backs up the files that would not otherwise be backed up:

- A file named C:\pub\doc is moved to or installed in C:\spec\doc. Here, the archive bit is unchanged for files and subdirectories inside that directory, but C:\pub\doc is new in the C:\spec\ directory and is backed up.
- A directory named C:\security\dev\ is renamed as C:\security\devices\. Here, the archive bit is unchanged for files and subdirectories inside that directory, but C:\security\devices\ is a new directory and is backed up.

The following table lists the files that are backed up in the C:\user\doc directory during a series of backups between December 1, 2012 and December 4, 2012:

Table M-1 Example table of files backed up because true image restore is enabled

Day	Type of backup	Backed up files in C:\user\doc	Backed up files in C:\user\doc	Backed up files in C:\user\doc	Backed up files in C:\user\doc	Backed up files in C:\user\doc	Backed up files in C:\user\doc
December 1, 2012	Full	file1	file2	dirA\fileA	dirB\fileB	file3	
December 2, 2012	Incremental	file1	file2	dirA\fileA	-----	-----	
December 3, 2012	Incremental	file1	file2	dirA\fileA	-----	-----	
December 4, 2012	Incremental	file1	file2	-----	-----	-----	file4

Note: Dashes (-----) indicate that the file was deleted before this backup.

Assume that you want to restore the December 4, 2012 version of the C:\user\doc directory.

You perform a regular restore of the full backup set followed by a regular restore of subsequent incremental backup sets. The restored directory contains all files and directories that ever existed in C:\user\doc from December 1, 2012 (last full backup) through December 4, 2012.

For example, the following files and directories are included:

- file1
- file2
- dirA\fileA
- dirB\fileB

- file3
- file4

In a true image restore of the December 4, 2012 backup, the restored directory has only the files and directories that existed at the time of the incremental backup on December 4, 2012.

The following list includes the files and directories that existed:

- file1
- file2
- file4

Backup Exec does not restore any of the files that were deleted before the December 4, 2012 incremental backup.

The restored directory does not include the 'dirA' subdirectories, even though they were backed up on December 4, 2012. Backup Exec does not restore these directories because they did not exist at the time of the incremental backup, which was the reference for the true image restore.

A true image restore preserves the files that are currently in the directory but were not present when the backup was completed. Assume that you created a file named file5 after the incremental backup that occurred on December 4, 2012, but before doing the restore.

In this case, the directory contains the following files after the restore:

- file1
- file2
- file4
- file5

See [“How to use synthetic backups in place of recurring full backups”](#) on page 1156.

See [“Configuring off-host backup options for a backup job”](#) on page 1157.

How to use off-host backup to process remote computer backups on the Backup Exec server

Off-host backup enables Backup Exec to move backup processing from the host computer to the Backup Exec server. The off-host backup creates a snapshot of the volume or volumes that are selected for backup on the remote computer. The snapshots are then imported to the Backup Exec server, where they are backed up.

After the backup, the snapshots are exported from the Backup Exec server and mounted back on the remote computer and resynchronized with the source volume. This process requires solutions from the hardware providers that can support transportable snapshots. Transportable snapshots are snapshots you can import to and export from the Backup Exec server. The Microsoft Volume Shadow Copy Services (VSS) provider that you select is used for each volume in the off-host backup. An off-host backup job is performed on one remote computer at a time.

You can find a list of hardware snapshot providers at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

Off-host backup supports the following:

- Microsoft Volume Shadow Copy Service (VSS).
- Backups for the NTFS volumes that use the full, incremental, and differential backup methods.
- SQL Agent backups for Microsoft SQL Server 2000 and later databases.
- Exchange Agent backups for Microsoft Exchange Server 2007/2010 instances that run on Windows Server 2003. Support for the option to use Backup Exec Granular Recovery Technology for Exchange Agent backups is included.

Requirements for off-host backup

Additionally, the following are requirements for off-host backup:

Table M-2 Off-host backup requirements

Item	Description
Backup Exec server	The Advanced Disk-based Backup Option must be installed.
Remote computer	The Agent for Windows must be installed on the remote computer.
Backup Exec server and the remote computer	<p>The following must be installed on both the Backup Exec server and on the remote computer:</p> <ul style="list-style-type: none"> ■ The same operating system; Microsoft Windows Server 2003 SP2/2008/2012. ■ The most recent Volume Shadow Copy Services (VSS) patches. ■ The Microsoft VSS hardware snapshot provider that you want to use. Otherwise, the snapshots of the volumes cannot be exported to the Backup Exec server. ■ Ability to access the SAN shared storage or logical disk that is shared between the Backup Exec server and the remote computer.

Table M-2 Off-host backup requirements (*continued*)

Item	Description
GRT-enabled off-host backup of Exchange Server resources	<p>Either of the following must be installed on the Exchange Server:</p> <ul style="list-style-type: none"> ■ Microsoft Exchange Server 2007 (Service Pack 3) instances that run on Windows Server 2003. ■ Microsoft Exchange Server 2010 (Service Pack 1) instances that run on 64-bit Windows Server 2008/2008 R2.
Central Admin Server Option	<p>If the Central Admin Server Option (CASO) is installed, do not let the central administration server delegate the job. It can delegate the job to a managed Backup Exec server that does not have the off-host capability. You must manually select the storage device for the CASO jobs that use the off-host backup method.</p>

Advanced Disk-based Option off-host backup does not support the following:

- The option **Checkpoint Restart**.
See [“Configuring checkpoint restart”](#) on page 543.
- Volumes that run Windows BitLocker Drive Encryption.
- The backup method for files **Using catalogs**.
See [“How Backup Exec determines if a file has been backed up”](#) on page 199.

Best practices for off-host backup

The following best practices are recommended:

- Do not allow source volumes and snapped volumes to share the same physical disks. Otherwise, any attempt to split the snapshot volume from the original volume fails.
- Most hardware and software providers have some limitation about the types of volumes that are transportable. Therefore, Symantec recommends that you use off-host backup jobs only for backing up data for which all dependent volumes can be imported and exported.
- Ensure that the volume that you select to back up can be imported and exported and that the VSS hardware provider is on the Symantec-approved compatibility list. Otherwise, the off-host backup fails. You can choose to continue the backup if the off-host backup fails.

You can find a list of compatible types of storage at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

- The Hitachi Raid Manager log cannot be on a volume that is being snapped. Hitachi executes I/O to its Raid Manager log file during the snapshot commit process, and the VSS coordinator blocks I/O to any drive being snapped. Therefore, if the log directory for Raid Manager is on the volume that is being snapped, then log I/O is blocked and the snap process is deadlocked.
- If the Central Admin Server Option (CASO) is installed, you must manually select the storage for the off-host backup. Otherwise, the job may be delegated to a Backup Exec server that does not have off-host capability.
See [“How to use Backup Exec server pools in CASO ”](#) on page 1136.
- When you run an off-host backup that uses a VSS hardware provider in a Microsoft Cluster (MSCS) environment, the Backup Exec server and the remote computer must not be in the same cluster group. The cluster applications cannot support the devices’ logical unit numbers (LUNs) that have duplicate signatures and partition layouts. The snapshots containing the LUNs must be transported to a host computer that is outside the cluster.

See [“Backing up data”](#) on page 158.

See [“Configuring off-host backup options for a backup job”](#) on page 1157.

See [“Setting default backup options for the Advanced Disk-based Backup Option”](#) on page 1159.

See [“Troubleshooting off-host backups”](#) on page 1166.

Troubleshooting off-host backups

Off-host backup requires that the VSS providers and the volumes that are to be transported are set up correctly. Not all arrays are supported with the Advanced Disk-based Option.

You can find a list of compatible types of storage at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

To troubleshoot off-host backup issues, Symantec recommends that you use the tools that are available from the VSS provider to verify the required setup for off-host backup.

The minimum setup requirements are as follows:

- Volumes that you want to back up are snappable.
- Volumes are shared between the remote computer and the Backup Exec server.
- An off-host backup job can only contain the volumes that can be transported to the Backup Exec server for backup.

Other factors to consider are as follows:

- Microsoft Windows Server 2003 with Service Pack 2 or Windows Server 2008 with Service Pack 2 must be installed on both the Backup Exec server and on the remote computer. Both computers must have the most recent Volume Shadow Copy Services (VSS) patches.
- Microsoft XML Core Services (MSXML4) must be installed and running on both the Backup Exec server and the remote computer.

Troubleshooting off-host backup issues depends to an extent on the VSS provider that is used for the snapshots.

The following setup issues that are common to all providers may cause off-host backup to fail:

Table M-3 Common setup issues for off-host backup

Issue	Solution
The volumes are not shared.	You must ensure that all of the volumes reside on the disks that are shared between the remote computer and the Backup Exec server. If the volumes are not shared, the import operation fails. You may need to clean up the snapshots and resynchronize the volumes manually.
The VSS provider is not installed on the Backup Exec server and the remote computer.	The provider that is used for the snapshot must be installed on both the Backup Exec server and on the remote computer. If the provider is not installed on the Backup Exec server, the import operation fails. You may need to clean up the snapshots and resynchronize the volumes manually.
All volumes are not transportable.	All of the volumes that are selected for backup must be transportable to the Backup Exec server. If Microsoft SQL, Microsoft Exchange, or other database applications are selected for backup, make sure that the databases and log files reside on transportable volumes.
The VSS provider cannot snap all of the selected volumes.	All of the volumes that are selected for backup must be transportable to the Backup Exec server. All volumes that you select for backup must be snappable by the same provider. You must ensure that the same VSS provider supports all of the volumes in a backup job.
The log path location is incorrect.	If the provider or the supporting application creates log files during normal snapshot operation, the log files should not reside on any of the volumes that are being snapped. VSS cannot flush the write buffers, and the snapshot times out. Change the log path to another volume.

Table M-3 Common setup issues for off-host backup *(continued)*

Issue	Solution
The provider or VSS services are not started.	The provider service should be running and the Microsoft Windows "Volume Shadow Copy" service should not be disabled.
The credentials are incorrect.	The machine-level credentials that are used for the job should be the same on both the Backup Exec server and the remote computer. Incorrect credentials cause snapshots or the backup to fail.
The VSS provider is not installed on all Backup Exec servers in a Central Admin Server Option (CASO) environment.	If you configure a backup job in a CASO environment, you must send the job to managed Backup Exec servers on which the selected VSS provider is installed. You should not let the central administration server delegate the job. Otherwise, the job may be delegated to a managed Backup Exec server that does not have off-host capability.
The Backup Exec server and the remote computer are in the same cluster group.	<p>For an off-host backup in a Microsoft Cluster environment, the Backup Exec server and the remote computer must not be in the same cluster group. The cluster applications cannot support the devices' logical unit numbers (LUNs) that have duplicate signatures and partition layouts. Therefore, you must transport the snapshots that contain the LUNs to a Backup Exec server that is outside the cluster in which the host cluster resides.</p> <p>See "How Backup Exec works in a Microsoft Cluster Server environment" on page 711.</p> <p>If you use a Hitachi 9970 and attempt to back up Microsoft Cluster data by using the Advanced Disk-based Backup Option, you may receive the following error message:</p> <p>The job failed with the following error: A failure occurred querying the Writer status.</p> <p>To correct this problem, ensure that the RM Shadow Copy Provider for Volume Snapshot Service is present and running. If the service is not running, run RMVSSPRV.exe from c:\horcm\tool. If the service is still not running, contact Hitachi for support.</p>

Off-host backup issues with hardware providers

Hardware disk array vendors may support VSS snapshots and the transporting of volumes to the Backup Exec server for backup in a SAN environment. Using hardware providers requires a sound understanding of how disk arrays are

configured for shared access between the remote computer and the Backup Exec server in a SAN.

Consult the documentation for your hardware disk array on how to set up such disk arrays for off-host backup . Specifically, note any limitations on using the disk arrays in context with VSS snapshots, and note how to verify that the volumes are transportable. Symantec recommends that you make use of any tools that vendors provide to help verify the setup and to troubleshoot issues.

An off-host backup issue that can occur when you use Hitachi hardware may be that Hitachi supports only basic disks for off-host backup. The off-host backup feature is not supported if a computer uses a combination of dynamic and basic disks, and the Hitachi provider.

See [“About the Advanced Disk-based Backup Option”](#) on page 1155.

See [“How to use off-host backup to process remote computer backups on the Backup Exec server”](#) on page 1163.

Symantec Backup Exec NDMP Option

This appendix includes the following topics:

- [Features of the NDMP Option](#)
- [Requirements for using the NDMP Option](#)
- [About installing the NDMP Option](#)
- [Adding NDMP servers to Backup Exec](#)
- [Sharing the tape drives on NDMP servers between multiple Backup Exec servers](#)
- [NDMP backup options for NDMP servers](#)
- [NDMP server backup selections](#)
- [How to use patterns to exclude files and directories from an NDMP server backup](#)
- [Supported configurations for duplicating data from NDMP servers](#)
- [About restoring and redirecting restore data for NDMP servers](#)
- [Setting the default backup options for the NDMP Option](#)
- [Viewing the properties of an NDMP server](#)
- [Viewing storage properties for an NDMP server](#)

Features of the NDMP Option

The Symantec Backup Exec NDMP Option uses the Network Data Management Protocol (NDMP) to back up and restore network-attached storage NDMP servers.

You can use the following configurations to back up data from a NDMP server:

- A direct-attached configuration in which a tape device or a virtual tape library is directly connected to the NDMP server.
- A three-way configuration in which a tape device or a virtual tape library is connected to another NDMP server.
- A remote configuration in which a storage device is attached to the Backup Exec server.
- A configuration in which any storage device is attached to the same SAN as the Backup Exec server.

Note: You cannot back up data from an NDMP server to a simulated tape library or to a tape device that is attached to a Backup Exec Remote Media Agent for Linux.

You cannot redirect backup sets from an NDMP server to a computer that runs the Windows or Linux operating systems.

See [“Requirements for using the NDMP Option”](#) on page 1171.

See [“Sharing the tape drives on NDMP servers between multiple Backup Exec servers”](#) on page 1174.

Requirements for using the NDMP Option

To use the NDMP Option, the following items are required:

- Backup Exec must be installed on a server running Windows Server 2003/Server 2008/Server 2008 R2 operating systems.
- The network-attached storage NDMP server must run version 4 of the Network Data Management Protocol.

You can find a list of compatible types of storage at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

See [“About installing the NDMP Option”](#) on page 1172.

See [“Adding NDMP servers to Backup Exec”](#) on page 1172.

About installing the NDMP Option

The NDMP Option is installed locally on the Backup Exec server as a separate add-on component of Backup Exec. No files are copied to the network-attached storage NDMP server.

See [“Installing additional Backup Exec options to the local Backup Exec server”](#) on page 76.

See [“Requirements for using the NDMP Option”](#) on page 1171.

Adding NDMP servers to Backup Exec

You can add a network-attached storage NDMP server to Backup Exec to back it up and to use the storage devices that are attached to it.

If you add an NDMP server in a Central Admin Server Option environment, you must add it to one of the following:

- The central administration server.
- A managed Backup Exec server on which the storage and media database is located.

Table N-1 Methods to add an NDMP server to Backup Exec

Task	Method
To add an NDMP server for backup	Use the Add Server wizard and add the NDMP server as a file server or NDMP data server. The NDMP server is added to the list of servers on the Backup and Restore tab. See “About the list of servers on the Backup and Restore tab” on page 152.
To add an NDMP server for backup and use its attached storage devices	Use the Configure Storage wizard to add the NDMP server to the list of servers on the Backup and Restore tab. Then, after the Backup Exec services are restarted, Backup Exec automatically discovers any storage devices that are attached to the NDMP server. The storage devices appear on the Storage tab. See “To add the NDMP server for backup and use its attached storage devices” on page 1173.

To add the NDMP server for backup and use its attached storage devices

1 On the **Storage** tab, in the **Configure** group, click **Configure Storage**.

2 Do one of the following:

If the Central Admin Server Option is not installed

Click **Network storage**, and then click **Next**.

If the Central Admin Server Option is installed

Do the following in the order listed:

- Select the Backup Exec server on which you want to configure storage, and then click **Next**.
- Click **Network storage**, and then click **Next**.

3 Click **NDMP Storage**, and then click **Next**.

4 Enter a host name or fully qualified domain name of the NDMP server and a description, and then click **Next**.

5 Specify information for the following fields:

Port number

Specify the port number to be used for communication between the Backup Exec server and the NDMP server.

Logon account

Select the name of the logon account for the NDMP server. You can add a new logon account or edit an existing account.

Have Backup Exec use ICMP ping operations to detect the server

Select this option to enable Backup Exec to use ping to locate the NDMP server.

6 Click **Next**.

7 Review the summary, and then do one of the following:

To change the configuration

Do the following in the order listed:

- Click the heading that contains the items that you want to change.
- Make any changes, and then click **Next** until the summary appears.
- Click **Finish**.

To configure the NDMP server

Click **Finish**.

- 8 Click the **Backup and Restore** tab, and view the NDMP server in the list of servers.

Sharing the tape drives on NDMP servers between multiple Backup Exec servers

To share the tape drives that are attached to a network-attached storage NDMP server, the NDMP Option and the Central Admin Server Option (CASO) must be installed. CASO is installed as part of the Enterprise Server Option (ESO).

You can then select which Backup Exec servers can share the tape drives. The Backup Exec server that you used to add the NDMP server is automatically selected for sharing.

Note: If you upgraded from an earlier version of Backup Exec, your existing configuration is preserved. You do not have to set up sharing for existing configurations.

To share a tape drive on NDMP servers between multiple Backup Exec servers

- 1 On the **Storage** tab, right-click the storage device that you want to share.
- 2 Click **Share**.
- 3 Check the Backup Exec servers or managed Backup Exec servers that you want to share this storage device.
- 4 Click **OK**.

See [“About the Central Admin Server Option”](#) on page 1096.

NDMP backup options for NDMP servers

When you create a backup job for a network-attached storage NDMP server, you can use the options that are appropriate for the job and the type of NDMP server.

The following backup options are available for NetApp, IBM, and Fujitsu NDMP servers:

Table N-2 NDMP backup options for NetApp/IBM/Fujitsu

Item	Description
Back up Access Control Lists	Backs up NetApp Access Control Lists.
Enable file history	<p>Enables the generation of file history data. File history is used to optimize recovery of selected backup sets, however, file history generation and processing increase the backup time. Disabling this option improves backup time.</p> <p>If the file history is not generated, and you must restore data later, restore the entire volume.</p>
Backup method	Specifies the backup level. Level 0 provides a full backup. Level 1 provides a differential backup that backs up new or modified files since level 0. Levels 2 through 9 back up new or modified files since the previous level backup. For example, the level 2 backup method backs up new or modified files since the level 1 backup. Level 3 backup backs up new or modified files since the level 2 backup, and so on.

The following backup options are available for EMC NDMP servers

Table N-3 NDMP backup options for EMC

Item	Description
Backup type	<p>Determines the backup type for this backup job.</p> <p>The following backup types are available:</p> <ul style="list-style-type: none">■ Dump■ VBB
Back up with integrated checkpoints (SnapSure)	Enables Backup Exec to create a backup set that uses the EMC SnapSure feature. For more information about SnapSure, see your EMC documentation.

Table N-3 NDMP backup options for EMC (continued)

Item	Description
Enable file history	Enables the generation of file history data. File history is used to optimize the recovery of selected backup sets. File history generation and processing increase the backup time. Disabling this option improves backup time. If file history is made unavailable and you must restore data later, restore the entire volume.
Backup method	Specifies the backup level. Level 0 provides a full backup. Level 1 provides a differential backup that backs up new or modified files since level 0. Levels 2 through 9 back up new or modified files since the previous level backup. For example, the level 2 backup method backs up new or modified files since the level 1 backup. Level 3 backup backs up new or modified files since the level 2 backup, and so on.

For NDMP servers other than NetApp, IBM, Fujitsu, and EMC, Backup Exec displays the appropriate options. You can change the values for most options, such as the backup level and file history option. You can change some options to a predefined value, some to any value, and some options you cannot change.

Symantec has tested the options that appear for each NDMP server. The NDMP server may support additional options but Symantec does not support them at this time. Backup Exec does not validate the values that you enter for the options, so ensure that you enter the correct values. See the documentation for your NDMP server for information about the options that appear.

NDMP server backup selections

When you create a backup, you can select a network-attached storage NDMP server. On the NDMP server, you can select volumes or directories, or you can specify the data to include or exclude.

The following limitations apply when you select volumes or directories on NDMP servers for backup:

- You can include entire volumes for any NDMP server.
- You can include subfolders for NetApp/IBM/Fujitus NDMP servers only.

- You cannot exclude files or directories.

See [“Backing up data”](#) on page 158.

If you want to specify what to include or exclude in the backup job, you can do the following:

- Select specific directories to include in the backup job.
- Type the names of the specific directories and files that you want to exclude from the backup job in the **Include/Exclude** dialog box.

Table N-4 What you can include and exclude for NDMP server backups

Type of NDMP server	Include	Exclude
NetApp/IBM/Fujitsu	Single or multiple directories	Directories and files, based on an exclusion pattern.
EMC	Single directory	Directories and files, but only if you select the Dump backup type.
Other	Single or multiple directories, depending on the NDMP server.	Directories and files. Refer to your NDMP server documentation for details.

See [“How to use patterns to exclude files and directories from an NDMP server backup”](#) on page 1177.

How to use patterns to exclude files and directories from an NDMP server backup

When you exclude files and directories from a backup for a NetApp/IBM/Fujitsu or EMC Celerra network-attached storage NDMP server, you must use patterns. You should enter patterns carefully to ensure that you exclude the correct files and directories. Backup Exec does not verify the validity of exclude patterns. If you enter an invalid pattern, the pattern is ignored and therefore the files or directories are not excluded.

For details about how to use patterns, see your NDMP server documentation.

The following example shows a pattern to exclude files and directories from a backup selection for a NetApp/IBM/Fujitsu NDMP server:

Table N-5 Example pattern for NetApp/IBM/Fujitsu NDMP servers

Pattern	Example
tmp	Excludes all files and directories that have the name "tmp".
*.core	Excludes all files and directories that end with ".core".

To exclude directories for an EMC Celerra NDMP server, do not include the name of the EMC Celerra server or the name of the file system in the pattern. The names of the NDMP server and the file system are already included in the **Resource name** text box on the **Include/Exclude** dialog box. If you repeat the name of the NDMP server and the file system in the pattern, the EMC Celerra NDMP server ignores the exclusion. Type the path from the root directory to the directory that you want to exclude. Do not include an initial forward slash (/).

The following example shows a pattern to exclude directories from a backup selection for an EMC Celerra NDMP server:

Table N-6 Example pattern to exclude directories for an EMC Celerra NDMP server

Pattern	Description
test_exclusion/subdir1	Excludes only the "subdir1" directory on the file system that is listed in the Resource name text box.
test_exclusion/d*	Excludes all directories that start with the letter "d" under the directory "/test_exclusion".

The following example shows a pattern to exclude files from a backup selection for an EMC Celerra NDMP server:

Table N-7 Example pattern to exclude files for an EMC Celerra NDMP server

Pattern	Description
*.mp3	Excludes all files that end with ".mp3".
temp	Excludes all files that have the name "temp".

See ["NDMP server backup selections "](#) on page 1176.

See ["About selecting data to back up "](#) on page 169.

Supported configurations for duplicating data from NDMP servers

With the NDMP Option, you can duplicate the backup data from a network-attached storage NDMP server to storage that is attached to a Backup Exec server or to another NDMP server.

Backup Exec supports the following configurations for duplicating backup data from NDMP servers:

- Two storage devices that are attached locally to the Backup Exec server.
- Two tape drives that are attached locally to an NDMP server.
- One tape drive that is attached locally to an NDMP server and one tape drive that is attached locally to another NDMP server.
- A storage device that is attached locally to a Backup Exec server and one tape drive that is attached locally to an NDMP server.

The procedure to duplicate backup data from NDMP servers is the same as the procedure to duplicate any other type of data. For NetApp/IBM/Fujitsu NDMP servers, you must also select the logon credential for the source NDMP server.

See [“Duplicating backup sets or a job history manually”](#) on page 221.

About restoring and redirecting restore data for NDMP servers

With the NDMP Option, you can use the Restore Wizard on the **Backup and Restore** tab to restore data for a network-attached storage NDMP server. During the restore process, you can select individual files for restore if file history was enabled for the backup job. You cannot exclude files and directories when you restore to an NDMP server. Excluded directories and files are restored.

Backup Exec cannot gather sufficient file and directory information on an NDMP server restore job to accurately populate the **Backup Set Summary** and **Backup Set Information** sections of the job history. Therefore, the number of files, directories, files skipped, corrupt files, and files in use always appears as 0.

Note: Ensure that the default catalog option **Use storage-based catalogs** is selected. Otherwise, NDMP server backup sets cannot be cataloged.

About redirecting restore data for NDMP servers

You can use the NDMP Option to redirect restore data from one network-attached storage NDMP server to another NDMP server.

When you redirect NDMP server data, be aware of the following limitations:

- You cannot redirect NDMP server data to a computer that runs the Windows or Linux operating systems.
- You cannot redirect non-NDMP server data, such as NTFS or SQL data, to an NDMP server.
- The NDMP server to which you want to redirect the restored data must be from the same vendor as the NDMP server from which the data was backed up.

Backup Exec cannot gather sufficient file and directory information on an NDMP server backup to accurately populate the **Job Summary Information** and the **Set Detail Information** sections of the job history. Therefore, for restore and verify operations, the number of files, directories, skipped files, corrupt files, and files in use always appears as 0. Job summary and set detail information display for backup and duplicate operations.

Note: Verify operations are only supported for NetApp servers.

See [“Configuring default options for catalogs”](#) on page 249.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

See [“NDMP server restore options”](#) on page 1180.

NDMP server restore options

With the NDMP Option, you can create a restore job for a network-attached storage NDMP server. The options that appear in the Restore Wizard vary depending on the type of NDMP server.

See [“About restoring and redirecting restore data for NDMP servers”](#) on page 1179.

Table N-8 NDMP Option restore options for NetApp/IBM/Fujitsu

Item	Description
Restore Access Control Lists	Restores NetApp Access Control Lists.

Table N-8 NDMP Option restore options for NetApp/IBM/Fujitsu *(continued)*

Item	Description
Enable Direct Access Recovery	<p>Enables Backup Exec to use Direct Access Recovery (DAR) during the restore job. With DAR-enabled recovery, Backup Exec can specify the exact location of a file in a backed up data stream. The NDMP server can then read the data that applies to the single file that you want to restore. This practice reduces the amount of information that is processed and significantly reduces recovery time.</p> <p>If you do not select this option, the restore may take significantly longer.</p> <p>Note: Not all vendors provide Direct Access Recovery.</p>
Restore without writing data to disk (Verify data without doing a restore)	<p>Tests the validity of the data that you selected for the restore job. Backup Exec does not restore the data.</p> <p>For NetApp/IBM NDMP servers, you should use this option to verify data instead of the Verify option on the backup definition..</p>
Recreate the directory structure from the backup when the data is restored; otherwise, all data is restored without any directory structure	<p>Restores the data with its original directory structure intact.</p>

Table N-9 NDMP Option restore options for EMC

Item	Description
Enable Direct Access Recovery	<p>Enables Backup Exec to use Direct Access Recovery (DAR) during the restore job. With DAR-enabled recovery, Backup Exec can specify the exact location of a file in a backed up data stream. The NDMP server can then read the data that applies to the single file that you want to restore. This practice reduces the amount of information that is processed and significantly reduces recovery time.</p> <p>If you do not select this option, the restore may take significantly longer.</p> <p>Note: Not all vendors provide Direct Access Recovery.</p>
Recreate the directory structure from the backup when the data is restored; otherwise, all data is restored without any directory structure	<p>Restores the data with its original directory structure intact.</p>
Restore over existing files	<p>Overwrites the files on the restore destination that have the same name as files that are restored. Use this option only when you are sure that you want to restore an older version of a file.</p>

For NDMP servers other than NetApp, IBM, Fujitsu, and EMC, Backup Exec displays the appropriate variables and default values for the specific type of NDMP server. You can change the values as needed. Variables that begin with the prefix "@@" are specific to Backup Exec rather than to a specific NDMP server. Symantec has tested the variables that appear for each NDMP server. The NDMP server may support additional variables but Symantec does not support them at this time. Backup Exec does not validate the values for the variables that you enter, so you should ensure that you enter the values correctly. See the documentation for your NDMP server for information about the values to use.

Setting the default backup options for the NDMP Option

You can use the defaults for all NDMP server backup jobs that Backup Exec sets for the NDMP Option during installation, or you can choose your own defaults. You can also change the defaults for any specific backup job.

To set default backup options for the NDMP Option

- 1 Click the Backup Exec button, and then select **Configuration and Settings**.
- 2 Select **Job Defaults**, and then select a backup option.
- 3 In the left pane, select **NDMP**.
- 4 Do either of the following:
 - Select the appropriate options for NetApp, IBM, and Fujitsu NDMP servers:

Back up Access Control Lists

Backs up NetApp Access Control Lists.

Enable file history

Enables the generation of file history data. File history is used to optimize recovery of selected backup sets, however, file history generation and processing increase the backup time. Disabling this option improves backup time.

If the file history is not generated, and you must restore data later, restore the entire volume.

Backup method

Specifies the backup level. Level 0 provides a full backup. Levels 1 through 9 provide various levels of incremental backups. Level 1 backup method backs up new or modified files since the level 0 backup. Level 2 backup method backs up new or modified files since the level 1 backup, and so on.

- Select the appropriate options for EMC NDMP servers

Backup type

Determines the backup type for this backup job.

The following backup types are available:

- Dump
- VBB

Back up with integrated checkpoints (SnapSure)	Enables Backup Exec to create a backup set that uses the EMC SnapSure feature. For more information about SnapSure, see your EMC documentation.
Enable file history	Enables the generation of file history data. File history is used to optimize the recovery of selected backup sets. File history generation and processing increase the backup time. Disabling this option improves backup time. If file history is made unavailable and you must restore data later, restore the entire volume.
Backup method	Specifies the backup level. Level 0 provides a full backup. Levels 1 through 9 provide various levels of incremental backups. Level 1 backup method backs up new or modified files since the level 0 backup. Level 2 backup method backs up new or modified files since the level 1 backup, and so on.

- 5 Click **Apply**.

Viewing the properties of an NDMP server

You can view the properties of a network-attached storage NDMP server that you back up.

To view the properties of an NDMP server

- 1 On the **Backup and Restore** tab, double-click the NDMP server.
- 2 In the left pane, select **Properties**.

See [“NDMP server properties”](#) on page 1184.

NDMP server properties

You can view the following properties for a network-attached storage NDMP server:

See [“Viewing the properties of an NDMP server”](#) on page 1184.

Table N-10 NDMP server properties

Item	Description
Name	Indicates the name of the NDMP server.
Description	Shows the user-defined description of the NDMP server.
Logon account	Indicates the name of the logon account for the NDMP server. You can add a new logon account or edit an existing account.

Viewing storage properties for an NDMP server

You can view the storage properties of a network-attached storage NDMP server.

To view storage properties for an NDMP server

- 1 On the **Storage** tab, double-click the NDMP server.
- 2 In the left pane, select **Properties**.

See [“Storage properties for an NDMP server”](#) on page 1185.

Storage properties for an NDMP server

You can view the following storage properties for a network-attached storage NDMP server.

See [“Viewing storage properties for an NDMP server”](#) on page 1185.

Table N-11 NDMP server storage properties

Item	Description
Server name	Indicates the name of the NDMP server.
Description	Shows the user-defined description of the server.
State	Indicates the status of the NDMP server storage. See “Backup Exec server and storage device states” on page 475.
Port	Lists the port that is used for communications between the Backup Exec server and the NDMP server.

Table N-11 NDMP server storage properties (*continued*)

Item	Description
Use ICMP ping operations to detect the server	Indicates whether ICMP ping is enabled. ICMP ping enables Backup Exec to use ping to locate the NDMP server.
Logon account	Indicates the name of the logon account for the NDMP server. You can add a new logon account or edit an existing account.
Host ID	Displays the identifier number that the NDMP server generates.
System version	Indicates the software version that is installed on the NDMP server.

Symantec Backup Exec Agent for Linux

This appendix includes the following topics:

- [About the Agent for Linux](#)
- [About open files and the Agent for Linux](#)
- [Requirements for the Agent for Linux](#)
- [About installing the Agent for Linux](#)
- [About establishing trust for a remote Linux computer in the Backup Exec list of servers](#)
- [Adding additional Backup Exec servers to which the Agent for Linux can publish information](#)
- [About configuring the Agent for Linux](#)
- [About excluding files and directories from backup jobs for Linux computers](#)
- [Editing configuration options for Linux computers](#)
- [About backing up a Linux computer by using the Agent for Linux](#)
- [About restoring data to Linux computers](#)
- [Edit the default backup job options for Linux computers](#)
- [Uninstalling the Agent for Linux](#)
- [Starting the Agent for Linux daemon](#)
- [Stopping the Agent for Linux daemon](#)

- [Troubleshooting the Agent for Linux](#)

About the Agent for Linux

The Backup Exec Agent for Linux (Linux Agent) is installed as a separate add-on component. The Linux Agent enables network administrators to perform backup and restore operations on Linux servers that are connected to the network. The Linux Agent must be installed on the Linux servers before you can perform backup or restore operations.

See [“About open files and the Agent for Linux”](#) on page 1188.

See [“Requirements for the Agent for Linux”](#) on page 1188.

See [“About installing the Agent for Linux”](#) on page 1189.

About open files and the Agent for Linux

The Agent for Linux uses advanced open file and image technologies that are designed to alleviate the issues that are sometimes encountered during backup operations, such as backing up open files.

After you make file and folder selections and submit the job for backup, the Linux Agent automatically makes a snapshot of the volume or volumes. Making snapshot of a volume provides a point-in-time record of the data. When it creates a snapshot, the Linux Agent uses snapshot technologies to momentarily suspend write activity to a volume so that a snapshot of the volume can be created. During the backup, files can be open and data can be changed.

The Linux Agent supports Simple, Logical Volume Manager (LVM), and RAID volume configurations.

See [“Requirements for the Agent for Linux”](#) on page 1188.

Requirements for the Agent for Linux

The following items are required to install the Agent for Linux (Linux Agent):

- The Backup Exec server must have TCP/IP installed.
- You must have a root logon account on the Linux servers.
- You must have the Backup Exec installation media.
- You must enter a license for the Linux Agent on the Backup Exec server.

Note: Some versions of Linux may require that you install the `libstdc++.so.5` package.

See [“Troubleshooting the Agent for Linux”](#) on page 1215.

Symantec recommends that you use the Secure Shell (SSH) protocol when you push-install the Linux Agent from one Linux server to another Linux server. You must enable SSH before you push-install the Linux Agent.

Backup Exec automatically installs the Remote Media Agent for Linux when it installs the Agent for Linux on a Linux server. However, you must enter a separate license for the Remote Media Agent for Linux before it is available for use.

See [“About the Remote Media Agent for Linux ”](#) on page 1241.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

See [“About installing the Agent for Linux”](#) on page 1189.

See [“Requirements for backing up Novell Open Enterprise Server on SUSE Linux Enterprise Server”](#) on page 1207.

About installing the Agent for Linux

Use the Backup Exec installation media to do the following:

- Install the Agent for Linux (Linux Agent) on a local Linux server.
- Push-install the Linux Agent from one Linux server to other remote Linux servers. If you push-install the Linux Agent, the RSH (Remote Shell) is used by default. Symantec recommends that you use SSH (Secure Shell) instead. To use SSH, you must enable the SSH before you install the Linux Agent. Refer to your operating system documentation for more information about SSH.

When you install the Linux Agent, Backup Exec creates the beoper group and adds root as a member. The beoper group contains the names of the users who have permission to back up and restore the Linux servers. However, if Backup Exec detects an NIS server during the Linux Agent installation, then the beoper group is not created. You must create the beoper group manually on the Linux servers on which you want to install the Linux Agent.

When the installation is complete, Backup Exec saves the install log file to the following location on the server on which the Linux Agent is installed:

`/var/tmp/vxif/installralus<summary file number>/installralus.log`

See [“Installing the Agent for Linux ”](#) on page 1190.

See [“Troubleshooting the Agent for Mac”](#) on page 1238.

Installing the Agent for Linux

You can install the Agent for Linux (Linux Agent) on a local Linux server. You can then push-install the Linux Agent from the local Linux server to one or more remote Linux servers.

See [“About installing the Agent for Linux”](#) on page 1189.

To install the Agent for Linux

- 1 At a Linux server, place the Backup Exec installation media in the appropriate drive.
- 2 Log on as root on the server on which you want to install the Linux Agent.
- 3 Navigate to the following directory on the installation media.

<LinuxMac>

- 4 Copy the **RALUS_RMALS_RAM***S_<version number>.gz* file in this directory to a directory on the local computer.

- 5 Unzip the file.

For example:

```
gunzip RALUS_RMALS_RAMS_<version number>.gz
```

- 6 Untar the file.

For example:

```
tar -xf RALUS_RMALS_RAMS_<version number>.tar
```

- 7 Do one of the following:

To install the Linux Agent on the local Linux server Start the **installralus** script.

For example: `./installralus`

To install the Linux Agent from the local Linux server to one remote Linux server

Do the following in the order listed:

- Start the **installralus** script using the -SSH switch.

For example: `./installralus -usessh`

- Type the name, IP address, or fully qualified domain name of a Linux server.

To install the Linux Agent from the local Linux server to multiple remote Linux servers

Do the following in the order listed:

- Start the **installralus** script using the -SSH switch.

For example: `./installralus -usessh`

- Type the names, IP addresses, or fully qualified domain names of the Linux servers. Leave a space between each identifier.

- 8 After the installer checks for a valid Linux operating system during the initial system check, press **Enter**.
- 9 Review the package installation summary, and then press **Enter**.
- 10 After the system installation requirements check completes, press **Enter**.
- 11 Start the prerequisites check by pressing **Enter**.
- 12 Start the NIS server scan by pressing **Enter**.
- 13 Examine the results of the NIS server scan, and then do one of the following:

If an NIS server is detected

The Linux Agent installer cannot create the beoper group. You must create it manually after the Linux Agent installation is complete.

Continue with the next step.

If an NIS server is not detected

Use the installer to create the beoper group.

Do the following in the order listed:

- To let the installer create the beoper group, type **y**.
- To select the next available Group ID, type **n**.
- To add the root user account to the beoper group, type **y**.
- Continue with the next step.

- 14 Start the installation by pressing **Enter**.
- 15 After installation completes, press **Enter** to start the post-installation configurations and installation of SymSnap drivers.

- 16 Press **Y** to automatically start the Beremote service; otherwise, press **N** to start the service later.
- 17 After the configuration process completes, press **Enter** to save the installation log to the following file:
`/var/tmp/vxif/installralussummary file number/installralus.log`
- 18 If the Linux Agent installer did not create a beoper group, you must create it.
See [“Creating the Backup Exec operators group manually”](#) on page 1192.
- 19 Configure the Agent for Linux as appropriate.
See [“About configuring the Agent for Linux”](#) on page 1195.
- 20 If the Beremote service is not running, start the Agent for Linux daemon.
See [“Starting the Agent for Linux daemon”](#) on page 1214.

About the Backup Exec operators group for the Agent for Linux

The Backup Exec operators (**beoper**) group contains the names of the users who have permission to back up and restore the Linux servers.

When you install the Agent for Linux (Linux Agent), Backup Exec creates the **beoper** group and adds root as a member. Any Linux user that you add to the **beoper** group gets the necessary permissions to back up and restore the servers.

However, if an NIS server is detected during the Linux Agent installation, Backup Exec cannot create the **beoper** group. You must create the **beoper** group manually on the Linux servers on which you want to install the Linux Agent. You must create the **beoper** group before you start backup and restore operations. Otherwise, connections fail between the Linux servers and the Backup Exec server.

Before the members of the **beoper** group can perform backup or restore operations, they must have a Backup Exec logon account.

See [“Creating the Backup Exec operators group manually”](#) on page 1192.

See [“Backup Exec logon accounts”](#) on page 610.

Creating the Backup Exec operators group manually

You must create a beoper group on each server on which you want to install the Agent for Linux (Linux Agent).

See [“About the Backup Exec operators group for the Agent for Linux”](#) on page 1192.

Note: Ensure that you understand how to set security for groups on Linux servers before you assign a Group ID for the beoper group.

Table O-1 How to manually create the beoper group

Step	Action	More Information
Step 1	Navigate to the Linux server on which you want to install the Linux Agent. If the Linux server is in an NIS domain, navigate to the NIS domain's group file.	Refer to the NIS documentation for information on how to add a group to an NIS domain group file.
Step 2	Create a group with the following case-sensitive name: beoper	See the operating system's documentation for more information about how to create a group.
Step 3	In the beoper group, add the users that you want to have permission to back up and restore the Linux server.	See the operating system's documentation for more information about how to add users to a group.
Step 4	Create a Backup Exec logon account for each user that you add to the beoper group.	See "Backup Exec logon accounts" on page 610.

About establishing trust for a remote Linux computer in the Backup Exec list of servers

When you connect to a Linux computer from the Backup Exec server, you must establish trust between the Backup Exec server and the remote Linux computer. You must also establish trust if you want to configure a remote Linux computer to perform client-side deduplication.

See ["Establishing a trust between the Backup Exec server and a remote computer"](#) on page 780.

See ["Establishing trust and adding a remote Linux computer to the Backup Exec list of servers"](#) on page 1193.

Establishing trust and adding a remote Linux computer to the Backup Exec list of servers

You can add one or more remote Linux computers to the list of servers that appear on the **Backup and Restore** tab. When you add remote Linux computers, you must

establish a trust between the Backup Exec server and the remote Linux computers to ensure secure communication.

To establish trust and add a remote Linux computer to the Backup Exec list of servers

- 1 On the **Backup and Restore** tab, in the **Servers** group, click **Add**.
- 2 Click **Linux computer**.
- 3 Follow the on-screen prompts.

See [“Adding additional Backup Exec servers to which the Agent for Linux can publish information”](#) on page 1194.

See [“About configuring the Agent for Linux”](#) on page 1195.

Adding additional Backup Exec servers to which the Agent for Linux can publish information

You can specify additional Backup Exec servers to which the Agent for Linux (Linux Agent) can publish information.

Each Backup Exec server to which the Linux Agent publishes information appears in the Backup Exec **Servers** list.

To add additional Backup Exec servers to which the Agent for Linux can publish information

- 1 Use a text editor to open the following file:

```
/etc/VRTSralus/ralus.cfg
```

- 2 Add the following string:

```
Software\Symantec\Backup Exec For Windows\Backup  
Exec\Engine\Agents\Agent Directory List unique identifier number = IP address  
or DNS name of Backup Exec server
```

- 3 Save and close the file.
- 4 Move to the Backup Exec server to which the Linux Agent is publishing itself and add the Linux server to the **Servers** list.

See [“About the list of servers on the Backup and Restore tab”](#) on page 152.

About configuring the Agent for Linux

Backup Exec creates a file named `ralus.cfg` on each Linux server on which the Agent for Linux (Linux Agent) is installed. You can edit the strings, identifiers, and variables in this file to add or edit options for the Linux Agent.

Options that you can edit in the `ralus.cfg` file include the following:

- The port to which the Linux Agent must send publishing messages.
- The logging level for Oracle database operations that use the Backup Exec Linux Agent Utility, and for NDMP information.
- The settings to allow the Linux Agent to publish to one or more Backup Exec servers.
- The files and directories on Linux servers that you want to exclude from backups.
- The setting for a Target Service Agent File System backup for Novell OES.

The `ralus.cfg` format contains three components. The first component (A) in the following example is a required string.

The second component (B) is a unique identifier followed by an equal sign (=). A unique identifier can consist of sequential numbers, letters, or alpha-numeric characters. For example, 1, 2, 3 or A, B, C. You can also use AA, BB, CC, or A1, A2, B1, B2.

The third component of the `ralus.cfg` format is the NetBIOS name, fully qualified domain name, or IP address of the Backup Exec server.

The `ralus.cfg` includes a registry key that works with the Linux Agent's open file technology. The name of the key is `DisableOFO` and appears in the `ralus.cfg` file in the following form:

```
Software\Symantec\Backup Exec for Windows\Backup  
Exec\Engine\RALUS\DisableOFO=0
```

By default, the `DisableOFO` key is set to 0, meaning that the Linux Agent is active, letting the Linux Agent back up the open files that it encounters. However, you can disable the open file technology by changing the value of the key to "1", and then restarting the Linux Agent daemon.

Figure O-1 Example of the `ralus.cfg` file

A

B

C

Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Agent Directory List 1=srv.mycompany.com
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Agent Directory List 2=datasrv
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Agent Directory List 3=66.35.250.151

A = Required string
B = Required and unique identifier (the order or appearance is irrelevant)
C = File or directory to be excluded

- See [“Editing configuration options for Linux computers”](#) on page 1196.
- See [“Configuration options for Linux computers”](#) on page 1197.
- See [“Stopping the Agent for Linux daemon”](#) on page 1215.
- See [“Starting the Agent for Linux daemon”](#) on page 1214.

About excluding files and directories from backup jobs for Linux computers

You can exclude specific files and directories on the Linux computers from all backup jobs. Edit the `ralus.cfg` file to specify the excluded files.

The following is an example of strings in the `ralus.cfg` file that excludes files and directories from all backup jobs.

Figure O-2 Example of file and directory exclusions in the `ralus.cfg` format

A

B

C

Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\RALUS\SystemExclude1=/dev/*.*
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\RALUS\SystemExclude2=/proc/*.*
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\RALUS\SystemExclude3=/mnt/nss/pools/
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\RALUS\SystemExclude4=/mnt/nss/.pools/

A = Required string
B = Required and unique identifier (the order or appearance is irrelevant)
C = File or directory to be excluded

- To exclude files and directories for specific backup jobs, specify the exclusions in the backup job properties.
- See [“Editing configuration options for Linux computers”](#) on page 1196.

Editing configuration options for Linux computers

- You can edit configuration options for the Agent for Linux.
- See [“About configuring the Agent for Linux”](#) on page 1195.

To edit configuration options for Linux computers

- 1 Use a text editor to open the following file:
- /etc/VRTSralus/ralus.cfg
- 2 Change the appropriate string in the file.
- See “[Configuration options for Linux computers](#)” on page 1197.

Configuration options for Linux computers

You can edit options to configure the Agent for Linux (Linux Agent).

See “[Editing configuration options for Linux computers](#)” on page 1196.

Table 0-2 Configuration options for Linux computers

String and default values	Description
Software\Symantec\Backup Exec For Windows\Agent Browser\TcpIp\AdvertisementPort=6101	Lists the port to which the Linux Agent must send publish and purge messages.
Software\Symantec\Backup Exec for Windows\Backup Exec\Debug\AgentConfig=0	Enables logging for the Linux Agent utility that Oracle operations use. Values include the following: <ul style="list-style-type: none">■ 0 Logging is not enabled.■ 1 Logging is enabled. Backup Exec automatically generates the log file. This option does not apply to the Agent for Mac.
Software\Symantec\Backup Exec for Windows\Backup Exec\Debug\VXBSAlevel=0	Enables logging for the Linux Agent for Oracle operations. Values include the following: <ul style="list-style-type: none">■ 0 Logging is not enabled.■ 5 Normal logging is enabled.■ 6 Advanced logging is enabled. Large log files may be created. This option does not apply to the Agent for Mac.

Table 0-2 Configuration options for Linux computers *(continued)*

String and default values	Description
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\Agents\Advertise All=1	<p>Enables the Linux Agent to publish information to all of the Backup Exec servers that are listed in the \Agents\Agent Directory List strings.</p> <p>Values include the following:</p> <ul style="list-style-type: none">■ 1 The Linux Agent publishes information to every Backup Exec server in the Agent Directory List.■ 0 The Linux Agent publishes information to the first Backup Exec server in the Agent Directory List. If the attempt is successful, the Linux Agent does not publish information to any other Backup Exec servers. If the attempt is not successful, the Linux Agent attempts to publish information to the next Backup Exec server in the list. Attempts continue until the Linux Agent reaches the end of the list.

Table 0-2 Configuration options for Linux computers *(continued)*

String and default values	Description
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Advertise Now=0	<p>Enables the Linux Agent to start a new publishing cycle after you add or edit any settings in the ralus.cfg file.</p> <p>Values include the following:</p> <ul style="list-style-type: none">■ 0 The Linux Agent publishes information according to its regular cycle, set in the string \Agents\Advertising Interval Minutes. Any changes to the ralus.cfg file take effect when a new publishing cycle begins.■ 1 The Linux Agent starts a new publishing cycle. Any changes to the ralus.cfg file take effect immediately. <p>If the Backup Exec server does not receive the publishing information, the Linux Agent makes 10 more attempts. Each attempt to publish information to the Backup Exec server is one minute apart. If the information is not sent at the end of the 10 attempts, the Linux Agent skips that Backup Exec server until the next publishing cycle. The publishing cycle is the number of minutes set in the string \Agents\Advertising Interval Minutes.</p>

Table O-2 Configuration options for Linux computers (*continued*)

String and default values	Description
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Advertisement Purge=0	<p>Lets the Linux Agent send a purge message to all of the Backup Exec servers in the string \Agents\Advertisement Purge. When a Backup Exec server receives a purge message, it removes the Linux Agent from Backup Exec's list of available servers. The Linux Agent continues to function.</p> <p>Values include the following:</p> <ul style="list-style-type: none">■ 0 Do not purge the Linux Agent from any Backup Exec servers that are listed in the \Agents\Advertisement Purge string.■ 1 Purge the Linux Agent from one or more Backup Exec servers in the \Agents\Advertisement Purge string.
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Advertising Disabled=0	<p>Enables the Linux Agent to publish to Backup Exec servers.</p> <p>Values include the following:</p> <ul style="list-style-type: none">■ 0 The Linux Agent attempts to publish information to the Backup Exec servers that are listed in the string \Agents\Agent Directory List.■ 1 The Linux Agent does not publish information to Backup Exec servers.
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Advertising Interval Minutes=240	<p>Sets the number of minutes that the Linux Agent must wait between publishing cycles. The default number of minutes is 240. The range of minutes is from 1 minute to 720 minutes.</p>
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Agent Directory List_1=<Backup Exec server name>	<p>Displays the list of NetBIOS names, fully qualified domain names, or IP addresses to which the Linux Agent publishes information.</p> <p>The Backup Exec server from which the Linux Agent is push-installed is added to the Agent Directory List by default.</p>

Table 0-2 Configuration options for Linux computers (*continued*)

String and default values	Description
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Auto Discovery Enabled=1	<p>Adds a Backup Exec server to the string \Agents\Agent Directory List if the Backup Exec server performs a backup job with which the Linux Agent is associated.</p> <p>Values include the following:</p> <ul style="list-style-type: none">■ 1 Adds the Backup Exec server that performs the backup job to the Agent Directory List. The Linux Agent can publish information to the Backup Exec server.■ 0 The Backup Exec server that performs the backup job is not added to the Agent Directory List.
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\Logging\RANT NDMP Debug Level=0	<p>Displays the level of verbosity for logging NDMP information for the Linux Agent.</p> <p>Values include the following:</p> <ul style="list-style-type: none">■ 0 Logs only the NDMP errors.■ 1 Logs the NDMP errors and warnings.■ 2 Logs the NDMP errors, warnings, and message information that is sent between the remote computer and the Backup Exec server.
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RALUS\Encoder=	<p>Displays the encoder that you can add if the default encoder incorrectly displays characters on the user interface.</p>

Table O-2 Configuration options for Linux computers (*continued*)

String and default values	Description
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RALUS\ShowTSAFS=	<p>Lets you perform a Target Service Agent file system (TSAFS) backup for applications on Novell Open Enterprise Services. By default, this option is not enabled.</p> <p>The Linux Agent backs up all file systems using the Root object. If ShowTSAFS is enabled, the Novell Open Enterprise Services resource appears in the backup selection list. If you select the whole computer for backup, then redundant backups are performed. Symantec recommends that you do not enable this option.</p> <p>Values include the following:</p> <ul style="list-style-type: none">■ Blank or 0 The file system TSA does not appear for backup selection.■ 1 The file system TSA resource appears for backup selection. <p>This option does not apply to the Agent for Mac.</p>
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RALUS\SystemExclude1=	<p>Lists the files that you want to exclude from all Linux Agent backup jobs.</p> <p>See “About excluding files and directories from backup jobs for Linux computers” on page 1196.</p>
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RALUS\SystemFSTypeExclude1	<p>Lists the type of file system that you want to exclude from the Linux Agent backup.</p>

Table 0-2 Configuration options for Linux computers (continued)

String and default values	Description
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RMAL\DisableRMAL=0	<p>Lets you use the Remote Media Agent for Linux to back up the Linux server on which it is installed. By default, this option is not enabled.</p> <p>If you install RMAL to an unsupported version of Linux, RMAL is unavailable for use. You cannot create the jobs that run on the devices that are attached to the Linux server. However, you can back up the Linux server by using the Agent for Linux component. This component is installed with RMAL. You must change the value of this string to 1 to use the Agent for Linux component.</p> <p>Values include the following:</p> <ul style="list-style-type: none">0 You can create backup, restore, and utility jobs on the Backup Exec server that run on the Linux server's storage devices.1 You can only use the Agent for Linux component to back up the Linux server on which it is installed. <p>See "Troubleshooting the Agent for Linux" on page 1215.</p>
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RALUS\DisableOFO=0	<p>Lets you disable or enable the Linux Agent's open file technology.</p> <p>By default, the DisableOFO key is set to 0, meaning that the Linux Agent is active, letting the Linux Agent back up the open files that it encounters. However, you can disable the open file technology by changing the default value of the key to 1 and then restarting the Linux Agent daemon.</p>

Table 0-2 Configuration options for Linux computers (*continued*)

String and default values	Description
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RALUS\AOFOL\CacheFileMountPoint =	<p>The alternate cache file location for Advanced Open File Option (AOF). This location should be the mount point.</p> <p>By default, snapshot cache file is created on the Volume mount point.</p> <p>If you specify the location, the Linux Agent uses this location to create the snapshot cache file.</p>
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RALUS\AOFOL\CacheFileSize =	<p>The alternate cache file size. The size should be at least 10% of the volume size that you plan to back up and the size value must be multiples of 4-KB blocks.</p> <p>The Linux Agent accepts the alternate cache file location and alternate cache file size when values for both are correctly configured.</p>

About backing up a Linux computer by using the Agent for Linux

The following backup methods appear when you use the Agent for Linux (Linux Agent) to back up data:

- Full - Using modified time
- Differential - Using modified time
- Incremental - Using modified time

However, on the Novell OES server, only full backups are supported. If you select any other backup method, a full backup runs instead. The Agent for Linux also supports synthetic full backup. But, Novell OES does not support synthetic full backup.

See [“Linux backup options”](#) on page 1205.

See [“Adding a stage to a backup definition”](#) on page 219.

See [“Editing a stage”](#) on page 221.

See [“About backing up Novell Open Enterprise Server \(OES\) components ”](#) on page 1206.

Linux backup options

The following Agent for Linux (Linux Agent) options are available when you back up Linux computers.

See [“About backing up a Linux computer by using the Agent for Linux”](#) on page 1204.

Table 0-3 Backup job options for Linux computers

Item	Description
Preserve file and directory timestamps during backups	<p>Prevents the Linux Agent from changing an object's attributes when a backup occurs. An object is a file or a directory.</p> <p>This option is not selected by default.</p> <p>During a backup, Backup Exec preserves an object's last access timestamp by resetting the last access timestamp to the value before the backup occurred. When Backup Exec modifies the object's last access timestamp, the operating system internally updates the object's ctime.</p> <p>An object's ctime is the time when an object's attributes such as permissions and timestamps are modified. If the Linux Agent does not change the attributes after a backup, the object's ctime is not changed.</p> <p>This option does not affect the object attributes that are set during restore operations.</p>
Follow local mount points	<p>Lets Backup Exec follow local mount points when it backs up data.</p> <p>This option is enabled by default.</p> <p>For more information on local mount points, see your operating system's documentation.</p>

Table O-3 Backup job options for Linux computers (*continued*)

Item	Description
Follow remote mount points	<p>Lets Backup Exec follow remote mount points when it backs up data.</p> <p>This option is not selected by default.</p> <p>When you use this option, the following limitations apply:</p> <ul style="list-style-type: none">■ The data that is mounted must reside on a computer type that Backup Exec supports. <p>You can find a list of supported operating systems, platforms, and applications at the following URL:</p> <p>http://entsupport.symantec.com/umi/V-269-1</p> <ul style="list-style-type: none">■ If the mount point leads to an operating system that Backup Exec does not support, contact the operating system's vendor to resolve any issues. <p>For more information on remote mount points, see your operating system's documentation.</p>
Lock remote files to prevent applications from modifying them during backups	<p>Lets the Linux Agent have exclusive access to the files on the remote servers that are connected through Network File System (NFS). Locking remote files prevents other applications from modifying the data during the backup.</p>
Backup method for eDirectory	<p>Deletes a backup method for backing up eDirectory data for Novell OES on SUSE Linux Enterprise Server.</p> <p>Note: Backup Exec does not support this option for Macintosh computers.</p> <p>See “About restoring Novell OES components” on page 1208.</p>

About backing up Novell Open Enterprise Server (OES) components

Before you can back up Novell OES components, the Agent for Linux must be installed on the server where the Novell OES components reside.

Backup Exec supports the following Novell Open Enterprise Server (OES) components:

- Novell iFolder
- Novell eDirectory
- Novell GroupWise
- Novell Storage Services (NSS)

Note: Backup Exec 2014 does not support GroupWise12.

While backing up Novell OES components, even if you select an incremental or differential backup option, Backup Exec backs up the full data. However, in the restore view, even though the backup instance available for restore displays incremental or differential, selecting such a backup data, restores the full backup.

See [“Requirements for backing up Novell Open Enterprise Server on SUSE Linux Enterprise Server”](#) on page 1207.

Requirements for backing up Novell Open Enterprise Server on SUSE Linux Enterprise Server

Backup Exec requires the following to back up Novell OES:

- To protect Novell eDirectory, iFolder, and GroupWise, the necessary Target Service Agents (TSA) must be loaded. See the Novell documentation for details on loading the TSA for these Novell OES components.
- A local UNIX user name that is the equivalent of the admin-level eDirectory user in the beoper group. Backup Exec does not support eDirectory users.
See [“About the Backup Exec operators group for the Agent for Linux”](#) on page 1192.
- A Backup Exec logon account that contains the credentials for the equivalent admin-level eDirectory user must exist before you can perform backup jobs for eDirectory.

See [“Backing up data”](#) on page 158.

See [“About backing up a Linux computer by using the Agent for Linux”](#) on page 1204.

About restoring data to Linux computers

You can specify restore job options to restore Linux computers.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

See [“Restore job options for Linux computers”](#) on page 1208.

About restoring Novell OES components

While backing up Novell OES components, even if you select an incremental or differential backup option, Backup Exec backs up the full data. However, in the restore view, even though the backup instance available for restore displays incremental or differential, selecting such a backup data, restores the full backup.

Backup Exec does not support the redirected restore of Novell OES components iFolder and eDirectory from an OES environment to a non-OES environment.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

Restore job options for Linux computers

See [“About restoring data to Linux computers”](#) on page 1207.

Table 0-4 Restore job options for Linux computers

Item	Description
Lock remote files if the mount points have neccessary permissions	Lets Backup Exec have exclusive access to the files on the remote computers that are connected through the Network File System (NFS). This option is enabled by default.
Restore DIB set	Restores the Directory Information Base (DIB), also known as the Novell directory services (NDS) database.
Activate DIB after verify	Lets Backup Exec rename the database from .RST to .NDS after the verification process completes successfully. If the verify operation fails, the .RST file is deleted and the original .NDS file is kept intact. If you do not select this option, after the database is restored, the .RST file is available for you to perform manual activation or manual disaster recovery.
Open database when finished	Lets Backup Exec open the database after the restore completes. If you want to perform maintenance tasks before the database opens, do not select this option.

Table O-4 Restore job options for Linux computers (*continued*)

Item	Description
Verify database after restore	Lets Backup Exec verify the database after the restore completes.
Roll forward log directory	Displays the location of the roll forward log directory.
Leave backup file on disk	Keeps the Novell DIB fileset on the hard drive. See “About restoring Novell OES components” on page 1208.

Edit the default backup job options for Linux computers

You can edit the existing default options for all backup and restore jobs for Linux systems.

To edit default backup job options for Linux systems

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Job Defaults**.
- 2 Select either **Back Up to Disk** or **Back Up to Tape**, and then select **Linux and Macintosh**.
- 3 Set the appropriate options.
See [“Default backup job options for Linux computers”](#) on page 1209.
- 4 Click **OK**.

Default backup job options for Linux computers

You can set default backup job properties for all jobs on Linux computers.

See [“Edit the default backup job options for Linux computers”](#) on page 1209.

You can find a list of supported operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

Table 0-5 Default backup job options for Linux computers

Item	Description
Preserve file and directory timestamps during backups	<p>Prevents the Linux Agent from changing an object's attributes when a backup occurs. An object is a file or a directory.</p> <p>This option is not selected by default.</p> <p>During a backup, Backup Exec preserves an object's last access timestamp by resetting the last access timestamp to the value before the backup occurred. When Backup Exec modifies the object's last access timestamp, the operating system internally updates the object's ctime.</p> <p>An object's ctime is the time when an object's attributes such as permissions, timestamps, etc., have been modified. If the Linux Agent does not change the attributes after a backup, the object's ctime is not changed.</p> <p>This option does not affect the object attributes that are set during restore operations.</p>
Follow local mount points	<p>Lets Backup Exec follow local mount points when it backs up data.</p> <p>This option is enabled by default.</p> <p>For more information on local mount points, see your operating system's documentation.</p>
Follow remote mount points	<p>Lets Backup Exec follow remote mount points when it backs up data.</p> <p>This option is not selected by default.</p> <p>When you use this option, the following limitations apply:</p> <ul style="list-style-type: none">■ The data that is mounted must reside on an operating system that Backup Exec supports.■ If the mount point leads to an operating system that Backup Exec does not support, contact the operating system's vendor to resolve any issues. <p>For more information on remote mount points, see your operating system's documentation.</p>
Backup method for eDirectory	<p>Displays a backup method for backing up eDirectory data for Novell OES on SUSE Linux Enterprise Server.</p> <p>Note: This option is not supported for Macintosh computers.</p> <p>See "About backing up Novell Open Enterprise Server (OES) components" on page 1206.</p>

Table 0-5 Default backup job options for Linux computers (continued)

Item	Description
Lock remote files to prevent applications from modifying them during backups	Lets the Linux Agent have exclusive access to the files on the remote servers that are connected through Network File System (NFS). Locking remote files prevents other applications from modifying the data during the backup or restore job.

Uninstalling the Agent for Linux

An automated uninstall process for the Agent for Linux (Linux Agent) is available on the Backup Exec installation media.

/opt/VRTS/install/logs/uninstallralus<summary file number>.summary

To uninstall the Agent for Linux

- On the Linux server, place the Backup Exec installation media in the appropriate device.
- Log on as root to the server from which you want to uninstall the Linux Agent.
- Navigate to the following directory on the Backup Exec installation media:
<LinuxMac>
- Start the **uninstallralus** script.
For example:

```
./uninstallralus
```
- Do one of the following:

To uninstall the Linux Agent from one server	Type the name, IP address, or fully qualified domain name of a Linux server.
To uninstall the Linux Agent from multiple servers	Type the names, IP addresses, or fully qualified domain names of the Linux servers. Leave a space between each identifier.
- Press **Enter**.
- After the Linux Agent package check completes successfully, press **Enter**.
- When you are prompted to uninstall the RALUS packages, press **Enter**.

- 9 When you are prompted to uninstall the SymSnap driver, press **Enter**.
- 10 To save the uninstall summary to the following location on the Linux server, press **Enter**:

`/opt/VRTS/install/logs/uninstallralus<summary file number>.summary`

See [“Installing the Agent for Linux”](#) on page 1190.

Manually uninstalling the Agent for Linux

You can manually uninstall the Agent for Linux (Linux Agent).

To manually uninstall the Agent for Linux

- 1 Use a terminal session to connect to the Linux server as the root user.
- 2 Change to the following directory:

`/opt/VRTSralus/bin`

For example:

```
cd /opt/VRTSralus/bin
```

- 3 Delete the following line if it is found in the `/etc/inittab` file:

`/opt/VRTSralus/bin/VRTSralus.init`

For example:

```
rm -r /opt/VRTSralus/bin/VRTSralus.init
```

- 4 Copy the `RALUS_RMALS_RAMs_<version number>.gz` file in this directory to a directory on the local computer.

- 5 Unzip the file using the following command:

```
gunzip RALUS_RMALS_RAMs_<version number>.gz
```

- 6 Untar the file using the following command:

```
tar -xf RALUS_RMALS_RAMs_<version number>.tar
```

- 7 Stop the Linux Agent daemon.

See [“Stopping the Agent for Linux daemon”](#) on page 1215.

- 8 Remove the Linux Agent package from the Linux server.

For example:

Debian GNU/Linux, Ubuntu	<code>dpkg -r VRTSralus</code>
--------------------------	--------------------------------

Linux	<code>rpm -e VRTSralus</code>
-------	-------------------------------

9 Change back to the root directory.

For example:

```
cd /
```

10 Remove the following files:

```
/etc/VRTSralus
```

```
/opt/VRTSralus
```

```
/var/VRTSralus
```

For example:

```
rm -r /etc/VRTSralus /opt/VRTSralus /van/VRTSralus
```

11 Type **y** if you are prompted to descend into the directories.

12 Type **y** if you are prompted to delete a directory.

13 Remove run-time scripts if they are present.

See [“Run-time scripts to remove when manually uninstalling the Agent for Linux”](#) on page 1213.

Run-time scripts to remove when manually uninstalling the Agent for Linux

When you manually uninstall the Agent for Linux (Linux Agent), remove the following run-time scripts if they are present.

Table 0-6 Run-time scripts to remove when manually uninstalling the Linux Agent

Operating system	Run-time scripts to remove
Debian, Ubuntu	<div><div>/etc/rc5.d/S95VRTSralus.init</div><div>/etc/rc3.d/S95VRTSralus.init</div><div>/etc/rc2.d/S95VRTSralus.init</div><div>/etc/init.d/VRTSralus.init</div><div>For example:</div><div>rm /etc/rc5.d/S95VRTSralus.init</div></div>

Table O-6

Run-time scripts to remove when manually uninstalling the Linux Agent *(continued)*

Operating system	Run-time scripts to remove
Red Hat Linux, Asianux	<div>/etc/rc.d/rc5.d/S95VRTSralus.init</div> <div>/etc/rc.d/rc3.d/S95VRTSralus.init</div> <div>/etc/rc.d/rc2.d/S95VRTSralus.init</div> <div>/etc/rc.d/init.d/VRTSralus.init</div> <div>For example:</div> <div>rm /etc/rc.d/rc5.d/S95VRTSralus.init</div>
Novell Open Enterprise Server 1.0/ SUSE Linux Enterprise Server 9 (32-bit only)	<div>/etc/init.d/rc5.d/SxxVRTSralus.init</div> <div>/etc/init.d/rc3.d/SxxVRTSralus.init</div> <div>/etc/init.d/rc2.d/SxxVRTSralus.init</div> <div>/etc/init.d/VRTSralus.init</div> <div>For example:</div> <div>rm /etc/init.d/rc5.d/SxxVRTSralus.init</div>
Novell Open Enterprise Server 2.0/ SUSE Linux Enterprise Server 10 (32-bit and 64-bit)	<div>/etc/init.d/VRTSralus.init,start=2,3,5</div> <div>/etc/init.d/VRTSralus.init</div> <div>For example:</div> <div>rm /etc/init.d/VRTSralus.init</div>

See [“Manually uninstalling the Agent for Linux ”](#) on page 1212.

Starting the Agent for Linux daemon

If necessary, you can start the Agent for Linux (Linux Agent) daemon after the operating system starts.

See [“Stopping the Agent for Linux daemon”](#) on page 1215.

To start the Agent for Linux daemon

- 1 Use a terminal session to connect to the Linux server as the root user.
- 2 Navigate to the following directory:

```
/etc/init.d/
```

For example:

```
cd /etc/init.d/
```

- 3 Start the Linux Agent daemon.

For example:

```
/etc/init.d/VRTSralus.init start
```

Stopping the Agent for Linux daemon

You can stop the Agent for Linux (Linux Agent) daemon.

See [“Starting the Agent for Linux daemon”](#) on page 1214.

To stop the Agent for Linux daemon

- 1 Use a terminal session to connect to the Linux server as the root user.
- 2 Navigate to the following directory:

```
/etc/init.d/
```

For example:

```
cd /etc/init.d/
```

- 3 Stop the Linux Agent daemon:

For example:

```
/etc/init.d/VRTSralus.init stop
```

- 4 Restart the daemon when necessary.

Troubleshooting the Agent for Linux

If you experience problems with the Agent for Linux (Linux Agent) review the following questions and answers.

See [“About the Agent for Linux”](#) on page 1188.

Table 0-7 Troubleshooting the Linux Agent

Question	Answer
Some characters do not appear correctly in the terminal session during the installation. What should I do?	This error occurs when the system location uses a non-English language character-set on the computer on which you install the Linux Agent. You can switch to another location setting of the same language to try to resolve this issue.
The Linux Agent installer is unable to install the Linux Agent. The following error is reported in the installralus log file: VxIF::Error:: Unable to compress files. Hash(0x8711e8)->({GUNZIP} not found on <hostname>	To support the uncompressing of the Linux Agent platform-specific packages, you can install the GNU data compression utility. Install this utility on the computer on which you want to install the Linux Agent. The utility is available at the following URL: http://www.gzip.org
The Agent for Linux is installed on a Linux server in an NIS domain. Backup Exec is unable to browse resources on the server. What should I do?	Verify if the group line and the password line in the nsswitch.conf file are set to compatibility mode. If they are, then you must configure the /etc/passwd and /etc/group files. Refer to the nsswitch.conf man pages for additional information on how to configure the nsswitch.conf to use compatibility mode. Alternatively, change the password line and the group line to NIS files so that the Linux server validates the user through NIS. If the NIS server is unavailable or if the user is not found, the local files are used for validation.

Table O-7 Troubleshooting the Linux Agent (*continued*)

Question	Answer
<p>I cannot load the Linux Agent. When I attempt to load the Linux Agent in console mode, /beremote --log-console shows the following message:</p> <p>ACE_SV_Semaphore_Complex: no space left on device.</p> <p>What should I do?</p>	<p>This issue occurs when the computer reaches its maximum limit on allowable semaphores. It can occur after an unexpected termination of the Linux Agent. When the Linux Agent unexpectedly terminates, it is unable to clean up some of the semaphore resources that it used. Other processes may have caused the use of semaphores to reach the limit. You must restart the computer to safely recover it from this condition.</p> <p>If other processes are running, it may not be feasible to restart the computer. Instead, you can use the commands that let you list and then remove all semaphores in use by the operating system. Be careful when you select semaphores to remove. Semaphores that are in use by the Linux Agent cannot be identified. If you remove semaphores of other programs that are in use, those programs can become unstable.</p> <p>To list semaphores, you can type the following command:</p> <pre>ipcs -a</pre> <p>To remove semaphores for each identifier that is listed, you can type the following command:</p> <pre>ipcrm -s <id></pre>
<p>I cannot load the Linux Agent. When I attempt to load the Linux Agent in console mode, /beremote --log-console shows the following message:</p> <p>Error while loading shared libraries: libstdc++.so.5: cannot open shared object file: No such file or directory.</p> <p>What should I do?</p>	<p>This error indicates that the libstdc++.so.5 library is not in the /usr/lib directory. This library is necessary to let the Linux Agent start and function. To resolve this issue, install the libstdc++5 package.</p> <p>You can install this package from the media on which your copy of Linux was provided. Or, you can run the following command from a computer that has Internet access:</p> <pre>apt-get install libstdc++5</pre> <p>For SUSE Linux Enterprise Server 11, run the following command:</p> <pre>zypper install libstdc++5</pre>

Symantec Backup Exec Agent for Mac

This appendix includes the following topics:

- [About the Agent for Mac](#)
- [Requirements for the Agent for Mac](#)
- [About the Backup Exec admin group on Macintosh systems](#)
- [About installing the Agent for Mac](#)
- [Uninstalling the Agent for Mac](#)
- [About configuring the Agent for Mac](#)
- [Starting the Agent for Mac](#)
- [Stopping the Agent for Mac](#)
- [About establishing trust for a remote Macintosh system in the Servers list](#)
- [Adding additional Backup Exec servers to which the Agent for Mac can publish information](#)
- [About backing up data by using the Agent for Mac](#)
- [About restoring Macintosh systems](#)
- [Troubleshooting the Agent for Mac](#)

About the Agent for Mac

The Agent for Mac (Mac Agent) is installed as a separate add-on component. The Mac Agent enables Windows Servers network administrators to perform backup and restore operations on Macintosh systems that are connected to the network. The Mac Agent must be installed on the Macintosh systems before you can perform backup or restore operations.

See [“Requirements for the Agent for Mac”](#) on page 1219.

See [“About installing the Agent for Mac”](#) on page 1221.

Requirements for the Agent for Mac

The following are required to install the Agent for Mac (Mac Agent):

- The Backup Exec server must have TCP/IP installed.
- You must be a member of the admin group on the Macintosh system on which you want to install the Mac Agent.
- You must have the Backup Exec installation media.
- You must enter a license for the Mac Agent on the Backup Exec server.

Symantec recommends that you use the Secure Shell (SSH) protocol when you push-install the Mac Agent from a local Macintosh system to other remote Macintosh systems. You must enable SSH before you install the Mac Agent.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

See [“About installing the Agent for Mac”](#) on page 1221.

See [“About the Backup Exec admin group on Macintosh systems”](#) on page 1219.

About the Backup Exec admin group on Macintosh systems

The Backup Exec admin group contains the name of the users who have permission to back up and restore the Macintosh systems.

When you install the Agent for Mac (Mac Agent), Backup Exec creates the admin group and adds root as a member. Any Macintosh user that you add to the admin group gets the necessary permission to back up and restore the Macintosh systems.

However, if an NIS server is detected during the Mac Agent installation, Backup Exec cannot create the admin group. After the installation, you must create the admin group manually on the Macintosh system. You must create the admin group before you start backup and restore operations. Otherwise, connections fail between the Macintosh system and the Backup Exec server .

Before the members of the admin group can perform backup or restore operations, they must have a Backup Exec logon account.

See [“Creating the Backup Exec admin group manually on Macintosh systems”](#) on page 1220.

See [“Backup Exec logon accounts”](#) on page 610.

Creating the Backup Exec admin group manually on Macintosh systems

You must create an admin group on each Macintosh system on which you want to install the Agent for Mac (Mac Agent).

See [“About the Backup Exec admin group on Macintosh systems”](#) on page 1219.

Note: Ensure that you understand how to set security for groups on Macintosh systems before you assign a Group ID for the admin group.

Table P-1 How to manually create the admin group

Step	Action	More Information
Step 1	Navigate to the Macintosh system on which you want to install the Mac Agent. If the Macintosh system is in a NIS domain, navigate to the NIS domain's group file.	See the NIS documentation for information on how to add a group to a NIS domain group file.
Step 2	Create a group with the following case-sensitive name: admin	See the Macintosh documentation for more information about how to create a group.
Step 3	In the admin group, add the users that you want to have permission to back up and restore the Macintosh system.	See the Macintosh documentation for more information about how to add users to a group

Table P-1 How to manually create the admin group (*continued*)

Step	Action	More Information
Step 4	Create a Backup Exec logon account for each user that you add to the admin group.	See “Backup Exec logon accounts” on page 610.

About installing the Agent for Mac

Use the Backup Exec installation media to do the following:

- Install the Agent for Mac (Mac Agent) on a local Macintosh system.
- Push-install the Mac Agent from a local Macintosh system to other remote Macintosh systems. If you push-install the Mac Agent, the RSH (Remote Shell) is used by default. Symantec recommends that you use SSH (Secure Shell) instead. To use SSH, you must enable it before you install the Mac Agent. See your Macintosh documentation for more information about SSH.

When the installation is complete, Backup Exec saves the install log file to the following location on the system on which the Mac Agent is installed:

/var/tmp/vxif/installrams <unique identifier number> for installs

Note: Some characters may not appear correctly in the terminal session during the installation. This error occurs when the system location uses a non-English language character-set on the computer on which you install the Mac Agent. You can switch to another location setting of the same language to try to resolve this issue.

See [“Installing the Agent for Mac”](#) on page 1221.

Installing the Agent for Mac

You can install the Agent for Mac (Mac Agent) on a local Macintosh system. You can also push-install the Mac Agent from a local Macintosh system to one or more remote Macintosh systems.

See [“About installing the Agent for Mac”](#) on page 1221.

Note: You must unzip the RALUS_RMALS_RAMs_<version number>.gz file on a Macintosh system. The installation does not run if it is unzipped on a computer that runs the Windows operating system.

To install the Agent for Mac

- 1 At a Macintosh system, place the Backup Exec installation media in the appropriate drive.
- 2 Navigate to the following directory on the installation media:
<LinuxMac>
- 3 Copy the RALUS_RMALS_RAMs_<version number>.gz file in this directory to a directory on the local system.
- 4 Unzip the file.
For example:

```
gunzip RALUS_RMALS_RAMs_<version number>.gz
```
- 5 Untar the file.
For example:

```
tar -xf RALUS_RMALS_RAMs_<version number>.tar
```
- 6 Open **Finder**, and then browse to **Applications > Utilities**.
- 7 Open **Terminal**.
- 8 Do one of the following:

To install the Mac Agent on a local system Do the following in the order listed:

- Start the **installrams** script by typing the following command:

```
sudo ./installrams
```
- Press **Enter**.

To install the Mac Agent from a local Macintosh computer to a remote Macintosh system

Do the following in the order listed:

- Start the **installrams** script using the **-usessh** switch by typing the following command:

```
sudo ./installrams -usessh
```
- Press **Enter**.
- Type the name, IP address, or fully qualified domain name of a Macintosh system.

To install the Mac Agent from a local Macintosh computer to multiple remote Macintosh systems

Do the following in the order listed:

- Start the **installrams** script using the **-usessh** switch by typing the following command:

```
sudo ./installrams -usessh
```

- Press **Enter**.
- Type the names, IP addresses, or fully qualified domain names of the Macintosh systems. Leave a space between each identifier.

- 9 Enter the password for the user name that is currently logged on.
- 10 Press **Enter**.
- 11 After the installer checks for a valid Macintosh system operating system during the initial system check, press **Enter**.
- 12 Review the package installation summary, and then press **Enter**.
- 13 After the system installation requirements check completes, press **Enter**.
- 14 Start the prerequisites check by pressing **Enter**.
- 15 Press **Enter** to begin the installation.
- 16 After a message appears stating that the installation has completed successfully, press **Enter**.
- 17 When the installer prompts to start the Agent for Mac, click **Yes**.
See [“Starting the Agent for Mac”](#) on page 1232.
- 18 Create the admin group if the installation did not create it automatically.
See [“Creating the Backup Exec admin group manually on Macintosh systems”](#) on page 1220.
- 19 Perform additional configuration as appropriate.
See [“About configuring the Agent for Mac”](#) on page 1226.

Uninstalling the Agent for Mac

An automated uninstall process for the Agent for Mac (Mac Agent) is available on the Backup Exec installation media.

You can also manually uninstall the Mac Agent.

See [“Manually uninstalling the Agent for Mac”](#) on page 1225.

The uninstall summary is saved to the following location on the Macintosh system:

`/var/tmp/vxif/uninstallrams<unique identifier number>.summary`

The uninstall log file is saved to the following location on the Macintosh system:

`/opt/VRTS/install/logs/uninstallrams<summary file number>.log`

After the log files are saved, the uninstall process is complete.

To uninstall the Agent for Mac

- 1 On a Macintosh system, place the Backup Exec installation media in the appropriate drive.
- 2 On the Macintosh system from which you want to uninstall the Mac Agent, log on using Admin privileges.
- 3 Navigate to the following directory on the Backup Exec installation media:
<LinuxMac>
- 4 Copy the RALUS_RMALS_RAMs_<version number>.gz file in this directory to a directory on the local system.

- 5 Unzip the file.

For example:

```
gunzip RALUS_RMALS_RAMs_<version number>.gz
```

- 6 Untar the file.

For example:

```
tar -xf RALUS_RMALS_RAMs_<version number>.tar
```

- 7 Open **Finder**, and then browse to **Applications > Utilities**.
- 8 Open **Terminal**.
- 9 Start the **uninstallrams** script.

For example:

```
./uninstallrams
```

10 Do one of the following:

To uninstall the Mac Agent from one system

Type the name, IP address, or fully qualified domain name of the Macintosh system.

To uninstall the Mac Agent from multiple systems

Type the names, IP addresses, or fully qualified domain names of the Macintosh systems. Leave a space between each identifier.

11 Press **Enter**.

12 After the Mac Agent package check completes successfully, press **Enter**.

13 When you are prompted to uninstall the RALUS packages, press **Enter**.

14 When the uninstall process is complete, press **Enter**.

Manually uninstalling the Agent for Mac

You can manually uninstall the Agent for Mac (Mac Agent) from Macintosh systems.

You can also use the Backup Exec installation media to uninstall the Mac Agent.

See [“Uninstalling the Agent for Mac”](#) on page 1223.

To manually uninstall the Agent for Mac

1 Use a logon account with Admin privileges to log on to a terminal session to connect to the Macintosh system.

2 Change to the following directory:

```
/opt/VRTSralus/bin
```

For example:

```
cd /opt/VRTSralus/bin
```

3 Delete the following line if it is found in the `/etc/inittab` file:

```
/opt/VRTSralus/bin/VRTSralus.init
```

For example:

```
rm -r /opt/VRTSralus/bin/VRTSralus.init
```

4 Stop the Mac Agent daemon.

See [“Stopping the Agent for Mac”](#) on page 1233.

5 Remove the Mac Agent package from the Linux server.

- 6 Change back to the root directory.

For example:

```
cd /
```

- 7 Remove the following files:

/etc/VRTSralus

/opt/VRTSralus

/var/VRTSralus

For example:

```
rm -r /etc/VRTSralus /opt/VRTSralus /var/VRTSralus
```

- 8 Type **y** if you are prompted to descend into the directories.

- 9 Type **y** if you are prompted to delete a directory.

- 10 Remove the /Library/StartupItems/VRTSrams folder.

For example:

```
rm -r /Library/StartupItems/VRTSrams
```

- 11 Type **y** if you are prompted to delete a directory.

About configuring the Agent for Mac

Backup Exec creates a file named `ralus.cfg` on each Macintosh system on which the Agent for Mac (Mac Agent) is installed.

You can edit the following strings, identifiers, and variables for the Mac Agent in the `ralus.cfg` file:

- The port to which the Mac Agent must send publishing messages.
- The settings to allow the Mac Agent to publish to one or more Backup Exec servers.
- The files and directories on Macintosh systems that you want to exclude from backups.

The `ralus.cfg` format contains three components. The first component (A) in the following example is a required string.

The second component (B) is a unique identifier followed by an equal sign (=). A unique identifier can consist of sequential numbers, letters, or alpha-numeric characters. For example, 1, 2, 3 or A, B, C. You can also use AA, BB, CC, or A1, A2, B1, B2.

The third component of the ralus.cfg format is the NetBIOS name, fully qualified domain name, or IP address of the Backup Exec server.

Figure P-1 Example of the ralus.cfg file

A	B	C
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Agent	Directory List 1=	srv.mycompany.com
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Agent	Directory List 2=	datasrv
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Agent	Directory List 3=	66.35.250.151

A = Required string

B = Required and unique identifier (the order or appearance is irrelevant)

C = File or directory to be excluded

See [“Editing the default backup options for Macintosh systems”](#) on page 1235.

See [“About the list of servers on the Backup and Restore tab”](#) on page 152.

See [“About excluding files and directories from backup jobs for Linux computers”](#) on page 1196.

Editing configuration options for Macintosh computers

You can edit configuration options for the Agent for Mac (Mac Agent).

See [“About excluding files and directories from backup jobs for Macintosh computers”](#) on page 1237.

See [“Configuration options for Macintosh computers”](#) on page 1227.

To edit configuration options for Macintosh computers

- 1 Use a text editor to open the following file:
/etc/VRTSralus/ralus.cfg
- 2 Change the appropriate string in the file.

Configuration options for Macintosh computers

You can edit options to configure the Agent for Mac (Mac Agent).

Table P-2

Item	Description
Software\Symantec\Backup Exec For Windows\Agent Browser\TcpIp\AdvertisementPort=6101	Lists the port to which the Mac Agent must send publish and purge messages.

Table P-2 (continued)

Item	Description
Software\Symantec\Backup Exec for Windows\Backup Exec\Debug\VBBSAlevel=0	<p>Enables logging for the Linux agent for Oracle operations.</p> <p>Values include the following:</p> <ul style="list-style-type: none">■ 0 Logging is not enabled.■ 5 Normal logging is enabled.■ 6 Advanced logging is enabled. Large log files may be created. <p>This option does not apply to the Agent for Mac.</p>
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\Agents\Advertise All=1	<p>Enables the Mac Agent to publish information to all of the Backup Exec servers that are listed in the \Agents\Agent Directory List strings.</p> <p>Values include the following:</p> <ul style="list-style-type: none">■ 1 The Mac Agent publishes information to every Backup Exec server in the Agent Directory List.■ 0 The Mac Agent publishes information to the first Backup Exec server in the Agent Directory List. If the attempt is successful, the Mac Agent does not publish information to any other Backup Exec servers. If the attempt is not successful, the Mac Agent attempts to publish information to the next Backup Exec server in the list. Attempts continue until the Mac Agent reaches the end of the list.

Table P-2 (continued)

Item	Description
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Advertise Now=0	<p>Enables the Mac Agent to start a new publishing cycle after you add or edit any settings in the <code>ralus.cfg</code> file.</p> <p>Values include the following:</p> <ul style="list-style-type: none">■ 0 The Mac Agent publishes information according to its regular cycle, set in the string <code>\Agents\Advertising Interval Minutes</code>. Any changes to the <code>ralus.cfg</code> file take effect when a new publishing cycle begins.■ 1 The Mac Agent starts a new publishing cycle. Any changes to the <code>ralus.cfg</code> file take effect immediately. If the Backup Exec server does not receive the publishing information, the Mac Agent makes 10 more attempts. Each attempt to publish information to the Backup Exec server is one minute apart. If the information is not sent at the end of the 10 attempts, the Mac Agent skips that Backup Exec server until the next publishing cycle. The publishing cycle is the number of minutes set in the string <code>\Agents\Advertising Interval Minutes</code>.

Table P-2 (continued)

Item	Description
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Advertisement Purge=0	<p>Lets the Mac Agent send a purge message to all of the Backup Exec servers in the string \Agents\Advertisement Purge. When a Backup Exec server receives a purge message, it removes the Mac Agent from Backup Exec's list of available servers. The Mac Agent continues to function.</p> <p>Values include the following:</p> <ul style="list-style-type: none">■ 0 Do not purge the Mac Agent from any Backup Exec servers that are listed in the \Agents\Advertisement Purge string.■ 1 Purge the Mac Agent from one or more Backup Exec servers in the \Agents\Advertisement Purge string.
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Advertising Disabled=0	<p>Enables the Mac Agent to publish to Backup Exec servers.</p> <p>Values include the following:</p> <ul style="list-style-type: none">■ 0 The Mac Agent attempts to publish information to the Backup Exec servers that are listed in the string \Agents\Agent Directory List.■ 1 The Mac Agent does not publish information to Backup Exec servers.
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Advertising Interval Minutes=240	<p>Sets the number of minutes that the Mac Agent must wait between publishing cycles. The default number of minutes is 240.</p> <p>The range of minutes is from 1 minute to 720 minutes.</p>

Table P-2 (continued)

Item	Description
Software\Symantec\Backup Exec for Windows\Backup Exec\Agents\Agent Directory List 1=<Backup Exec server name>	<p>Displays the list of NetBIOS names, fully qualified domain names, or IP addresses to which the Mac Agent publishes information.</p> <p>The Backup Exec server from which the Mac Agent is push-installed is added to the Agent Directory List by default.</p>
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\Agents\Auto Discovery Enabled=1	<p>Adds a Backup Exec server to the string \Agents\Agent Directory List if the Backup Exec server performs a backup job with which the Mac Agent is associated.</p> <p>Values include the following:</p> <ul style="list-style-type: none">■ 1 Adds the Backup Exec server that performs the backup job to the Agent Directory List. The Mac Agent can publish information to the Backup Exec server.■ 0 The Backup Exec server that performs the backup job is not added to the Agent Directory List.
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\Logging\RANT NDMP Debug Level=0	<p>Displays the level of verbosity for logging NDMP information for the Mac Agent.</p> <p>Values include the following:</p> <ul style="list-style-type: none">■ 0 Logs only the NDMP errors.■ 1 Logs the NDMP errors and warnings.■ 2 Logs the NDMP errors, warnings, and message information that is sent between the remote computer and the Backup Exec server.
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RALUS\Encoder=	<p>Displays the encoder that you can add if the default encoder incorrectly displays characters on the user interface.</p>

Table P-2 (continued)

Item	Description
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RALUS\ShowTSAFS=	<p>Lets you perform a Target Service Agent file system (TSAFS) backup for applications on Novell Open Enterprise Services. By default, this option is not enabled.</p> <p>The Mac Agent backs up all file systems using the Root object. If ShowTSAFS is enabled, the Novell Open Enterprise Services resource appears in the backup selection list. If you select the whole computer for backup, then redundant backups are performed. Symantec recommends that you do not enable this option.</p> <p>Values include the following:</p> <ul style="list-style-type: none">■ Blank or 0 The file system TSA does not appear for backup selection.■ 1 The file system TSA resource appears for backup selection. <p>This option does not apply to the Agent for Mac.</p>
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RALUS\SystemExclude1=	<p>Lists the files that you want to exclude from all Mac Agent backup jobs.</p> <p>See “About excluding files and directories from backup jobs for Macintosh computers” on page 1237.</p>
Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RALUS\SystemFSTypeExclude1	<p>Lists the type of file system that you want to exclude from the Mac Agent backup.</p>
Software\Symantec\Backup Exec for Windows\Backup Exec\ Engine\RALUS\vmPath= /opt/VRTSralus/VRTSvxms	<p>Displays the path to the Veritas Mapping Service libraries that the Mac Agent uses.</p>

Starting the Agent for Mac

You can manually start the Agent for Mac (Mac Agent).

See [“Stopping the Agent for Mac”](#) on page 1233.

To manually start the Agent for Mac

- 1 Use a terminal session to connect to the target Macintosh system as the root user.
- 2 From the root prompt, start the VRTSrams service.

For example:

```
SystemStarter start VRTSrams
```

Stopping the Agent for Mac

You can manually stop the Agent for Mac (Mac Agent).

See [“Starting the Agent for Mac”](#) on page 1232.

To manually stop the Agent for Mac

- 1 Use a terminal session to connect to the target Macintosh system as the root user.
- 2 From the root prompt, stop the VRTSrams service:

For example:

```
SystemStarter stop VRTSrams
```

About establishing trust for a remote Macintosh system in the Servers list

When you connect to a Macintosh system from the Backup Exec server, you must establish trust between the Backup Exec server and the remote Macintosh computer.

See [“Establishing a trust between the Backup Exec server and a remote computer”](#) on page 780.

See [“Establishing trust and adding a remote Macintosh computer to the Backup Exec Servers list”](#) on page 1233.

Establishing trust and adding a remote Macintosh computer to the Backup Exec Servers list

You can add one or more remote Macintosh computers to the list of servers on the **Backup and Restore** tab. When you add remote Macintosh computers, you must

establish a trust between the Backup Exec server and the remote Macintosh computers to ensure secure communication.

See [“About establishing trust for a remote Macintosh system in the Servers list”](#) on page 1233.

To establish trust and add a remote Macintosh computer to the Backup Exec Servers list

- 1 On the **Backup and Restore** tab, in the **Servers** group, click **Add**.
- 2 Click **Macintosh computer**.
- 3 Follow the on-screen prompts.

See [“About the list of servers on the Backup and Restore tab”](#) on page 152.

See [“Adding additional Backup Exec servers to which the Agent for Mac can publish information”](#) on page 1234.

Adding additional Backup Exec servers to which the Agent for Mac can publish information

You can specify additional Backup Exec servers to which the Agent for Mac (Mac Agent) can publish information.

Each Backup Exec server to which the Mac Agent publishes information appears in the **Backup and Restore** tab, under the **Servers** list.

To add additional Backup Exec servers to which the Agent for Mac can publish information

- 1 Use a text editor to open the following file:
`/etc/VRTSralus/ralus.cfg`
- 2 Add the following string:
`Software\Symantec\Backup Exec For Windows\Backup
Exec\Engine\Agents\Agent Directory List unique identifier number = IP address
or DNS name of Backup Exec server`
- 3 Save and close the file.
- 4 Move to the Backup Exec server to which the Mac Agent publishes itself and add the Macintosh computer to the **Servers** list.

See [“About the list of servers on the Backup and Restore tab”](#) on page 152.

About backing up data by using the Agent for Mac

When you use the Agent for Mac (Mac Agent) to back up data, only the following backup methods are supported for Macintosh systems:

- Full - Using modified time
- Differential - Using modified time
- Incremental - Using modified time

See [“Backing up data”](#) on page 158.

See [“Adding a stage to a backup definition”](#) on page 219.

See [“Editing a stage”](#) on page 221.

Editing the default backup options for Macintosh systems

You can use the existing defaults for all backup jobs for Macintosh systems, or you can edit the defaults.

To edit default backup options for Macintosh computers

- 1 Click the Symantec Backup Exec button, select **Configuration and Settings**, and then select **Backup Job Defaults**.
- 2 Select the **Back Up to Disk** option.
- 3 In the left pane, click **Linux and Macintosh**.
- 4 Select the default backup options for the Agent for Mac (Mac Agent).
- 5 Set the appropriate options.
See [“Default backup job options for Macintosh systems”](#) on page 1235.
- 6 Click **OK**.

Default backup job options for Macintosh systems

When you back up a Macintosh system, you should consider the following options.

See [“Editing the default backup options for Macintosh systems”](#) on page 1235.

Table P-3 Default backup job options for Macintosh systems

Item	Description
Preserve file and directory timestamps during backups	<p>Prevents the Agent for Mac (Mac Agent) from changing an object's attributes when a backup occurs. An object is a file or a directory.</p> <p>This option is not selected by default.</p> <p>During a backup, Backup Exec preserves an object's last access timestamp by resetting the last access timestamp to the value before the backup occurred. When Backup Exec modifies the object's last access timestamp, the operating system internally updates the object's ctime.</p> <p>An object's ctime is the time when an object's attributes, such as permissions and timestamps, have been modified. If the Mac Agent does not change the attributes after a backup, the object's ctime is not changed.</p> <p>This option does not affect the attributes of the object that are set during restore operations.</p>
Follow local mount points	<p>Lets Backup Exec follow local mount points to back up data.</p> <p>This option is enabled by default.</p> <p>For more information on local mount points, see your operating system's documentation.</p>

Table P-3 Default backup job options for Macintosh systems (*continued*)

Item	Description
Follow remote mount points	<p>Lets Backup Exec follow remote mount points to back up data.</p> <p>This option is not selected by default.</p> <p>When you use this option, the following limitations apply:</p> <ul style="list-style-type: none">■ The data that is mounted must reside on a system that Backup Exec supports. You can find a list of supported operating systems, platforms, and applications at the following URL: http://entsupport.symantec.com/umi/v-269-1■ If the mount point leads to an operating system that Backup Exec does not support, contact the operating system's vendor to resolve any issues. <p>For more information on remote mount points, see your operating system's documentation.</p>
Lock remote files to prevent applications from modifying them during backups	<p>Lets the Mac Agent have exclusive access to the files on the remote servers that are connected through Network File System (NFS).</p> <p>Locking remote files prevents other applications from modifying the data during the backup or restore job.</p>

About excluding files and directories from backup jobs for Macintosh computers

You can exclude specific files and directories on Macintosh computers from all backup jobs. Edit the `ralus.cfg` file to specify the excluded files.

The following is an example of strings in the `ralus.cfg` file that excludes files and directories from all backup jobs.

Figure P-2 Example of file and directory exclusions in the ralus.cfg format

	A	B	C
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\RALUS\SystemExclude1=	/dev/*.*		
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\RALUS\SystemExclude2=	/proc/*.*		
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\RALUS\SystemExclude3=	/mnt/nss/pools/		
Software\Symantec\Backup Exec For Windows\Backup Exec\Engine\RALUS\SystemExclude4=	/mnt/nss/.pools/		

A = Required string
B = Required and unique identifier (the order or appearance is irrelevant)
C = File or directory to be excluded

To exclude files and directories for specific backup jobs, specify the exclusions in the backup job properties.

See [“Editing configuration options for Macintosh computers”](#) on page 1227.

About restoring Macintosh systems

You can specify restore job properties to restore Macintosh systems.

See [“Macintosh restore options”](#) on page 1238.

Macintosh restore options

When you restore Macintosh systems, you can enable the option to **Lock remote files**. This option allows exclusive access to the files on the remote systems that are connected through the Network File System (NFS). This option is enabled by default.

See [“About restoring Macintosh systems”](#) on page 1238.

Troubleshooting the Agent for Mac

If you experience problems with the Agent for Mac (Mac Agent), read the following questions and answers.

Table P-4 Troubleshooting the Agent for Mac

Question	Answer
The Mac Agent is installed on a Macintosh system in a NIS domain, but Backup Exec is unable to browse resources on the system. What do I do?	<p>If the group line and the password line in the nsswitch.conf file are set to compatibility mode, additional configuration is necessary. Refer to the nsswitch.conf man pages for additional information on configuring nsswitch.conf to use compatibility mode.</p> <p>Alternately, change the password line and the group line to NIS files so that the Macintosh system validates the user through NIS. If the NIS server is unavailable or the user is not found, the local files are used for validation.</p>
<p>I cannot load the Mac Agent. When I attempt to load the Mac Agent in console mode, <code>"/.beremote --log-console"</code> shows the following message:</p> <p>"ACE_SV_Semaphore_Complex: no space left on device."</p> <p>What should I do?</p>	<p>This issue occurs when the computer reaches its maximum limit on allowable semaphores. It can occur after an unexpected termination of the Mac Agent. When the Mac Agent unexpectedly terminates, it is unable to clean up some of the semaphore resources that it used. Other processes may have caused the use of semaphores to reach the limit. You must restart the computer to safely recover it from this condition.</p> <p>If other processes are running, it may not be feasible to restart the computer. Instead, you can use the commands that let you list and then remove all semaphores in use by the operating system. Be careful when you select semaphores to remove. Semaphores that are in use by the Mac Agent cannot be identified. If you remove semaphores of other programs that are in use, those programs can become unstable.</p>

Symantec Backup Exec Remote Media Agent for Linux

This appendix includes the following topics:

- [About the Remote Media Agent for Linux](#)
- [How the Remote Media Agent for Linux works](#)
- [Requirements for the Remote Media Agent for Linux](#)
- [About open files and the Remote Media Agent for Linux](#)
- [About installing the Remote Media Agent for Linux](#)
- [Uninstalling the Remote Media Agent for Linux](#)
- [Starting the Remote Media Agent for Linux daemon](#)
- [Stopping the Remote Media Agent for Linux daemon](#)
- [About establishing trust for a Remote Media Agent for Linux computer in the Backup Exec list of servers](#)
- [Adding additional Backup Exec servers to which Remote Media Agent for Linux can publish](#)
- [Finding simulated tape library files](#)
- [About the Backup Exec operators group for the Remote Media Agent for Linux](#)
- [About adding a Linux server as a Remote Media Agent for Linux](#)

- [Changing the port for communications between the Backup Exec server and the Remote Media Agent for Linux](#)
- [About creating storage device pools for devices attached to the Remote Media Agent for Linux](#)
- [Editing properties for the Remote Media Agent for Linux](#)
- [Deleting a Remote Media Agent for Linux from the Backup Exec list of servers](#)
- [Sharing a Remote Media Agent for Linux between multiple Backup Exec servers](#)
- [About backing up data by using the Remote Media Agent for Linux](#)
- [About restoring data by using the Remote Media Agent for Linux](#)
- [About the Tape Library Simulator Utility](#)
- [Creating a simulated tape library](#)
- [Viewing simulated tape libraries properties](#)
- [Deleting a simulated tape library](#)
- [Managing simulated tape libraries from the command line](#)
- [Command line switches for the Tape Library Simulator Utility](#)
- [Troubleshooting the Remote Media Agent for Linux](#)

About the Remote Media Agent for Linux

The Remote Media Agent for Linux (RMAL) lets you back up data from remote computers to the following devices:

- The storage devices that are directly attached to a Linux server.
- A simulated tape library on a Linux server.

You can add a Linux server to a Backup Exec server as a RMAL. You can then back up data from the Linux server or from supported remote computers to the devices that are attached to the Linux server. You can also create a virtual device on a server on which RMAL is installed. This virtual device emulates a SCSI tape library.

See [“How the Remote Media Agent for Linux works”](#) on page 1242.

See [“About the Tape Library Simulator Utility”](#) on page 1258.

How the Remote Media Agent for Linux works

From the Backup Exec server, you can add a Linux server as a Remote Media Agent for Linux (RMAL). RMAL establishes a data connection to the remote computer on which a supported agent is installed. You can then create backup, restore, and utility jobs on the Backup Exec server that run on the Linux server's storage devices.

If you use Backup Exec Central Admin Server Option, you can share a RMAL computer between multiple Backup Exec servers. Sharing can be enabled when you add RMAL. You can select new Backup Exec servers to share RMAL or remove the sharing ability from the Backup Exec servers at any time.

See [“Sharing storage devices”](#) on page 448.

Job performance increases because data travels from the remote computers to the devices that are attached to the Linux server. This increase is especially apparent if the Backup Exec server is located at a different site than the RMAL computer and the remote computers.

RMAL does not have a user interface. You use the administration console on the Backup Exec server to manage the jobs and devices on RMAL. The Backup Exec server maintains job logs, catalogs, job histories, alerts, and notifications.

See [“Requirements for the Remote Media Agent for Linux”](#) on page 1242.

See [“About installing the Remote Media Agent for Linux”](#) on page 1243.

See [“About adding a Linux server as a Remote Media Agent for Linux”](#) on page 1252.

See [“About the Tape Library Simulator Utility”](#) on page 1258.

Requirements for the Remote Media Agent for Linux

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

You can find a list of compatible devices at the following URL:

<http://entsupport.symantec.com/umi/V-269-2>

You must have superuser privileges on the Linux servers before you install RMAL.

Note: RMAL does not support the Backup Exec File System Archiving Option or the Exchange Mailbox Archiving Option.

Symantec recommends that you use the Secure Shell (SSH) protocol when you push-install RMAL to remote servers. You must enable SSH before you install RMAL.

Note: Some versions of Linux may require that you install the libstdc++.so.5 package.

See [“About installing the Remote Media Agent for Linux”](#) on page 1243.

About open files and the Remote Media Agent for Linux

The Remote Media Agent for Linux (RMAL) uses advanced open file and image technologies that are designed to alleviate issues that are sometimes encountered during backup operations, such as backing up open files.

After you make file and folder selections and the job is submitted for backup, RMAL automatically makes a snapshot of the volume or volumes. Making snapshot of a volume provides a point-in-time record of the data. When it creates a snapshot, RMAL uses snapshot technologies to momentarily suspend write activity to a volume so that a snapshot of the volume can be created. During the backup, files can be open and data can be changed.

RMAL supports Simple, Logical Volume Manager (LVM), and RAID volume configurations.

See [“Requirements for the Remote Media Agent for Linux”](#) on page 1242.

About installing the Remote Media Agent for Linux

Use the Backup Exec installation media to do the following:

- Install Remote Media Agent for Linux (RMAL) on a local Linux server.
- Push-install RMAL to one or more remote Linux servers.
If you push-install RMAL, the RSH (Remote Shell) is used by default. Symantec recommends that you use SSH (Secure Shell) instead. To use SSH, you must enable it before you install RMAL. Refer to your operating system documentation for more information about SSH.

When you install RMAL, Backup Exec creates the beoper group and adds root as a member. Any Linux user that you add to the beoper group gets the permissions necessary to back up and restore the Linux servers.

However, if Backup Exec detects an NIS server during the RMAL installation, then the beeper group is not created. You must create the beeper group manually on the Linux servers.

After the installation completes, you must add the Linux server as a Remote Media Agent on the Backup Exec server. Then, you can send jobs to the devices that are attached to the Linux server.

See [“Installing the Remote Media Agent for Linux”](#) on page 1244.

See [“Creating the Backup Exec operators group manually for the Remote Media Agent for Linux”](#) on page 1251.

See [“About adding a Linux server as a Remote Media Agent for Linux”](#) on page 1252.

See [“About the Backup Exec operators group for the Remote Media Agent for Linux”](#) on page 1251.

Installing the Remote Media Agent for Linux

You can install Remote Media Agent for Linux (RMAL) on a local Linux server or push-install it to one or more remote Linux servers.

See [“About installing the Remote Media Agent for Linux”](#) on page 1243.

Note: You must unzip the RALUS_RMALS_RAMs_<version number>.gz file on a Linux server. The installation does not run if it is unzipped on a computer that runs the Windows operating system.

To install the Remote Media Agent for Linux

- 1 At a Linux server, place the Backup Exec installation media in the appropriate drive.
- 2 Log on as root on the server on which you want to install RMAL.
- 3 Navigate to the following path on the installation media:
`<LinuxMac>`
- 4 Copy the RALUS_RMALS_RAMs_<version number>.gz file in this directory to a directory on the local server.
- 5 Unzip the file.

For example:

```
gunzip RALUS_RMALS_RAMs_<version number>.gz
```


6 Untar the file.

For example:

```
tar -xf RALUS_RMALS_RAMs_<version number>.tar
```

7 Start the **installrmal** script.

For example:

```
./installrmal
```

8 Do one of the following:

To install on a local server

Press **Enter**.

To install to one remote server

Type the name, IP address, or fully qualified domain name of a Linux server.

To install to multiple remote servers

Type the names, IP addresses, or fully qualified domain names of the Linux servers. Leave a space between each identifier.

9 After the installer checks for a valid Linux operating system during the initial system check, press **Enter**.

10 Review the package installation summary, and then press **Enter**.

11 After the system installation requirements check completes, press **Enter**.

12 Start the prerequisites check by pressing **Enter**.

13 Type the name, IP address, or fully qualified domain name of the Backup Exec server (directory host) that you want to use this Remote Media Agent.

14 Type any additional names, IP addresses, or fully qualified domain names of Backup Exec servers that you want to use this Remote Media Agent.

15 Do one of the following:

If the server name, IP address, or fully qualified domain name is correct

Press **Enter** to continue the installation.

If you want to change a server name, IP address, or fully qualified domain name

Type **N**, press **Enter**, and then change the information.

16 Start the NIS server scan by pressing **Enter**.

17 Examine the results of the NIS server scan, and then do one of the following:

If an NIS server is detected	<p>The RMAL installer cannot create the beoper group. You must create it manually after RMAL installation is complete.</p> <p>Continue with the next step.</p>
If an NIS server is not detected	<p>Use the installer to create the beoper group.</p> <p>Do the following in the order listed:</p> <ul style="list-style-type: none">■ To let the installer create the beoper group, type y.■ To select the next available Group ID, type n.■ To add the root user account to the beoper group, type y.■ Continue with the next step.

- 18 Start the installation by pressing **Enter**.
- 19 After installation completes, press **Enter** to start the post-installation configurations and installation of SymSnap drivers.
- 20 Press **Y** to automatically start the Beremote service; otherwise, press **N** to start the service later
- 21 After the configuration process completes, press **Enter** to save the installation log to the following file:

`/var/tmp/vxif/installrmalsummary file number/installrmal.log`
- 22 If the RMAL installer did not create a beoper group, you must create it.

See [“Creating the Backup Exec operators group manually for the Remote Media Agent for Linux”](#) on page 1251.
- 23 Start the Agent for Linux daemon.

See [“Starting the Agent for Linux daemon”](#) on page 1214.
- 24 Add the Linux server as a Remote Media Agent.

See [“About adding a Linux server as a Remote Media Agent for Linux ”](#) on page 1252.

Uninstalling the Remote Media Agent for Linux

Before you uninstall the Remote Media Agent for Linux (RMAL), you should note the location of the simulated tape library files. Then, you can delete all of the

simulated tape library files after the uninstall operation completes. When you delete these files, you delete the backup data that you stored on the Linux server.

See [“Finding simulated tape library files”](#) on page 1250.

Note: You must have the Backup Exec installation media to uninstall RMAL.

To uninstall the Remote Media Agent for Linux

- 1 On the Linux server, place the Backup Exec installation media in the appropriate device.
- 2 Log on as root to the server from which you want to uninstall RMAL.
- 3 Navigate to the following path on the installation media:

 <LinuxMac>
- 4 Start the **uninstallrmal** script.

 For example:

 ./uninstallrmal
- 5 Do one of the following:

To uninstall RMAL from one server	Type the name, IP address, or fully qualified domain name of the Linux server.
To uninstall RMAL from multiple servers	Type the names, IP addresses, or the fully qualified domain names of the Linux servers. Leave a space between each identifier.
- 6 Press **Enter**.
- 7 After the RMAL package check completes successfully, press **Enter**
- 8 When you are prompted to uninstall the RMAL packages, press **Enter** to save the uninstall summary and log to the following location:

 /var/tmp/vxif/uninstallrmalsummary file number.log
- 9 Manually delete the simulated tape library files.

Starting the Remote Media Agent for Linux daemon

If necessary, you can start the Remote Media Agent for Linux (RMAL) daemon after the operating system starts.

See [“Stopping the Remote Media Agent for Linux daemon”](#) on page 1248.

To start the Remote Media Agent for Linux daemon

- 1 Use a terminal session to connect to the Linux server as the root user.
- 2 Navigate to the following directory:

/etc/init.d/

For example:

```
cd /etc/init.d/
```

- 3 Start the RMAL daemon.

For example:

```
/etc/init.d/VRTSralus.init start
```

Stopping the Remote Media Agent for Linux daemon

You can stop the Remote Media Agent for Linux (RMAL) daemon.

See [“Starting the Remote Media Agent for Linux daemon”](#) on page 1247.

To stop the Remote Media Agent for Linux daemon

- 1 Use a terminal session to connect to the Linux server as the root user.
- 2 Navigate to the following directory:

/etc/init.d/

For example:

```
cd /etc/init.d/
```

- 3 Stop the RMAL daemon.

For example:

```
/etc/init.d/VRTSralus.init stop
```

- 4 Restart the RMAL daemon when necessary.

About establishing trust for a Remote Media Agent for Linux computer in the Backup Exec list of servers

When you connect to a Remote Media Agent for Linux (RMAL) computer from the Backup Exec server, you must establish trust between the Backup Exec server and the RMAL computer. You must also establish trust if you want to configure a remote RMAL computer to perform client-side deduplication.

See [“Establishing a trust between the Backup Exec server and a remote computer”](#) on page 780.

See [“Establishing trust and adding a Remote Media Agent for Linux computer to the Backup Exec list of servers”](#) on page 1249.

Establishing trust and adding a Remote Media Agent for Linux computer to the Backup Exec list of servers

You can add one or more Remote Media Agent for Linux (RMAL) computers to the Backup Exec list of servers. When you add RMAL, you must establish a trust between the Backup Exec server and the remote Linux computers to ensure secure communication.

To establish trust and add a Remote Media Agent for Linux computer to the Backup Exec list of servers

- 1 On the **Storage** tab, click **Configure Storage**, and then select **Network storage**.
- 2 Click **Next**.
- 3 Select **Backup Exec Remote Media Agent for Linux**, and then click **Next**.
- 4 Follow the on-screen prompts.

See [“About establishing trust for a Remote Media Agent for Linux computer in the Backup Exec list of servers”](#) on page 1248.

Adding additional Backup Exec servers to which Remote Media Agent for Linux can publish

You can specify additional Backup Exec servers to which the Remote Media Agent for Linux (RMAL) can publish information.

Each Backup Exec server to which RMAL publishes information appears in the Backup Exec list of servers.

To add additional Backup Exec servers to which the Remote Media Agent for Linux can publish information

- 1 Use a text editor to open the following file:
`/etc/VRTSralus/ralus.cfg`
- 2 Add the following string:
`Software\Symantec\Backup Exec For Windows\Backup
Exec\Engine\Agents\Agent Directory List unique identifier number = IP address
or DNS name of Backup Exec server`
- 3 Save and close the file.
- 4 Move to the Backup Exec server to which the Remote Media Agent for Linux is publishing itself and add the RMAIL server to the **Servers** list.

See [“About the list of servers on the Backup and Restore tab”](#) on page 152.

Finding simulated tape library files

Before you uninstall RMAIL, you should note the location of the simulated tape library files. Then, after you uninstall RMAIL, you can delete all of the simulated tape library files. When you delete these files, you delete the backup data that you stored on the Linux server.

See [“Uninstalling the Remote Media Agent for Linux”](#) on page 1246.

See [“About the Tape Library Simulator Utility”](#) on page 1258.

To find simulated tape library files

- 1 Log on as root to the server on which you want to find the simulated tape library files.
- 2 Navigate to the following directory that contains the Tape Library Simulator:
`/opt/VRTSralus/bin`
For example:
`cd /opt/VRTSRALUS/bin`
- 3 Start the **mktls** utility to list the simulated tape library files and folders.
For example:
`/opt/VRTSralus/bin/mktls -l`
- 4 Write down the locations of the directories for the simulated tape library files.

About the Backup Exec operators group for the Remote Media Agent for Linux

The Backup Exec operators (**beoper**) group contains the names of the users who have permission to back up and restore the Linux servers.

When you install the Remote Media Agent for Linux (RMAL), Backup Exec creates the **beoper** group and adds root as a member. Any Linux user that you add to the beoper group gets the necessary permission to back up and restore the Linux servers.

However, if an NIS server is detected during RMAL installation, Backup Exec cannot create the **beoper** group. You must create the **beoper** group manually on the Linux servers on which you want to install RMAL. You must create the **beoper** group before you start backup and restore operations. Otherwise, connections fail between the Linux servers and the Backup Exec server.

Before the members of the **beoper** group can perform backup or restore operations, they must have a Backup Exec logon account.

See [“Creating the Backup Exec operators group manually for the Remote Media Agent for Linux”](#) on page 1251.

See [“Backup Exec logon accounts”](#) on page 610.

Creating the Backup Exec operators group manually for the Remote Media Agent for Linux

If the Remote Media Agent for Linux (RMAL) installation detects an NIS server during installation, you must create a beoper group on each Linux server on which you install RMAL.

See [“About the Backup Exec operators group for the Remote Media Agent for Linux”](#) on page 1251.

Note: Ensure that you understand how to set security for groups on Linux servers before you assign a Group ID for the beoper group.

Table Q-1 How to manually create the beoper group

Step	Action	More Information
Step 1	Navigate to the Linux server on which you want to install RMAL. If the Linux server is in a NIS domain, navigate to the NIS domain's group file.	See the NIS documentation for information on how to add a group to a NIS domain group file.
Step 2	Create a group with the following case-sensitive name: beoper	See the operating system documentation for more information about how to create a group.
Step 3	In the beoper group, add the users that you want to have permission to back up and restore the Linux server.	See the operating system documentation for more information about how to add users to a group
Step 4	Create a Backup Exec logon account for each user that you add to the beoper group.	See "Backup Exec logon accounts" on page 610.

About adding a Linux server as a Remote Media Agent for Linux

After you add the Linux server as a Remote Media Agent for Linux (RMAL), and if you purchased and installed the Central Admin Server Option, you can share the storage devices that are connected to the RMAL computer with other Backup Exec servers.

See ["Adding a Linux server as a Remote Media Agent for Linux"](#) on page 1252.

See ["Sharing storage devices"](#) on page 448.

See ["About the Central Admin Server Option"](#) on page 1096.

Adding a Linux server as a Remote Media Agent for Linux

Use the following steps to add a Linux server as a Remote Media Agent for Linux (RMAL).

See ["About adding a Linux server as a Remote Media Agent for Linux"](#) on page 1252.

To add a Linux server as a Remote Media Agent for Linux

- 1 On the **Storage** tab, click **Configure Storage**.
 - 2 Select **Network storage**, and then click **Next**.
 - 3 Select **Backup Exec Remote Media Agent for Linux**, and then click **Next**.
 - 4 Follow the on-screen prompts to trust the RMAL and to restart the services.
- See [“Sharing storage devices”](#) on page 448.

Remote Media Agent for Linux options

You must provide information when you add a Linux server as a Remote Media Agent for Linux (RMAL) to a Backup Exec server.

See [“About adding a Linux server as a Remote Media Agent for Linux”](#) on page 1252.

Table Q-2 Add Remote Media Agent for Linux options

Item	Description
Server name	<p>Specifies the name of the Linux server that you want to add as a Remote Media Agent for Linux.</p> <p>If the Backup Exec Central Admin Server Option is installed in your environment, use the host name or fully qualified domain name of the Linux server. That is, use the name of the Linux computer that appears when you browse for backup selections. If you use the IP address, Backup Exec cannot distinguish which device path to use for jobs.</p>
Port number	<p>Lists the port to use for communications between the Backup Exec server and RMAL. If you change the port number, you must edit the services file in the /etc directory on the Linux server, and update the NDMP entry.</p> <p>See “Changing the port for communications between the Backup Exec server and the Remote Media Agent for Linux” on page 1254.</p> <p>Ensure that this port is open in any firewalls that exist between RMAL and the Backup Exec server. Use a port number that is not in use by another application or service.</p> <p>The default port is 10000.</p>

Table Q-2 Add Remote Media Agent for Linux options (continued)

Item	Description
Description	Displays a description that you choose.
Logon Account	Indicates the logon account for RMAL. The default logon account is the system logon account for the Backup Exec server.
Have Backup Exec use ICMP ping operations to detect the server	Lets the Backup Exec server use ICMP ping operations to locate the Linux server. You can turn off this option in environments where ping requests are blocked. This option is enabled by default.
Logon account	Indicates the Backup Exec logon account that you want to use to log on this server. See “Backup Exec logon accounts” on page 610.

See [“About creating storage device pools for devices attached to the Remote Media Agent for Linux ”](#) on page 1255.

See [“About backing up data by using the Remote Media Agent for Linux”](#) on page 1257.

Changing the port for communications between the Backup Exec server and the Remote Media Agent for Linux

You can change the port that Backup Exec uses to communicate with the Remote Media Agent for Linux (RMAL).

To change the port for communications between the Backup Exec server and the Remote Media Agent for Linux

- 1 On the computer on which RMAL is installed, use a text editor to open the services file in the /etc directory.

For example:

```
vi /etc/services
```

- 2 Search the file for an entry that is similar to the following:

ndmp 10000/tcp

3 Do one of the following:

If this entry exists

Change the port number to the port number that you want to use.

If this entry does not exist

Do the following in the order listed:

- At the end of the file, type `ndmp`, and then press **Tab**.
- Type the port number that you want NDMP to use, and then type `/tcp`.
- Press **Enter**.

4 Save the file, and then exit the editor.

5 Restart the Agent for Linux daemon.

See [“Starting the Remote Media Agent for Linux daemon”](#) on page 1247.

About creating storage device pools for devices attached to the Remote Media Agent for Linux

Remote Media Agents may reside in different physical locations. To reduce network traffic and increase job performance, you can create separate storage device pools for Remote Media Agents that are located at different sites.

See [“Creating storage device pools”](#) on page 416.

See [“About the Tape Library Simulator Utility”](#) on page 1258.

Editing properties for the Remote Media Agent for Linux

You can edit the properties for a Remote Media Agent for Linux (RMAL).

To edit properties for the Remote Media Agent for Linux

- 1 On the **Storage** tab, right-click a RMAL server.
- 2 Click **Details**.

See [“Remote Media Agent for Linux properties”](#) on page 1255.

Remote Media Agent for Linux properties

You can view properties for a Remote Media Agent for Linux server (RMAL).

See “Editing properties for the Remote Media Agent for Linux ” on page 1255.

The following table lists Remote Media Agent for Linux server properties:

Table Q-3 Remote Media Agent for Linux properties

Item	Description
Name	Displays the name, IP address, or fully qualified domain name of RMAL.
Port	Displays the port that is used for communications between the Backup Exec server and RMAL.
Backup Exec server status	Displays the status of the Backup Exec server. Backup Exec server status includes Online, Pause, Unavailable, and Offline.
Description	Displays a description of RMAL. You can edit this description.
Enable ICMP ping operations for Backup Exec to detect the Remote Media Agent	Lets Backup Exec communicate with RMAL. You can turn off this option in environments where ping requests are blocked. This option is enabled by default.
Host ID	Displays the identifier number that RMAL generates.
System version	Displays the version of the operating system that runs on RMAL.
Logon account	Indicates the logon account for RMAL. Click Change to select or create another logon account.

Deleting a Remote Media Agent for Linux from the Backup Exec list of servers

Use the following to delete a Remote Media Agent for Linux (RMAL) from the Backup Exec list of servers.

To delete a Remote Media Agent for Linux from a Backup Exec

- 1 On the **Storage** tab, right-click a RMAL.
- 2 Click **Yes**.

See [“Establishing trust and adding a Remote Media Agent for Linux computer to the Backup Exec list of servers”](#) on page 1249.

Sharing a Remote Media Agent for Linux between multiple Backup Exec servers

If the Central Admin Server Option is installed, you can select Backup Exec servers to share a Remote Media Agent for Linux (RMAL) server. When you add an RMAL server, the Backup Exec server that you used to add the device is automatically selected for sharing.

See [“Sharing storage devices”](#) on page 448.

To share a Remote Media Agent for Linux between multiple Backup Exec servers

- 1 On the **Storage** tab, under **All Storage**, right-click a RMAL server that you want the Backup Exec servers to access.
- 2 Select **Share**.
- 3 Under **Server**, select the Backup Exec servers that you want to use with RMAL.
- 4 Click **OK**.
- 5 Restart the Backup Exec services on the Backup Exec servers that you selected in step 3.

About backing up data by using the Remote Media Agent for Linux

Create a backup job for Remote Media Agent for Linux (RMAL) from the Backup Exec server.

See [“Backing up data”](#) on page 158.

See [“Editing backup definitions”](#) on page 206.

See [“Adding a stage to a backup definition”](#) on page 219.

About restoring data by using the Remote Media Agent for Linux

Create a restore job for RMAL from the Backup Exec server.

Note: Use devices that are attached to the Backup Exec server to restore data from the tapes that other applications created. RMAL supports only Microsoft Tape Format (MTF) media.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

About the Tape Library Simulator Utility

The Tape Library Simulator Utility lets you create a virtual device on a hard disk or on any mounted volume on a Linux server. This virtual device emulates a SCSI tape library. The Remote Media Agent for Linux (RMAL) must be installed on the server.

When you run the Tape Library Simulator Utility, you are prompted for the following information:

- The number of slots that you want to allocate to this library.
- The location or path for the library.

The Tape Library Simulator Utility then creates the media for the simulated tape library. To ensure that each media has a unique name, the Tape Library Simulator Utility creates a bar code label for each media. You cannot rename these bar code labels. However, you can add a unique media description.

The simulated tape library emulates an Advanced Intelligent Tape (AIT) media type. This media type is seldom used, so it helps you distinguish between a physical robotic library and a simulated tape library. The simulated media also has an AIT media type label.

The format of the files that are written to the simulated tape library is similar to the file format of backup-to-disk files. However, you cannot copy or move files between simulated tape libraries and backup-to-disk folders.

You can add the simulated tape library to Backup Exec device pools.

See [“Creating storage device pools”](#) on page 416.

To use the Tape Library Simulator Utility, you must have a minimum of 500 MB of available space on the Linux server. The available space includes hard disk space, flash drives, and USB drives. If there is not enough space, the jobs fail with an end-of-media error. You must either create available disk space or you must direct the jobs to another volume, and then start the jobs again.

A simulated tape library does not support all of the tasks that are available for physical robotic libraries.

See [“About storage operation jobs”](#) on page 424.

See [“Creating a simulated tape library”](#) on page 1259.

Creating a simulated tape library

Create a simulated tape library on a server on which RMA is installed. You must create the simulated tape library on a hard disk or on a mounted volume.

See [“About the Tape Library Simulator Utility”](#) on page 1258.

To create a simulated tape library

- 1 At the RMA computer, stop the Agent for Linux daemon.
See [“Stopping the Agent for Linux daemon”](#) on page 1215.
- 2 Navigate to the following path that contains the Tape Library Simulator Utility:
`</opt/VRTSralus/bin>`
For example:

```
cd /opt/VRTSralus/bin
```
- 3 Start the **mktls** utility.
For example:

```
./mktls
```
- 4 Select **Create a new simulated tape library**, and then press **Enter**.
- 5 Enter the appropriate information.
See [“Simulated Tape Library options”](#) on page 1259.
- 6 Exit the utility.
- 7 Restart the Agent for Linux daemon.
See [“Starting the Agent for Linux daemon”](#) on page 1214.
- 8 On the Backup Exec server, restart the Backup Exec services.
See [“Starting and stopping Backup Exec services”](#) on page 621.

Simulated Tape Library options

When you create a simulated tape library, you must provide a directory path and the number of slots for the library.

See [“Creating a simulated tape library”](#) on page 1259.

Table Q-4 Simulated Tape Library options

Item	Description
Directory Path	Type the path of the directory for the simulated tape library. You can enter up to 512 characters. If the path does not exist, the Tape Library Simulator Utility creates it.
Number of Slots	Select the number of slots for this simulated tape library. The number of slots can range from 1 to 50. The default number of slots is 20.

See [“Viewing simulated tape libraries properties”](#) on page 1260.

Viewing simulated tape libraries properties

You can use the Symantec Tape Library Simulator Utility to view information about a simulated tape library and its contents.

To view simulated tape library properties

- 1 On the RMAI computer, stop the Agent for Linux daemon.
See [“Stopping the Agent for Linux daemon”](#) on page 1215.
- 2 Navigate to the following directory that contains the Tape Library Simulator Utility:

 /opt/VRTSralus/bin

 For example:

 `cd /opt/VRTSralus/bin`
- 3 Start the **mktls** utility.

 For example:

 `./mktls`
- 4 Select **View an existing simulated tape library**.
- 5 Move your cursor to the simulated tape library that you want to view, and then press **Enter**.
- 6 Press **Enter** again to view the simulated tape library properties.

 See [“Simulated tape library properties”](#) on page 1261.

- 7 Type **Q** to exit the utility.
- 8 Restart the Agent for Linux daemon.
- See “[Starting the Agent for Linux daemon](#)” on page 1214.

Simulated tape library properties

You can view the properties of a simulated tape library.

See “[Viewing simulated tape libraries properties](#)” on page 1260.

Table Q-5 Simulated tape library properties

Item	Description
Number of drives	Displays the number of drives for this simulated tape library. A simulated tape library can have only drive. This drive is not configurable.
Number of slots	Displays the number of slots for this simulated tape library. The number of slots can range from 1 to 50. The default number of slots is 20.
Tape capacity	Displays the tape capacity. The default capacity is 100 gigabytes.
Directory path	Displays the directory path where the simulated tape library exists.

Deleting a simulated tape library

You can use the Tape Library Simulator Utility to delete a simulated tape library. You must then manually delete the content of the simulated tape library files, and then delete the directories that contain these files.

To delete a simulated tape library

- 1 At the RMAL computer stop the Agent for Linux daemon.
See [“Stopping the Agent for Linux daemon”](#) on page 1215.
- 2 Navigate to the following directory that contains the Tape Library Simulator:
`/opt/VRTSralus/bin/`
For example:

```
cd /opt/VRTSralus/bin/
```
- 3 Start the **mktls** utility:
For example:

```
./mktls
```
- 4 Select **View an existing simulated tape library**.
- 5 Select the simulated tape library that you want to delete.
- 6 When you are prompted, delete the simulated tape library.
- 7 Exit the utility.
- 8 Restart the Agent for Linux daemon.
See [“Starting the Agent for Linux daemon”](#) on page 1214.
- 9 Find the simulated tape library files, and then manually delete them.
See [“About the Tape Library Simulator Utility”](#) on page 1258.
- 10 On the Backup Exec server, restart the Backup Exec services when it is convenient.
See [“Starting and stopping Backup Exec services”](#) on page 621.

Managing simulated tape libraries from the command line

You can use the command line to create a simulated tape library. Create a simulated tape library on a hard disk or on any mounted volume on the RMAL computer. From the command line, you can also view and delete simulated tape libraries.

To manage simulated tape libraries from the command line

- 1
- At the RMAL computer stop the Agent for Linux daemon.

See “[Stopping the Agent for Linux daemon](#)” on page 1215.
- 2
- Navigate to the following directory that contains the Tape Library Simulator Utility:

/opt/VRTSralus/bin

For example:

`cd /opt/VRTSralus/bin`
- 3
- Start the **mktls** utility with the appropriate parameter switches.

See “[Command line switches for the Tape Library Simulator Utility](#)” on page 1263.
- 4
- Start the Agent for Linux daemon.

See “[Starting the Agent for Linux daemon](#)” on page 1214.

Command line switches for the Tape Library Simulator Utility

You can use command line switches to manage simulated tape libraries. For example, the following command line creates a simulated tape library with 10 slots that is located at /TLS2/Testing.

```
./mktls -s10 -p/TLS2/Testing
```

See “[Managing simulated tape libraries from the command line](#)” on page 1262.

Table Q-6 Command line switches for the Tape Library Simulator Utility

Switch	Description
-p<path>	Specifies the path to the directory for the simulated tape library. If the path does not exist, the utility creates it. The maximum path size is 512 characters.
-s<number of slots>	Specifies the number of slots for this simulated tape library. The number of slots can range from one to 50. The default number is 20.
-r	Prevents the information from displaying.

Table Q-6 Command line switches for the Tape Library Simulator Utility
(continued)

Switch	Description
-l	Lists the simulated tape libraries that exist for RMAL.
-d -p<path>	Specifies the path of the simulated tape library that you want to delete.
-h	Displays the online Help.

Troubleshooting the Remote Media Agent for Linux

If there are issues with Remote Media Agent for Linux (RMAL), review the following questions and answers.

Table Q-7 Troubleshooting the RMAL

Question	Answer
RMAL does not detect my attached device. What should I do?	<p>First, ensure that Backup Exec and RMAL support the device.</p> <p>You can find a list of compatible devices at the following URL:</p> <p>http://entsupport.symantec.com/umi/V-269-2</p> <p>If the device is listed on the hardware compatibility list, ensure the following:</p> <ul style="list-style-type: none">■ The operating system detects the device■ The device is listed in /proc/scsi/scsi <p>If the operating system can detect the device, ensure that the device is listed in /etc/VRTSralus/TILDBG.TXT.</p>

Table Q-7 Troubleshooting the RMAL *(continued)*

Question	Answer
My Backup Exec server does not display the devices that are attached to my Remote Media Agent. What should I do?	<p>Try the following procedures:</p> <ul style="list-style-type: none">■ Ensure that the Agent for Linux daemon is running. If it is not running, start the daemon, and check that power for the server is enabled, and that all cables are properly attached.■ Ensure that RMAL properties are set to the correct port, and that ICMP ping operations are enabled.■ Ensure that the Backup Exec services are restarted after a Remote Media Agent is added to the Backup Exec server. The available devices should be displayed under RMAL node. <p>See “Editing properties for the Remote Media Agent for Linux ” on page 1255.</p> <p>See “Starting the Agent for Linux daemon” on page 1214.</p>
Why don't my remote devices appear in any of the storage device pools that are created by Backup Exec?	<p>By default, Backup Exec does not include remote devices in the storage device pools it creates. Symantec recommends that you create a separate storage device pool for the devices that are attached to each Remote Media Agent.</p> <p>See “About creating storage device pools for devices attached to the Remote Media Agent for Linux ” on page 1255.</p>

Table Q-7 Troubleshooting the RMAL (*continued*)

Question	Answer
<p>RMAL won't run on the remote computer. What should I do?</p>	<p>Ensure that RMAL is installed on a supported version of Linux.</p> <p>You can find a list of compatible operating systems, platforms, and applications at the following URL:</p> <p>http://entsupport.symantec.com/umi/V-269-1</p> <p>If you install RMAL to an unsupported version of Linux, RMAL is unavailable for use. You cannot create the jobs that run on the devices that are attached to the Linux server. However, you can back up the Linux server by using the Agent for Linux component. This component is installed with RMAL.</p> <p>To use the Agent for Linux component to back up the Linux server, do the following:</p> <ul style="list-style-type: none"> ■ Edit the <code>ralus.cfg</code> file. ■ In the string <code>Software\Symantec\Backup Exec for Windows\Backup Exec\Engine\RMAL\DisableRMAL=0</code>, change 0 to 1. <p>See “Editing configuration options for Linux computers” on page 1196.</p> <p>See “Running the begather utility to troubleshoot Backup Exec components on Linux servers” on page 706.</p>
<p>I cannot load RMAL. When I attempt to load RMAL in console mode, <code>/beremote --log-console</code> shows the following message:</p> <p>Error while loading shared libraries: libstdc++.so.5: cannot open shared object file: No such file or directory.</p> <p>What should I do?</p>	<p>This error indicates that the libstdc++.so.5 library is not in the <code>/usr/lib</code> directory. This library is necessary to let RMAL start and function. To resolve this issue, install the libstdc++.so.5 package. You can install this package from the media on which your copy of Linux was provided. Or, you can run the following command from a computer that has Internet access:</p> <pre>apt-get install libstdc++5</pre> <p>For SUSE Linux Enterprise Server 11, run the following command:</p> <pre>zypper install libstdc++5</pre>

Symantec Backup Exec Storage Provisioning Option

This appendix includes the following topics:

- [Features of the Storage Provisioning Option](#)
- [Requirements for the Storage Provisioning Option](#)
- [About installing the Storage Provisioning Option](#)
- [Viewing storage array components in Backup Exec](#)
- [Configuring a storage array by using the Configure Storage Wizard](#)
- [Viewing properties for a storage array and its physical disks](#)
- [About the Any Virtual Disk Storage device pool in the Storage Provisioning Option](#)
- [About virtual disks in the Storage Provisioning Option](#)
- [Adding a hot spare by using the Configure Storage Wizard](#)
- [Changing a hot spare by using the Configure Storage Wizard](#)
- [Detecting a new storage array](#)
- [Renaming a virtual disk or storage array](#)
- [Identifying the physical disks of a virtual disk by using the blink feature](#)
- [Troubleshooting the Storage Provisioning Option](#)

Features of the Storage Provisioning Option

The Storage Provisioning Option (SPO) lets you configure, manage, and monitor storage arrays that are attached to the Backup Exec server.

Note: If you use a Dell DL Appliance, do not use this appendix. See the *Dell™ PowerVault™ DL Backup to Disk Appliance and the Symantec Backup Exec Storage Provisioning Option* documentation that Dell provides with the appliance.

Table R-1 Features of the Storage Provisioning Option

Feature	Description
Discovery of new storage arrays, physical disks, and virtual disks	Backup Exec can discover new storage arrays, physical disks, and the virtual disks that you add to a storage array. If you create virtual disks by using storage array vendor tools or the Microsoft Storage Manager for SANs utility, Backup Exec also detects those virtual disks.
A wizard to help you configure a storage array for use with Backup Exec	Backup Exec provides the Configure Storage Wizard to help you configure virtual disks on a storage array. The virtual disks are added to the Any virtual disk storage device pool. Backup Exec then uses the virtual disks in the device pool as destination devices for jobs.
Trend analysis of disk space usage	Backup Exec collects statistical information to predict the amount of disk space that is required on the storage arrays. Alerts are sent if the available disk space does not meet the predicted amount of disk space that is needed. See “Storage trending statuses for disk storage and virtual disks” on page 316.
Alerts for low disk space	Backup Exec sends an alert when available disk space reaches each of three thresholds that you set for a virtual disk.

See [“Requirements for the Storage Provisioning Option”](#) on page 1269.

See [“Configuring a storage array by using the Configure Storage Wizard”](#) on page 1270.

Requirements for the Storage Provisioning Option

Do the following before you install the Backup Exec Storage Provisioning Option (SPO):

- Attach any storage arrays to the Backup Exec server.
- Install the storage array vendor's VDS hardware provider on the Backup Exec server.

Requirements for the Storage Provisioning Option in a CASO environment

The following are required to run the Storage Provisioning Option (SPO) in a Central Admin Server Option (CASO) environment:

- SPO must be installed on the Backup Exec server to which the storage array is attached.

If the storage array is attached to a managed Backup Exec server, install SPO on that managed Backup Exec server. You do not need to install SPO on the central administration server if the storage array is not attached to it.

- The Central Admin Server Option must use a centralized database.
See [“CASO distributed, centralized, and replicated catalog locations”](#) on page 1118.

See [“How to choose the location for CASO storage and media data”](#) on page 1099.

See [“About installing the Storage Provisioning Option”](#) on page 1269.

See [“Configuring a storage array by using the Configure Storage Wizard”](#) on page 1270.

About installing the Storage Provisioning Option

Install SPO on a local Backup Exec server as a separate add-on component of Backup Exec.

You can install SPO when you upgrade from a previous version of Backup Exec.

See [“Installing additional Backup Exec options to the local Backup Exec server”](#) on page 76.

See [“Requirements for the Storage Provisioning Option”](#) on page 1269.

Viewing storage array components in Backup Exec

After you install the Storage Provisioning Option (SPO), storage arrays appear in the **All Storage** view. After you use the **Configure Storage Wizard** to configure the storage array, virtual disks appear under the storage array to which they belong.

You can also view the properties of the physical disks that comprise a storage array's configured virtual disk or disks by viewing the details of the storage array itself.

To view storage array components in Backup Exec

- 1 On the **Storage** tab, under **All Storage**, expand a storage array.
- 2 View the storage array components.

See [“Properties for a storage array and its physical disks”](#) on page 1272.

See [“Configuring a storage array by using the Configure Storage Wizard”](#) on page 1270.

Configuring a storage array by using the Configure Storage Wizard

With the **Configure Storage Wizard**, you can:

- Specify a virtual disk size, and then let the storage array automatically choose the physical disks for the virtual disk's disk group.
With this option, the storage array attempts to spread the required disk space for the virtual disk equally across each physical disk it selects.
When you specify a virtual disk size, the **Configure Storage Wizard** presents you with the total amount of disk space available to create a virtual disk. The total available disk space is derived from the physical disks that are present. You can use the entire space to create one virtual disk, or you can specify smaller amounts of space and create multiple virtual disks. Each virtual disk you create is based on the amount of disk space you enter. To create multiple virtual disks, run the wizard multiple times, until you have the virtual disks you want, or until you exhaust the disk space.
- Create a disk group by manually selecting the physical disks, from which you can then choose the number of virtual disks to create.
With this option, the **Configure Storage Wizard** lets you customize the entire virtual disk creation process on the storage array.
- Configure hot spares.
You can also use this wizard to add or change hot spares for the disk groups that are already configured.

In addition, you can also use the **Configure Storage Wizard** to configure a virtual disk as deduplication disk storage. Deduplication reduces the disk storage and the network bandwidth that is used for backups by storing and sending only unique data.

Note: Deduplication may not be available for the storage array if deduplication disk storage presently exists on the Backup Exec server. Only one deduplication disk storage device per Backup Exec server is supported. You can delete the existing deduplication disk storage device and then rerun the **Configure Storage Wizard** to enable deduplication on the storage array.

Use the **Configure Storage Wizard** to configure a storage array for use with the Backup Exec Storage Provisioning Option (SPO).

When the wizard completes, it runs a utility job named Configure Storage Array. This utility job creates the virtual disks that you specified. Then, Backup Exec adds the virtual disks to a device pool named **Any virtual disk storage**. You can submit jobs to the **Any virtual disk storage** device pool, to the storage array, or to a specific virtual disk.

Note: In a Central Admin Server Option (CASO) environment, run the **Configure Storage Wizard** from the central administration server. You can run the **Configure Storage Wizard** for any managed Backup Exec server that has SPO installed. Managed Backup Exec servers can share a single storage array but cannot share a virtual disk on a storage array.

See [“Viewing storage array components in Backup Exec”](#) on page 1270.

See [“Requirements for the Deduplication Option”](#) on page 798.

To configure a storage array by using the Configure Storage Wizard

- 1 On the **Storage** tab, click **Configure Storage**.
- 2 Do one of the following:

In a non-CASO environment

Do the following in the order listed.

- Click **Disk-based Storage**, and then click **Next**.
- Click **Storage array**.
- Follow the on-screen directions.

In a CASO environment

Do the following in the order listed:

- Select a Backup Exec server for which you want to configure storage.
- Click **Disk-based Storage**, and then click **Next**.
- Click **Storage array**.
- Follow the on-screen directions.

Viewing properties for a storage array and its physical disks

Properties provide detailed information for the storage array and its physical disks.

To view properties for storage arrays

- 1 On the **Storage** tab, select a storage array.
- 2 Right-click the storage array, and then click **Details**.

See [“Properties for a storage array and its physical disks”](#) on page 1272.

Properties for a storage array and its physical disks

Properties provide detailed information for the storage array and its physical disks.

See [“Viewing properties for a storage array and its physical disks”](#) on page 1272.

Table R-2 Properties for storage arrays

Item	Description
Name	<p>Displays the default name that Backup Exec assigns to the storage array. The default name is Storage array x, where x is a number that increments each time that you add a storage array.</p> <p>You can change the default name of the storage array.</p> <p>See “Renaming a virtual disk or storage array” on page 1309.</p>
Description	<p>Displays a description that you use to describe a storage array.</p>

Table R-2 Properties for storage arrays (*continued*)

Item	Description
State	<p>Indicates the current state of the storage array.</p> <p>States are as follows:</p> <ul style="list-style-type: none">■ Pause The storage array is temporarily stopped. You can pause a storage array to perform maintenance activities. Active jobs are not affected if they start before the storage array is paused.■ Enable The storage array is available for use with Backup Exec. If the storage array is not enabled, it is available for use with other applications.■ Online The storage array is available for use.■ Offline Backup Exec cannot access the storage array. You can click this check box to try to bring the storage array online.■ Failed Backup Exec cannot access the storage array.
Identification	<p>Displays the information that the vendor provides to assist you with identification of the attached storage array hardware. Refer to the vendor documentation that is supplied with the storage array for information.</p>
Hardware name	<p>Displays the name that the storage array hardware or the vendor hardware provider assigns.</p>
Hot spare disk is configured	<p>Indicates whether or not a hot spare disk has been configured for the storage array.</p>

Table R-2 Properties for storage arrays (*continued*)

Item	Description
Total backup storage	<p>Displays the expected amount of total raw capacity of all of the physical disks in the storage array.</p> <p>If a storage array can read the total capacity, then it appears in this field. Otherwise, Backup Exec estimates the total capacity.</p>
Used capacity	<p>Displays the amount of backup storage capacity presently being used.</p>
Connection type	<p>Shows the location of the storage array in use with the Backup Exec server. Local appears when the storage array is physically connected to the Backup Exec server.</p>
Backup Exec service restarted needed	<p>Indicates if the Backup Exec services must be restarted to apply any changes that are made to this device.</p>
Hardware status	<p>Indicates the hardware status.</p> <p>The values for hardware status are as follows:</p> <ul style="list-style-type: none">■ OK The storage array is online.■ Offline The storage array and the virtual disks on the storage array are offline. Backup Exec cannot access them.■ Unknown The status of the storage array cannot be determined. <p>Refer to the vendor documentation and management software that are supplied with the storage array. After the storage array is brought online, the virtual disks are automatically brought online.</p>

Table R-2 Properties for storage arrays (*continued*)

Item	Description
Hardware health	<p>Indicates the hardware health.</p> <p>Values for hardware health are as follows:</p> <ul style="list-style-type: none">■ OK The storage array is online.■ Warning The storage array may fail or produce errors, but it is currently operational. The storage array and the virtual disks on the storage array are offline. Backup Exec cannot access them.■ Critical The storage array has failed. The storage array and the virtual disks on the storage array are offline. Backup Exec cannot access them. <p>Unknown The health of the storage array cannot be determined.</p> <p>Refer to the vendor documentation and management software that are supplied with the storage array. After the storage array is brought online, the virtual disks are automatically brought online.</p>
Total capacity	<p>Displays the expected amount of total raw capacity of all the of the physical disks in the storage array.</p> <p>If a storage array can read the total capacity, then it appears in this field. Otherwise, Backup Exec estimates the total capacity.</p>
Unconfigured capacity	<p>Displays the expected amount of total raw capacity of all of the unconfigured physical disks in the storage array.</p>
Unused configured capacity	<p>Displays the amount of total raw capacity of all of the configured physical disks in the storage array.</p>
Enclosure	<p>Identifies the enclosure that the physical disk is in.</p>

Table R-2 Properties for storage arrays (*continued*)

Item	Description
Slot	Identifies the slot that the physical disk occupies.
Capacity	Displays the total amount of available disk space on the physical disk in this slot.
State	<p>Displays the hardware status of a physical disk.</p> <p>Values for hardware status are as follows:</p> <ul style="list-style-type: none">■ OK The physical disk is online.■ Offline The physical disk is offline. The virtual disks that use this physical disk may also be offline. Backup Exec cannot access them.■ Failed The physical disk has failed. The virtual disks that use this physical disk may also fail. Backup Exec cannot access the virtual disks. If hot spares are configured, the virtual disk is automatically rebuilt. If your storage array does not support an automatic rebuild capability, you must use vendor tools to perform a manual rebuild of the virtual disks. Refer to your vendor documentation for the storage array for more information.■ Not configured The drive is not configured as a virtual disk. <p>To troubleshoot issues, refer to the vendor documentation and management software that are supplied with the storage array.</p>

Table R-2 Properties for storage arrays (*continued*)

Item	Description
Status	<p>Displays the status of a physical disk within an enclosure.</p> <ul style="list-style-type: none">■ OK The physical disk is online.■ Offline The physical disk is offline. The virtual disks that use this physical disk may also be offline. Backup Exec cannot access them.■ Failed The physical disk has failed. The virtual disks that use this physical disk may also fail. Backup Exec cannot access the virtual disks.■ Unknown The status of the physical disk cannot be determined.
Health	<p>Displays the hardware health of a physical disk.</p> <p>Values for hardware health are as follows:</p> <ul style="list-style-type: none">■ OK The physical disk is online.■ Warning The physical disk may fail or produce errors, but it is currently operational.■ Critical The physical disk may fail. You should replace the physical disk.■ Unknown The health of the physical disk cannot be determined. <p>To troubleshoot issues, refer to the vendor documentation and management software that are supplied with the storage array.</p>

About the Any Virtual Disk Storage device pool in the Storage Provisioning Option

After you install the Storage Provisioning Option (SPO), Backup Exec adds the **Any virtual disk storage** device pool to the list of storage device pools. The **Any virtual disk storage** device pool contains all virtual disks from all storage arrays on all computers in the Backup Exec environment.

See [“Creating storage device pools”](#) on page 416.

About virtual disks in the Storage Provisioning Option

A virtual disk is a logical disk that you create on a storage array to provide virtual storage to the Backup Exec server.

You can use any of the following to create a virtual disk:

- The Configure Storage Wizard
- The management tools that the vendor of the storage array provides
- The Microsoft Storage Manager for SANs management tool

If you create a virtual disk with a tool other than the **Configure Storage Wizard**, you must configure the virtual disk for use with Backup Exec. After you configure a virtual disk, Backup Exec uses it as a destination device for jobs. Backup Exec automatically adds configured virtual disks to the **Any virtual disk storage** device pool.

See [“Configuring a virtual disk on a storage array”](#) on page 1285.

In the **Configure Storage Wizard**, you specify the number of virtual disks to create from the physical disks that are in the storage array. The Backup Exec server cannot access the physical disks. The Backup Exec server can access only the virtual disks that you create.

Backup Exec uses a configured virtual disk in the same manner in which it uses a disk storage device.

Backup Exec does not assign a drive letter to the virtual disk. You cannot browse for a virtual disk or access it from a command prompt. Since you cannot browse to the virtual disk, you cannot back it up with Backup Exec. Symantec recommends that you create a duplicate backup data job to move the data from the virtual disk to another device. For example, you can move the data to a tape device or to another virtual disk on a separate storage array.

Backup Exec provides three low disk space thresholds for the virtual disks. As available disk space reaches each threshold, Backup Exec sends an alert. When

the available disk space on the virtual disk reaches the third threshold, the alert warns you to create more disk space immediately.

See [“Editing default options for a virtual disk on a storage array”](#) on page 1280.

Note: You cannot share a virtual disk between two computers.

Virtual disks can have the following states:

Table R-3 States for virtual disks








States for virtual disks	Icon	Description
Configured		The virtual disk is in use as a device or as a hot spare.
Online		The virtual disk is online.
Not configurable		The virtual disk cannot be configured because it is in a bad state, or it has failed.
Not configured		The virtual disk is available for configuration but has not yet been configured.
Offline		The virtual disk is offline. Backup Exec cannot access it.
Disabled		The virtual disk is disabled. Backup Exec cannot access it.

Table R-3 States for virtual disks (continued)

States for virtual disks	Icon	Description
Paused		The virtual disk is paused.

See [“Viewing storage array components in Backup Exec”](#) on page 1270.

See [“Editing default options for a virtual disk on a storage array”](#) on page 1280.

See [“Viewing properties for unconfigured virtual disks on a storage array”](#) on page 1285.

See [“About the Any Virtual Disk Storage device pool in the Storage Provisioning Option”](#) on page 1278.

Editing default options for a virtual disk on a storage array

You can set the default options that apply to individual virtual disks.

See [“Editing properties of virtual disks on storage arrays”](#) on page 1292.

To edit default options for a virtual disk on a storage array

- 1 On the **Storage** tab, select a storage array.
- 2 Expand the storage array, and then double-click the virtual disk that you want to view.
- 3 Change any of the following options:

Name	<p>Displays the default name that Backup Exec assigns to the unconfigured virtual disk.</p> <p>The default name is Virtual disk x, where x is a number that increments each time that you create a virtual disk.</p> <p>You can change the default name of the virtual disk.</p> <p>See “Renaming a virtual disk or storage array” on page 1309.</p>
Description	<p>Displays a description of the virtual disk.</p> <p>You can edit this field.</p>

Limit Backup Exec to read-only operations

Indicates whether or not you want Backup Exec to reclaim disk space from this virtual disk. If you reattach a storage device to Backup Exec after a specified number of days, and if any data has expired, then Backup Exec may attempt to reclaim the disk space. To prevent this situation, limit Backup Exec to read-only operations so that you can restore data from this storage without having the existing data overwritten.

If you enable this option, you can specify the number of days that the storage must be detached from Backup Exec before this option takes effect.

Maximum file size

Lets you set a maximum file size for the disk-based file that Backup Exec creates during a backup job.

When this option is enabled, you can set a file size limit for a file that is created during a backup job. If Backup Exec reaches the maximum file size limit during a job, it creates an additional backup file and job spanning occurs. The additional backup file size conforms to the maximum file size limit that you set in this option.

Preallocate disk space incrementally up to the maximum file size

Creates the file when the job starts by preallocating space incrementally, according to the size of the increment that you set in Preallocation increment. As the job uses the disk space, more disk space is preallocated up to the maximum file size. When the job completes, the file size is then reduced to the amount of disk space that the job actually used.

For example, if you enable preallocation and set the preallocation increment to 4 GB, then 4 GB of disk space is preallocated when the job starts. After the job uses 4 GB, then Backup Exec allocates another 4 GB. Disk space continues to be preallocated by 4 GB until the job completes. If the job only uses 13 GB of the 16 GB that was allocated, then the file size is reduced to 13 GB.

Choices are **Enabled** or **Disabled**.

Note: This option works with the option **Maximum file size**.

Preallocation increment

Displays the amount of disk space by which to increase the file size. The file size is increased by this increment as the job requires disk space, up to the maximum file size. The default is 1 GB.

Note: This option appears when you enable the option **Preallocate disk space incrementally up to the maximum file size**.

Auto detect block and buffer size

Indicates if Backup Exec automatically detects the preferred settings for block size, and for the read and write buffers for the virtual disk. The default value is Enabled.

When you disable this option, you can manually set the virtual block size and buffers size options as desired.

Choices are Enabled or Disabled.

Low disk space - Critical

Displays the critically low disk space threshold at which you want Backup Exec to send an alert. Backup Exec sends alerts when the amount of free disk space drops below the low disk space threshold, and again if it drops below the warning threshold. The amount of free disk space does not include the disk space that is reserved for non-Backup Exec operations.

You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes. This threshold must be less than the warning threshold.

The default is 5%.

Low disk space - Warning

Displays the low disk space threshold at which you want Backup Exec to send an alert. If free disk space drops below the warning threshold to the critical threshold, another alert is sent. The amount of free disk space does not include the disk space that is reserved for non-Backup Exec operations.

You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes. This threshold must be less than the low disk space threshold.

The default is 15%.

Low disk space

Displays the low disk space threshold at which you want Backup Exec to send an alert. If free disk space drops below this threshold to the amount specified in the warning threshold, another alert is sent. If free disk space drops below the warning threshold to the critical threshold, another alert is sent. The amount of disk space does not include the disk space that is reserved for non-Backup Exec operations.

You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes.

The default is 25%.

Auto detect settings

Indicates if Backup Exec automatically detects the preferred settings for read and write buffers for the virtual disk.

Buffered read

Indicates the following when the setting is enabled:

- You do not want Backup Exec to automatically detect settings for this virtual disk.
- You want this virtual disk to allow buffered reads, which is the reading of large blocks of data.

Enabling buffered reads may provide increased performance.

Buffered write

Indicates the following when the setting is enabled:

- You do not want Backup Exec to automatically detect settings for this virtual disk.

You want this virtual disk to allow buffered writes, which is the writing of large blocks of data.

Concurrent write operations

Displays the number of concurrent write operations that you want to allow to this virtual disk.

Configuring a virtual disk on a storage array

If you create a virtual disk with a tool other than Backup Exec, then you must configure the virtual disk to use it with Backup Exec. Backup Exec can only use configured virtual disks as destination devices for jobs. When you configure the virtual disk, Backup Exec submits a job named **Configure Disk Storage**. When the job completes successfully, the virtual disk is configured and added to the **Any virtual disk storage** device pool.

Note: Be careful when you select an unconfigured virtual disk. An unconfigured virtual disk may be in use as a Microsoft SQL Server database, an Exchange database, or a boot disk.

To configure a virtual disk on a storage array

- 1 On the **Storage** tab, select a storage array.
- 2 Expand the storage array, and then select the unconfigured virtual disk.
- 3 On the **Storage Array Operations** group, click **Configure Virtual Disk**.
- 4 Follow the on-screen instructions.

See [“Editing properties of virtual disks on storage arrays”](#) on page 1292.

Viewing properties for unconfigured virtual disks on a storage array

You can view the properties of an unconfigured virtual disk on a storage array.

Note: You must configure a virtual disk before Backup Exec can use it as a destination device for jobs.

See [“Configuring a virtual disk on a storage array”](#) on page 1285.

To view properties for unconfigured virtual disks on a storage array

- 1 On the **Storage** tab, select a storage array.
- 2 Expand a storage array, and then select an unconfigured virtual disk.
- 3 Right-click the unconfigured virtual disk, and then click **Details**.

See [“Properties for unconfigured virtual disks on storage arrays”](#) on page 1285.

Properties for unconfigured virtual disks on storage arrays

Properties for unconfigured virtual disks provide information on the name, status, and health of the disks.

See [“Viewing properties for unconfigured virtual disks on a storage array”](#) on page 1285.

Table R-4 Properties for unconfigured virtual disks on storage arrays

Item	Description
Name	<p>Displays the default name that Backup Exec assigns to the unconfigured virtual disk.</p> <p>The default name is Virtual disk x, where x is a number that increments each time that you create a virtual disk.</p> <p>You can change the default name of the virtual disk.</p> <p>See “Renaming a virtual disk or storage array” on page 1309.</p>
Description	<p>Displays a description of the virtual disk.</p> <p>You can edit this field.</p>

Table R-4 Properties for unconfigured virtual disks on storage arrays
(continued)

Item	Description
State	<p>Indicates the current state of the virtual disk.</p> <p>States are as follows:</p> <ul style="list-style-type: none">■ Pause The virtual disk is temporarily stopped. You can pause a virtual disk to perform maintenance activities. Active jobs are not affected if they start before the virtual disk is paused.■ Enable The virtual disk is available for use with Backup Exec. If the virtual disk is not enabled, it is available for use with other applications.■ Online The virtual disk is available for use.■ Offline Backup Exec cannot access the virtual disk. You can click this check box to try to bring the virtual disk online.■ Failed Backup Exec cannot access the virtual disk.■ Not configured The virtual disk has not been configured. You can configure the virtual disk by clicking Configure Virtual Disk in the Storage Array Operations group.
Hardware name	<p>Displays the name that you assign to a virtual disk if you use a vendor-specific tool to create the virtual disk.</p>

Table R-4 Properties for unconfigured virtual disks on storage arrays
(continued)

Item	Description
Hardware status	<p>Displays the hardware status.</p> <p>Values for the hardware status are as follows:</p> <ul style="list-style-type: none">■ OK The unconfigured virtual disk is online.■ Offline The unconfigured virtual disk is offline.■ Failed The unconfigured virtual disk has failed.■ Unknown The status of the hardware cannot be determined. <p>See your vendor's documentation for more information.</p>
Hardware health	<p>Displays the hardware health.</p> <p>Values for the hardware health are as follows:</p> <ul style="list-style-type: none">■ OK The unconfigured virtual disk is online.■ Warning The unconfigured virtual disk may fail or produce errors, but it is currently operational.■ Critical The unconfigured virtual disk has failed.■ Unknown The status of the hardware cannot be determined. <p>See your vendor's documentation for more information.</p>

Table R-4

Properties for unconfigured virtual disks on storage arrays

(continued)

Item	Description
Disk classification	

Table R-4 Properties for unconfigured virtual disks on storage arrays
(continued)

Item	Description
	<p>Displays the type of disk group that the unconfigured virtual disk is on.</p> <p>Disk classifications are as follows:</p> <ul style="list-style-type: none">■ Simple (RAID 0) A single physical disk, no striping, or parity. No redundancy.■ Span A set of multiple physical disks that are concatenated together. No striping or parity. No redundancy.■ Stripe A set of multiple physical disk extents with data striped across the physical disks. No redundancy.■ Mirror (RAID 1) A pair or multiple pairs of physical disks with the same data written to each physical disk of the pair. Provides for data redundancy.■ Stripe with parity (RAID 5 or RAID 6) Three or more physical disks with data striped across the physical disks, with one disk's worth of space that is used for parity. Provides for data redundancy.■ Unknown <p>Backup Exec creates only physical disk groups with a disk classification of Stripe with parity (RAID 5/RAID 6). If another disk classification appears, the disk group was created with a vendor-supplied tool other than the Storage Provisioning Option (SPO).</p> <p>Additional disk classifications may appear, including the following:</p> <ul style="list-style-type: none">■ RAID-10 Implements striping over mirroring technology.■ RAID-50 Implements a striped (RAID 0) array that is striped across a RAID 5 array.

Table R-4 Properties for unconfigured virtual disks on storage arrays
(continued)

Item	Description
	<ul style="list-style-type: none">■ RAID-60 Implements a combination of multiple RAID 6 sets with RAID 0 (striping) See your vendor documentation for more information about RAID technologies.
Total capacity	<p>Displays the total amount of storage space that is available on the virtual disk.</p> <p>If the virtual disk is not configured, the value shown in this field is 0.</p>
Total backup storage	<p>Displays the expected amount of total raw capacity of all of the physical disks in the storage array.</p> <p>If a storage array can read the total capacity, then it appears in this field. Otherwise, Backup Exec estimates the total capacity.</p> <p>If the virtual disk is not configured, this field is empty.</p>
Used capacity	<p>Displays the amount of raw capacity of all of the physical disks that are used in the storage array. Backup Exec calculates used capacity by subtracting available capacity from total capacity.</p> <p>If the virtual disk is not configured, this field is empty.</p>
Amount of data written	<p>Displays the amount of data that has been written to the disk. The amount of data written may differ from the used capacity due to the effects of data compression. Data compression tends to increase the amount of data written when compared to used capacity.</p> <p>If the virtual disk is not configured, this field is empty.</p>

Table R-4 Properties for unconfigured virtual disks on storage arrays
(continued)

Item	Description
Available capacity	<p>Displays the amount of expected raw capacity on the disk that remains unused. Available capacity is calculated by subtracting the amount of reserved disk space on the drive from the drive's total amount of space.</p> <p>If the virtual disk is not configured, the value shown in this field is 0.</p>
Compression ratio	<p>Displays the ratio of bytes written to used capacity. Compression ratio will show the overall effect that data compression and media flaws are having on the amount of data that is being stored on the virtual disk.</p> <p>If the virtual disk is not configured, this field is empty.</p>

Editing properties of virtual disks on storage arrays

You can edit properties for a virtual disk on a storage array.

To edit properties of virtual disks on storage arrays

- 1 On the **Storage** tab, select a storage array.
- 2 Expand the storage array, and then select the virtual disk on which you want to edit properties.
- 3 Right-click the virtual disk, and then click **Details**.
- 4 Edit the following properties as appropriate:

Name

Displays the default name that Backup Exec assigns to the virtual disk.

The default name is Virtual disk x, where x is a number that increments each time that you create a virtual disk.

You can change the default name of the virtual disk.

See [“Renaming a virtual disk or storage array”](#) on page 1309.

Limit Backup Exec to read-only operations

Indicates whether or not you want Backup Exec to reclaim disk space from this virtual disk. If you reattach a storage device to Backup Exec after a specified number of days, and if any data has expired, then Backup Exec may attempt to reclaim the disk space. To prevent this situation, limit Backup Exec to read-only operations so that you can restore data from this storage without having the existing data overwritten.

If you enable this option, you can specify the number of days that the storage must be detached from Backup Exec before this option takes effect.

Maximum file size

Lets you set a maximum file size for the disk-based file that Backup Exec creates during a backup job.

When this option is enabled, you can set a file size limit for a file that is created during a backup job. If Backup Exec reaches the maximum file size limit during a job, it creates an additional backup file and job spanning occurs. The additional backup file size conforms to the maximum file size limit that you set in this option.

Preallocate disk space incrementally up to the maximum file size

Creates the file when the job starts by preallocating space incrementally, according to the size of the increment that you set in Preallocation increment. As the job uses the disk space, more disk space is preallocated up to the maximum file size. When the job completes, the file size is then reduced to the amount of disk space that the job actually used.

For example, if you enable preallocation and set the preallocation increment to 4 GB, then 4 GB of disk space is preallocated when the job starts. After the job uses 4 GB, then Backup Exec allocates another 4 GB. Disk space continues to be preallocated by 4 GB until the job completes. If the job only uses 13 GB of the 16 GB that was allocated, then the file size is reduced to 13 GB.

Choices are **Enabled** or **Disabled**.

Note: This option works with the option **Maximum file size**.

Preallocation increment

Displays the amount of disk space by which to increase the file size. The file size is increased by this increment as the job requires disk space, up to the maximum file size. The default is 1 GB.

Note: This option appears when you enable the option **Preallocate disk space incrementally up to the maximum file size**.

Auto detect block and buffer size

Indicates if Backup Exec automatically detects the preferred settings for block size, and for the read and write buffers for the virtual disk. The default value is Enabled.

When you disable this option, you can manually set the virtual block size and buffers size options as desired.

Choices are Enabled or Disabled.

Low disk space - Critical

Displays the critically low disk space threshold at which you want Backup Exec to send an alert. Backup Exec sends alerts when the amount of free disk space drops below the low disk space threshold, and again if it drops below the warning threshold. The amount of free disk space does not include the disk space that is reserved for non-Backup Exec operations.

You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes. This threshold must be less than the warning threshold.

The default is 5%.

Low disk space - Warning

Displays the low disk space threshold at which you want Backup Exec to send an alert. If free disk space drops below the warning threshold to the critical threshold, another alert is sent. The amount of free disk space does not include the disk space that is reserved for non-Backup Exec operations.

You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes. This threshold must be less than the low disk space threshold.

The default is 15%.

Low disk space

Displays the low disk space threshold at which you want Backup Exec to send an alert. If free disk space drops below this threshold to the amount specified in the warning threshold, another alert is sent. If free disk space drops below the warning threshold to the critical threshold, another alert is sent. The amount of disk space does not include the disk space that is reserved for non-Backup Exec operations.

You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes.

The default is 25%.

Auto detect settings

Indicates if Backup Exec automatically detects the preferred settings for read and write buffers for the virtual disk.

Buffered read

Indicates the following when the setting is enabled:

- You do not want Backup Exec to automatically detect settings for this virtual disk.
- You want this virtual disk to allow buffered reads, which is the reading of large blocks of data.

Enabling buffered reads may provide increased performance.

Buffered write

Indicates the following when the setting is enabled:

- You do not want Backup Exec to automatically detect settings for this virtual disk.

You want this virtual disk to allow buffered writes, which is the writing of large blocks of data.

Concurrent write operations

Displays the number of concurrent write operations that you want to allow to this virtual disk.

See [“Properties for virtual disks on storage arrays”](#) on page 1297.

Properties for virtual disks on storage arrays

Properties provide information about virtual disks on storage arrays.

See [“Editing properties of virtual disks on storage arrays”](#) on page 1292.

Table R-5 Properties for virtual disks on storage arrays

Item	Description
Name	<p>Displays the default name that Backup Exec assigns to the virtual disk.</p> <p>The default name is Virtual disk x, where x is a number that increments each time that you create a virtual disk.</p> <p>You can change the default name of the virtual disk.</p> <p>See “Renaming a virtual disk or storage array” on page 1309.</p>
Description	Displays a description of the virtual disk.
State	<p>Displays the current status of the virtual disk.</p> <p>Statuses for a virtual disk are as follows:</p> <ul style="list-style-type: none">■ Pause The virtual disk is temporarily stopped.■ Enable The virtual disk is available for use with Backup Exec. If the virtual disk is disabled, it is available for use with other applications. Backup Exec does not monitor the low disk space thresholds for a disabled virtual disk.■ Online The virtual disk is available for use.■ Offline Backup Exec cannot access the virtual disk. You can check Offline to try to bring the storage array online.
Hardware name	Displays the name that the storage array hardware or the vendor hardware provider assigns.

Table R-5 Properties for virtual disks on storage arrays (*continued*)

Item	Description
Hardware status	<p>Displays the hardware status.</p> <p>Values for the hardware status are as follows:</p> <ul style="list-style-type: none">■ OK The virtual disk is online.■ Offline The virtual disk is offline. Backup Exec cannot access it. To bring the virtual disk online, refer to the vendor documentation and management software that are supplied with the storage array.■ Failed The virtual disk has failed. Backup Exec cannot access it. To troubleshoot the issue, refer to the vendor documentation and management software that are supplied with the storage array. After the issue is resolved, the virtual disk is automatically brought online. <p>Unknown The status of the hardware cannot be determined.</p> <p>See your vendor's documentation for more information.</p>

Table R-5 Properties for virtual disks on storage arrays (*continued*)

Item	Description
Hardware health	<p>Displays the hardware health.</p> <p>Values for the hardware health are as follows:</p> <ul style="list-style-type: none">■ OK The virtual disk is online.■ Warning The virtual disk may fail or produce errors, but it is currently operational.■ Critical The virtual disk has failed. Backup Exec cannot access it. To troubleshoot the issue, refer to the vendor documentation and management software that are supplied with the storage array.■ Unknown The status of the hardware health cannot be determined. <p>See your vendor's documentation for more information.</p>

Table R-5 Properties for virtual disks on storage arrays *(continued)*

Item	Description
Disk classification	

Table R-5 Properties for virtual disks on storage arrays (*continued*)

Item	Description
	<p>Displays the type of disk group that the virtual disk is on.</p> <p>Disk classifications are as follows:</p> <ul style="list-style-type: none">■ Simple (RAID 0) A single physical disk, no striping, or parity. No redundancy.■ Span A set of multiple physical disks that are concatenated together. No striping or parity. No redundancy.■ Stripe A set of multiple physical disk extents with data striped across the physical disks. No redundancy.■ Mirror (RAID 1) A pair or multiple pairs of physical disks with the same data written to each physical disk of the pair. Provides for data redundancy.■ Stripe with parity (RAID 5 or RAID 6) Three or more physical disks with data striped across the physical disks, with one disk's worth of space that is used for parity. Provides for data redundancy.■ Unknown <p>Backup Exec creates only physical disk groups with a disk classification of Stripe with parity (RAID 5/RAID 6). If another disk classification appears, the disk group was created with a vendor-supplied tool other than SPO.</p> <p>Additional disk classifications may appear, including the following:</p> <ul style="list-style-type: none">■ RAID-10 Implements striping over mirroring technology.■ RAID-50 Implements a striped (RAID 0) array that is striped across a RAID 5 array. <p>RAID-60</p>

Table R-5 Properties for virtual disks on storage arrays (*continued*)

Item	Description
	<ul style="list-style-type: none">■ Implements a combination of multiple RAID 6 sets with RAID 0 (striping) <p>See your vendor documentation for more information about RAID technologies.</p>
Limit Backup Exec to read-only operations	<p>Indicates whether or not you want Backup Exec to reclaim disk space from this virtual disk. If you reattach a storage device to Backup Exec after a specified number of days, and if any data has expired, then Backup Exec may attempt to reclaim the disk space. To prevent this situation, limit Backup Exec to read-only operations so that you can restore data from this storage without having the existing data overwritten.</p> <p>If you enable this option, you can specify the number of days that the storage must be detached from Backup Exec before this option takes effect.</p>
Maximum file size	<p>Lets you set a maximum file size for the disk-based file that Backup Exec creates during a backup job.</p> <p>When this option is enabled, you can set a file size limit for a file that is created during a backup job. If Backup Exec reaches the maximum file size limit during a job, it creates an additional backup file and job spanning occurs. The additional backup file size conforms to the maximum file size limit that you set in this option.</p>

Table R-5 Properties for virtual disks on storage arrays (*continued*)

Item	Description
Preallocate disk space incrementally up to the maximum file size	<p>Creates the file when the job starts by preallocating space incrementally, according to the size of the increment that you set in Preallocation increment. As the job uses the disk space, more disk space is preallocated up to the maximum file size. When the job completes, the file size is then reduced to the amount of disk space that the job actually used.</p> <p>For example, if you enable preallocation and set the preallocation increment to 4 GB, then 4 GB of disk space is preallocated when the job starts. After the job uses 4 GB, then Backup Exec allocates another 4 GB. Disk space continues to be preallocated by 4 GB until the job completes. If the job only uses 13 GB of the 16 GB that was allocated, then the file size is reduced to 13 GB.</p> <p>Choices are Enabled or Disabled.</p> <p>Note: This option works with the option Maximum file size.</p>
Preallocation increment	<p>Displays the amount of disk space by which to increase the file size. The file size is increased by this increment as the job requires disk space, up to the maximum file size. The default is 1 GB.</p> <p>Note: This option appears when you enable the option Preallocate disk space incrementally up to the maximum file size.</p>
Auto detect block and buffer size	<p>Indicates if Backup Exec automatically detects the preferred settings for block size, and for the read and write buffers for the virtual disk. The default value is Enabled.</p> <p>When you disable this option, you can manually set the virtual block size and buffers size options as desired.</p> <p>Choices are Enabled or Disabled.</p>

Table R-5 Properties for virtual disks on storage arrays (*continued*)

Item	Description
Low disk space - Critical	<p>Displays the critically low disk space threshold at which you want Backup Exec to send an alert. Backup Exec sends alerts when the amount of free disk space drops below the low disk space threshold, and again if it drops below the warning threshold. The amount of free disk space does not include the disk space that is reserved for non-Backup Exec operations.</p> <p>You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes. This threshold must be less than the warning threshold.</p> <p>The default is 5%.</p>
Low disk space - Warning	<p>Displays the low disk space threshold at which you want Backup Exec to send an alert. If free disk space drops below the warning threshold to the critical threshold, another alert is sent. The amount of free disk space does not include the disk space that is reserved for non-Backup Exec operations.</p> <p>You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes. This threshold must be less than the low disk space threshold.</p> <p>The default is 15%.</p>

Table R-5 Properties for virtual disks on storage arrays (*continued*)

Item	Description
Low disk space	<p>Displays the low disk space threshold at which you want Backup Exec to send an alert. If free disk space drops below this threshold to the amount specified in the warning threshold, another alert is sent. If free disk space drops below the warning threshold to the critical threshold, another alert is sent. The amount of disk space does not include the disk space that is reserved for non-Backup Exec operations.</p> <p>You can change the value of the threshold, and you can change the amount of disk space to megabytes or gigabytes.</p> <p>The default is 25%.</p>
Total capacity	Displays the size of the volume on the virtual disk.
Total backup storage	<p>Displays the expected amount of total raw capacity of all of the physical disks in the storage array.</p> <p>If a storage array can read the total capacity, then it appears in this field. Otherwise, Backup Exec estimates the total capacity.</p>
Used capacity	Displays the amount of raw capacity of all of the physical disks that are used in the storage array. Backup Exec calculates used capacity by subtracting available capacity from total capacity.
Amount of data written	Displays the amount of data that has been written to the disk. The amount of data written may differ from the used capacity due to the effects of data compression. Data compression tends to increase the amount of data written when compared to used capacity.

Table R-5 Properties for virtual disks on storage arrays (*continued*)

Item	Description
Available capacity	Displays the amount of expected raw capacity on the disk that remains unused. Available capacity is calculated by subtracting the amount of reserved disk space on the drive from the drive's total amount of space.
Compression ratio	Displays the ratio of the uncompressed size of a file over its compressed size.
Path	Displays the physical location of the virtual disk on the storage array on the Backup Exec server.
Connection type	Indicates if the virtual disk is located on disk that is local to the Backup Exec server or if it is on a remote disk.
Backup Exec service restart needed	Indicates if the Backup Exec services must be restarted to apply any changes that are made to the Backup Exec server.
Auto detect settings	Indicates if Backup Exec automatically detects the preferred settings for read and write buffers for the virtual disk.
Buffered read	<p>Indicates the following when the setting is enabled:</p> <ul style="list-style-type: none">■ You do not want Backup Exec to automatically detect settings for this virtual disk.■ You want this virtual disk to allow buffered reads, which is the reading of large blocks of data. <p>Enabling buffered reads may provide increased performance.</p>

Table R-5 Properties for virtual disks on storage arrays (*continued*)

Item	Description
Buffered write	<p>Indicates the following when the setting is enabled:</p> <ul style="list-style-type: none">■ You do not want Backup Exec to automatically detect settings for this virtual disk. <p>You want this virtual disk to allow buffered writes, which is the writing of large blocks of data.</p>
Concurrent write operations	<p>Displays the number of concurrent write operations that you want to allow to this virtual disk.</p>

Adding a hot spare by using the Configure Storage Wizard

If a storage array that has automatic rebuild capability loses virtual disk redundancy, it uses a physical disk as a hot spare to regain redundancy. If your storage array does not support an automatic rebuild capability, you must use vendor tools to manually rebuild the virtual disks. Refer to your storage vendor documentation for more information.

You can use the **Configure Storage Wizard** to add a hot spare in a storage array. When you complete this wizard, it submits a utility job named Configure Storage Array. When the job completes successfully, the hot spare has been added.

Before you specify a hot spare, refer to the following best practices:

- Specify at least one hot spare for each enclosure. Although you can specify only one hot spare for all of the enclosures, consider the risk if more than one physical disk fails.
- Specify the physical disks that are in slot 0 in the enclosures as hot spares. Then, you can quickly identify which disk is a hot spare.
- Specify a hot spare that is at least the same size as the physical disk that it replaces. If the hot spare is smaller than the physical disk, the storage array cannot rebuild the virtual disk.

For more recommendations, refer to your storage array vendor's documentation.

To add a hot spare by using the Configure Storage Wizard

- 1 On the **Storage** tab, click **Configure Storage Wizard**.
- 2 Click **Disk-based Storage**, and then click **Next**.
- 3 Select **Storage array**, and then click **Next**.
- 4 Select a storage array on which to configure a hot spare drive, and then click **Next**.
- 5 Select the option, **I only want to configure hot spares**, and then click **Next**.
- 6 Follow the on-screen instructions.

See “[Changing a hot spare by using the Configure Storage Wizard](#)” on page 1308.

See “[Configuring a storage array by using the Configure Storage Wizard](#)” on page 1270.

Changing a hot spare by using the Configure Storage Wizard

You can use the **Configure Storage Wizard** to select a different physical disk to use as a hot spare in a storage array. When you complete this wizard, it submits a utility job named **Configure Storage Array**. When the job completes successfully, the hot spare has been changed.

See “[Adding a hot spare by using the Configure Storage Wizard](#)” on page 1307.

To change a hot spare by using the Configure Storage Wizard

- 1 On the **Storage** tab, click **Configure Storage Wizard**.
- 2 Click **Disk-based Storage**, and then click **Next**.
- 3 Select **Storage array**, and then click **Next**.
- 4 Select a storage array on which to change a hot spare drive, and then click **Next**.
- 5 Click **I only want to configure hot spares**, and then click **Next**.
- 6 Do one of the following:

To designate a hot spare as an available physical disk Do the following in the order listed:

- In the **Hot Spares** list, select the hot spare that you want to return to the **Available Physical Disks** list.

To designate an available physical disk as a hot spare Do the following in the order listed:

- In the **Available Physical Disks** list, select one or more physical disks that you want to use as hot spares.

- 7 Click **Next**.
- 8 Review the **Summary** panel, and then click **Finish**.

See [“Viewing storage array components in Backup Exec”](#) on page 1270.

Detecting a new storage array

Backup Exec periodically searches for new storage arrays or new physical disks. If Backup Exec does not find a new storage array or physical disk that you added, then you should run a refresh operation. If a refresh operation does not discover the new devices, then restart the Backup Exec services.

After you restart the services, the new storage array appears on the **Storage** tab, under **All Storage**.

You must install the Storage Provisioning Option (SPO) before Backup Exec can detect a new storage array.

To detect a new storage array

- 1 Click the **Storage** tab.
- 2 If you recently added a new storage array or a physical disk, and neither appears under **All Storage**, press F5 to refresh the view.
- 3 If the refresh does not discover the storage array, restart the Backup Exec services.

See [“Starting and stopping Backup Exec services”](#) on page 621.

See [“Troubleshooting the Storage Provisioning Option”](#) on page 1311.

Renaming a virtual disk or storage array

You can rename a virtual disk or storage array. Names cannot exceed 128 characters. You cannot change the hardware name.

If you use a vendor-supplied tool to configure the storage array, the hardware name that you assign it in the vendor tool appears. To change the hardware name of the storage array, you must use the vendor-supplied tool or the Microsoft Storage Manager for SANs utility.

To rename a virtual disk or storage array

- 1 On the **Storage** tab, under **All Storage**, do one of the following:

To rename a virtual disk

Do the following in the order listed.

- Expand a storage array whose virtual disk you want to rename.
- Right-click a virtual disk, then select **Details**.
- In the **Name** field, type a new name for the virtual disk.
- Continue with step 2.

To rename a storage array

Do the following in the order listed.

- Right-click a storage array, and then select **Details**.
- In the **Name** field, type a new name for the storage array.
- Continue with step 2.

- 2 Click **Apply**.

See [“About virtual disks in the Storage Provisioning Option”](#) on page 1278.

Identifying the physical disks of a virtual disk by using the blink feature

Many storage array enclosures incorporate a set of physical disks that use small status lights to indicate the operational status of physical disks. The Storage Provisioning Option (SPO) uses these lights with its **blink** feature to help you quickly identify the physical disks that comprise a virtual disk. When you select the **blink** feature for a virtual disk, the status lights on the physical disks blink.

Note: Storage array support for the blink features depends on storage array hardware support for the feature. Not all storage array hardware supports blinking. See your storage array hardware documentation for more information.

You can use the blink feature in different ways. You can use it to assist with:

- Moving virtual disks from one storage array to another.
You can use the blink feature when you want to move a virtual disk from one enclosure to another. If you have many enclosures, you can use the blink feature

to identify the physical disks that comprise the virtual disk. Without it, determining the physical disks that comprise a virtual disk can be difficult.

- Identifying problematic physical disks.
When SPO generates an alert for a physical disk issue, you can use the blink feature to assist you in finding problematic physical disks.

When you use the blink feature, the following applies:

- The blink feature works on one virtual disk at a time.
You cannot use it to simultaneously identify the physical disks in multiple virtual disks.

Use the following steps to identify the physical disks of a virtual disk.

To identify the physical disks of a virtual disk

- 1 On the **Storage** tab, select a storage array.
- 2 Expand the storage array, and then select a virtual disk.
- 3 Right-click the virtual disk, and then click **Blink**
- 4 To turn off the blink feature, right-click the storage array and then click **Unblink**.

Troubleshooting the Storage Provisioning Option

If problems occur with the Storage Provisioning Option (SPO) or with storage array hardware, ensure the following:

- The vendor storage array and the vendor hardware provider are supported.
You can find a list of compatible devices at the following URL:
<http://entsupport.symantec.com/umi/v-269-2>
- The storage array has power and is turned on.
- All lights and indicators on the storage array appear normal.
- The storage array is properly zoned if it is on a SAN.
- The cables are plugged into the correct ports.
- The Microsoft DiskRAID command line tool or the Microsoft Storage Manager for SANs management tool can detect and exercise the storage array hardware.
- The Disk Manager can detect the unmasked virtual disks.
- The refresh operation has run to detect new virtual disks.

If you installed the trial version of SPO, ensure that the license is still within the trial period. When the trial period expires, the option functions in a very limited mode.

See “[Features of the Storage Provisioning Option](#)” on page 1268.

Symantec Backup Exec Archiving Option

This appendix includes the following topics:

- [About the Archiving Option](#)
- [Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option](#)
- [Installing the Backup Exec Archiving Option](#)
- [How the Archiving Option works](#)
- [Best practices for the Archiving Option](#)
- [About creating an archive job](#)
- [Editing default settings for the Archiving Option](#)
- [About single instance storage of archived items](#)
- [About synchronizing archive permissions and settings](#)
- [About vault stores in the Archiving Option](#)
- [About vault store partitions in the Archiving Option](#)
- [About archives in the Archiving Option](#)
- [About Archiving Option operation entries in the audit log](#)
- [About Exchange mailbox groups in archive jobs](#)
- [About deleting archived data from its original location](#)
- [About archive settings in the Archiving Option](#)

- [About managing index locations in the Backup Exec Archiving Option](#)
- [About restoring items from the archives](#)
- [About searching for data in the archives](#)
- [About deleting data from the archives](#)
- [About backing up and restoring the Archiving Option components from a remote Backup Exec server](#)
- [Preventing the deletion of expired archived items from an archive](#)
- [About backing up Archiving Option components](#)
- [About restoring an Archiving Option component](#)
- [About moving Archiving Option components to a new location](#)
- [Troubleshooting archive jobs](#)
- [Reports for the Archiving Option](#)
- [About Backup Exec Virtual Vault](#)

About the Archiving Option

The Archiving Option includes the following features that you can install separately or together:

- The File System Archiving Option, which archives the eligible Windows file system data.
- The Exchange Mailbox Archiving Option, which archives the eligible Exchange mail messages.

To find the data that is eligible for archiving, Backup Exec applies rules to the selected file system shares and folders and to the Exchange mailboxes. Data in the selections is eligible for archiving if it is backed up and meets the criteria that the rules specify. The archive job then sends the data to disk-based vault stores. The data is deleted from its original location on the source computer immediately after it is archived, or after you back up the Archiving Option components.

You can apply retention categories to the data that is archived that specify how long to keep data in the archives. Backup Exec automatically deletes the archived data that has expired retention dates.

By archiving data from the backup sets, Backup Exec eliminates additional querying and movement of data on the remote computers. After Backup Exec deletes the

archived data from its original location, you have more disk space, and Backup Exec requires less time for future backup jobs.

The Archiving Option uses Symantec Enterprise Vault technology to archive data. When you install the Archiving Option, some Enterprise Vault services are also installed.

When you install the Exchange Mailbox Archiving Option, the Virtual Vault feature is automatically installed. Virtual Vault lets end users view the archived email messages from Microsoft Outlook on their desktop and laptop computers. These desktop and laptop computers are the Virtual Vault clients, and they connect to the Backup Exec server for synchronization using a DNS alias. You must install the Backup Exec Outlook Add-In on the end users' computers to enable Virtual Vault.

You can set defaults for each archive job that you create. You can also set the defaults that apply to all archive jobs, and that let you manage other archive operations.

See [“About creating an archive job”](#) on page 1340.

See [“Editing default settings for the Archiving Option”](#) on page 1346.

See [“About Enterprise Vault services for the Archiving Option”](#) on page 1336.

See [“Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option”](#) on page 1314.

See [“Installing the Backup Exec Archiving Option”](#) on page 1332.

See [“How the Archiving Option works”](#) on page 1337.

See [“Best practices for the Archiving Option”](#) on page 1339.

See [“About Backup Exec Virtual Vault ”](#) on page 1389.

Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option

Some requirements apply to both the Exchange Mailbox Archiving Option and to the File System Archiving Option. You should also review the requirements that apply to each specific option.

See [“Requirements for the Exchange Mailbox Archiving Option”](#) on page 1316.

See [“Requirements for the File System Archiving Option”](#) on page 1319.

Following are the requirements that apply to both of the Archiving Options:

- The Backup Exec server must be part of a domain. You cannot install the Archiving Option on a server in a workgroup.

- Only common encryption keys can be used to decrypt a backup set during an archive job. If a restricted key is used, then eligible items in the backup set are not archived.
- The Backup Exec server must have enough space to store the Archiving Options index files. When you install the Archiving Options, you are prompted to provide a path where the index files are stored. The path must be on a local NTFS volume. Paths on FAT and FAT32 partitions are not supported.

Note: Symantec recommends that you have more RAM available in addition to Backup Exec's base requirements.

See “[System requirements for Backup Exec](#)” on page 63.

You can find a list of compatible operating systems, platforms, and applications at the following URL:

<http://entsupport.symantec.com/umi/V-269-1>

The Archiving Option does not support the following:

- The Backup Exec Central Admin Server Option.

Note: You can install the Archiving Option on a central administration server. However, distributed job management for archive jobs is not supported.

- Archiving from backup sets from the Backup Exec Remote Media Agent for Linux.
- Installation on clustered servers. Additionally, you cannot install Backup Exec on a cluster if you have also selected the Archiving Option for installation.
- Installation of the Exchange Mailbox Archiving Option on a computer on which Microsoft Exchange Server is installed.
- Server groups or multiple servers that are selected for backup.

See “[How to calculate disk space requirements for the Exchange Mailbox Archiving Option](#)” on page 1326.

See “[How to calculate disk space requirements for the File System Archiving Option](#)” on page 1329.

See “[About Enterprise Vault services for the Archiving Option](#)” on page 1336.

Requirements for the Exchange Mailbox Archiving Option

Requirements that apply to both of the Archiving Options are listed separately. Review those requirements in addition to the requirements that apply only to the Exchange Mailbox Archiving Option.

See [“Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option”](#) on page 1314.

If appropriate, you should also review Exchange Mailbox Archiving Option support for Exchange Server 2010.

See [“Requirements for Exchange Mailbox Archiving Option support for Exchange Server 2010”](#) on page 1318.

Following are the requirements that apply to the Exchange Mailbox Archiving Option:

- You must enter an Exchange Agent license on the Backup Exec server for each Exchange Server that you want to archive.
- You must install one of the following versions of Microsoft Outlook on the Backup Exec server before you install the Exchange Mailbox Archiving Option:
 - Outlook 2007 Service Pack 2 with hot fix kb968858 or later
 - Outlook 2003 Service Pack 3
- Microsoft Outlook must be the default mail client on the Backup Exec server.
- You must create a profile when you install Outlook on the Backup Exec server, and then connect to an Exchange Server mailbox. Outlook may display an error message about a conflicting program. If Outlook offers to fix the problem, choose to do so, and then follow the instructions that are given.
- The Exchange Server backups must have the Granular Recovery Technology (GRT) option enabled.
- The Exchange Server backups must be on one of the following types of storage:
 - Non-removable disk storage.
 - Deduplication disk storage.
 - A storage array in a Storage Provisioning Option environment.
- The Exchange Server backups must have a valid path configured on an NTFS volume that is local to the Backup Exec server for temporary storage of data. The default path is set to use C:\temp.
See [“Setting default Granular Recovery Technology \(GRT\) options”](#) on page 596.
- A mailbox must be configured for exclusive use by Backup Exec on each Exchange Server on which you want to select mailboxes to archive. Whenever you create an archive job for the Archiving Option for Exchange you are prompted

to enter the name of the system mailbox. The system mailbox is the mailbox that you configure for use by Backup Exec. You do not need to be name this mailbox as *system* mailbox.

- The following are restrictions for this mailbox:
 - The mailbox must not be used for any other purpose. The Exchange Mailbox Archiving Option requires exclusive access.
 - The mailbox must not be hidden from address lists.
 - The mailbox account must not be disabled.
- The Backup Exec server domain and the Exchange Server domains must trust the domain that the Backup Exec service account belongs to.
- You must grant permissions to the Backup Exec service account to the Exchange servers.
See [“About granting permissions on the Exchange Server for the Backup Exec service account in the Archiving Option ”](#) on page 1320.
- The Backup Exec service account must be a member of the Active Directory domain. Symantec recommends that you use a Backup Exec service account that has domain and local administrator rights on the Exchange Server. The Backup Exec service account does not need to be a domain administrator.

The components that are listed in the following table are required to install or upgrade the Exchange Mailbox Archiving Option.

Table S-1 Installed or required components for the Exchange Mailbox Archiving Option

Operating system	Installed or required components
Windows Server 2008	Backup Exec automatically installs IIS 7.0 and some IIS 7.0 role services on Windows Server 2008. See “About the Internet Information Services (IIS) 7.0 role services installed by Backup Exec” on page 1333.

Table S-1 Installed or required components for the Exchange Mailbox Archiving Option *(continued)*

Operating system	Installed or required components
Windows Server 2003	<p>Backup Exec automatically installs ASP.NET 2.0.</p> <p>On Windows 2003 x64, Backup Exec installs and configures ASP.NET 2.0 in 32-bit mode.</p> <p>You must install the following components manually:</p> <ul style="list-style-type: none"> ■ Internet Information Services (IIS) 6.0 ■ World Wide Web Publishing Service <p>See your Microsoft documentation for information about installing these components.</p>

Requirements for Exchange Mailbox Archiving Option support for Exchange Server 2010

You should also review all other requirements that may apply to the Exchange Mailbox Archiving Option.

See [“Requirements for the Exchange Mailbox Archiving Option”](#) on page 1316.

The following are requirements for Exchange Mailbox Archiving Option support for Exchange Server 2010:

- You must install Microsoft Outlook 2007 Service Pack 2 with hot fix kb968858 on the Backup Exec server.
- You must install the Exchange Management Tools for Exchange Server 2010 on the Backup Exec server.

The management tools on the Backup Exec server must be the same version or later as the management tools that are on the Exchange Server 2010. You can install the management tools when you perform a custom installation of Exchange Server 2010.

Note: If you install the management tools and Backup Exec together on a Backup Exec server, install the tools first. If you install Backup Exec before you install the management tools, you must restart the Backup Exec server when the tools installation is complete.

Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option

- If you select mailbox databases in a Microsoft Exchange Database Availability Group (DAG), you must configure a system mailbox for each member server in the DAG.
- You must grant the Backup Exec service account Send As permissions on a system mailbox
- You must grant permissions for Exchange Server 2007/2010 for the Archiving Option.

See [“Granting the Backup Exec service account Send As permissions on a system mailbox for Exchange Server 2010”](#) on page 1321.

See [“Granting permissions for Exchange Server 2007/2010 for the Archiving Option”](#) on page 1323.

Requirements for the File System Archiving Option

Some requirements that apply to both of the Archiving Options and to the File System Archiving Option are listed separately. Review those requirements in addition to the requirements that apply only to the File System Archiving Option.

See [“Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option”](#) on page 1314.

The following are requirements for the File System Archiving Option:

- The Backup Exec server domain, the file server domains, and the Exchange Server domains must trust the domain that the Backup Exec service account belongs to.
- The Backup Exec server domain must trust the domains that contain the accounts of users that access the file server shares.
- The Backup Exec service account must have local administrative rights on the file server.
- The Backup Exec service account must have Full Control share permissions on the share that is selected for archiving.
- The Backup Exec service account must be granted the following NTFS rights on the folders in the share that is selected for archiving:
 - **Modify**
 - **List Folder Contents**
 - **Read**
 - **Write**

- You must install one of the following Outlook versions on the Backup Exec server to provide full indexing of .msg files:
 - Outlook 2003 Service Pack 3
 - Outlook 2007 Service Pack 2 with hotfix kb968858 or later

About granting permissions on the Exchange Server for the Backup Exec service account in the Archiving Option

The Backup Exec service account must access mailboxes on the Exchange Servers that you want to archive. To gain this access, the Backup Exec service account must have permission to access the Exchange Servers.

You can use either of the following methods to grant the permissions that the Backup Exec service account needs to access the mailboxes on Exchange Servers:

- Grant permissions at the organization level or at the Administrative Group level. Permissions are then propagated automatically to any new Exchange Servers that you add under the level at which the permissions are assigned.

Note: You must have Exchange administrative permissions to grant permissions to other accounts.

- Grant permissions explicitly on each Exchange Server.
If you grant permissions explicitly and then add another Exchange Server, you must grant permissions explicitly on the added server as well.

The Backup Exec service account must also have the Send As permission on the mailbox that you create for Backup Exec's exclusive use. This mailbox, called the system mailbox, must be created on each Exchange Server on which you want to select mailboxes to archive.

See [“Granting permissions at the Organization level for Exchange Server 2007 for the Archiving Option”](#) on page 1323.

See [“Granting permissions explicitly on each Exchange Server 2007 for the Archiving Option”](#) on page 1324.

See [“Granting permissions at the Organization level for Exchange Server 2003 for the Archiving Option”](#) on page 1325.

See [“Granting permissions at the server level for Exchange Server 2003 for the Archiving Option”](#) on page 1326.

See [“Granting the Backup Exec service account Send As permissions on a system mailbox for Exchange Server 2010”](#) on page 1321.

See [“Granting permissions for Exchange Server 2007/2010 for the Archiving Option”](#) on page 1323.

Granting the Backup Exec service account Send As permissions on a system mailbox for Exchange Server 2010

The Backup Exec service account requires Send As permissions on the system mailbox on each Exchange server. You can set this permission manually on each account, or use the following procedure.

See [“Requirements for Exchange Mailbox Archiving Option support for Exchange Server 2010”](#) on page 1318.

To grant the Backup Exec service account Send As permissions on a system mailbox

- 1 Log on to the Exchange Server using an account with the management role of Active Directory Permissions.

By default, members of the Organization Management role group are assigned this role.
- 2 On the Exchange Server, click **Start > All Programs > Microsoft Exchange Server 2010 > Exchange Management Shell**.
- 3 Type the following command:

```
Add-ADPermission -Identity '<system mailbox name>' -User '<Domain Name\Backup Exec service account>' -AccessRights ExtendedRight -ExtendedRights "send as"
```

See [“Microsoft Exchange permissions granted to the Backup Exec service account”](#) on page 1321.

Microsoft Exchange permissions granted to the Backup Exec service account

The following table lists the permissions that the script SetEVEExchangePermissions.ps1 grants to the Backup Exec service account.

See [“Granting the Backup Exec service account Send As permissions on a system mailbox for Exchange Server 2010”](#) on page 1321.

Table S-2 Permissions granted to the Backup Exec service account

Path	Object	Permissions
CN=Configuration, CN=Services, CN=Microsoft Exchange, CN=Organization, CN=Administrative Groups, CN=AdminGroup	Permissions are granted to the CN=Database and descendant objects if Exchange Server 2010 exists in the environment. Permissions are granted to the CN=Servers and descendant objects if Exchange Server 2007 or previous exists in the environment.	The following permissions are granted: <ul style="list-style-type: none"> ■ Read ■ Administer Information Store ■ Open mail send queue ■ Receive as ■ Send as ■ View Information Store status
CN=Configuration, CN=Services, CN=Microsoft Exchange	Permissions are granted to the following object: CN=Organization	Read
CN=Configuration, CN=Services, CN=Microsoft Exchange, CN=Organization	Permissions are granted to the following objects: <ul style="list-style-type: none"> ■ CN=ELC Folders Container and descendant objects ■ CN=Global Settings and descendant objects ■ CN=Transport Settings 	Read
CN=Configuration, CN=Services, CN=Microsoft Exchange, CN=Organization, CN=Transport Settings	Permissions are granted to the following object: CN=Rules	Read
CN=Configuration, CN=Services, CN=Microsoft Exchange, CN=Organization, CN=Transport Settings, CN=Rules	Permissions are granted to the following objects: <ul style="list-style-type: none"> ■ CN=Journaling and descendant objects ■ CN=JournalingVersioned and descendant objects 	Read

Granting permissions for Exchange Server 2007/2010 for the Archiving Option

You can use a PowerShell script to assign the necessary permissions to the Backup Exec service account for Exchange Server 2007/2010 support.

You must run this script on Exchange Server 2010 or Exchange Server 2007. Note that the script assigns the permissions that all the Exchange Server versions in your environment require, including Exchange Server 2003.

To grant permissions for Exchange Server 2007/2010 for the Archiving Option

- 1 Log on to the Exchange Server using an account with the following management roles:
 - Active Directory Permissions
 - Exchange Servers
 - Organization ConfigurationBy default, members of the "Organization Management" role group are assigned these roles.
- 2 Copy the script `SetEVEExchangePermissions.ps1` to the Exchange Server.

The script `SetEVEExchangePermissions.ps1` is located on the Backup Exec installation path at `\enterprise vault\powershellscripts`.
- 3 On the Exchange Server, open the Exchange Management Shell.
- 4 Type the following command:

```
SetEVEExchangePermissions.ps1 -User "<Domain Name\Backup Exec service account>"
```
- 5 To force these changes to take effect immediately, restart the Exchange Information Store service on each Exchange mailbox server.

See ["Microsoft Exchange permissions granted to the Backup Exec service account"](#) on page 1321.

Granting permissions at the Organization level for Exchange Server 2007 for the Archiving Option

You can grant the **Full Control** permission for the Backup Exec service account at the Organization level.

See ["About granting permissions on the Exchange Server for the Backup Exec service account in the Archiving Option "](#) on page 1320.

Note: You must have Exchange administrative permissions to grant permissions to other accounts.

To grant permissions at the Organization level for the Backup Exec service account in the Archiving Option

- 1 On the Exchange Server, click **Start > All Programs > Microsoft Exchange Server 2007 > Exchange Management Shell**.

- 2 Type the following command:

```
Get-OrganizationConfig | Add-ADPermission -User '<Domain  
Name\Backup Exec service account>' -AccessRights GenericAll  
-InheritanceType All
```

- 3 Type the following command:

```
Add-ADPermission -Identity '<system mailbox name>' -User '<Domain  
Name\Backup Exec service account>' -ExtendedRights 'Send-as'
```

- 4 To grant **Send As** permission on the mailboxes that you created for Backup Exec's exclusive use, repeat the previous step on the appropriate Exchange Servers.

Granting permissions explicitly on each Exchange Server 2007 for the Archiving Option

You can grant the **Full Control** permission for the Backup Exec service account on each Exchange Server. Perform this procedure on each Exchange Server that you want to archive.

See [“About granting permissions on the Exchange Server for the Backup Exec service account in the Archiving Option ”](#) on page 1320.

Note: You must have Exchange administrative permissions to grant permissions to other accounts.

To grant permissions explicitly on each Exchange Server 2007 for the Archiving Option

- 1 On the Exchange Server, click **Start > All Programs > Microsoft Exchange Server 2007 > Exchange Management Shell**.

- 2 Type the following command:

```
Get-MailboxServer -Identity "<mailbox server name>" > |  
Add-ADPermission -User "<Domain name\Backup Exe service account>"  
-AccessRights GenericAll -InheritanceType All
```

- 3 Type the following command:

```
Add-ADPermission -Identity '<system mailbox name>' -User '<Domain  
Name\Backup Exec service account>' -ExtendedRights 'Send-as'
```

Granting permissions at the Organization level for Exchange Server 2003 for the Archiving Option

You can grant the **Full Control** permission for the Backup Exec service account at the Organization level.

See [“About granting permissions on the Exchange Server for the Backup Exec service account in the Archiving Option ”](#) on page 1320.

Refer to the Microsoft knowledge base for more information on this procedure.

To grant permissions at the Organization level for Exchange Server 2003 for the Archiving Option

- 1 Configure the **ShowSecurityPage** registry setting to enable the display of the **Security** page.
- 2 In the left pane of Microsoft Exchange System Manager, right-click the **Exchange Organization**, and then click **Properties**.
- 3 On the **Security** tab, click **Add**.
- 4 Select the Backup Exec service account to add it to the list.
- 5 Click **OK**.
- 6 In the **Name** list, select the Backup Exec service account.
- 7 In the **Permissions** list, ensure that all of the check boxes in the **Allow** column are selected.
- 8 Select any check boxes that are not selected.
- 9 Click **OK**.

Granting permissions at the server level for Exchange Server 2003 for the Archiving Option

You can grant permissions for the Backup Exec service account at the server level for Exchange Server 2003. Perform this procedure on each Exchange Server that you want to archive.

See [“About granting permissions on the Exchange Server for the Backup Exec service account in the Archiving Option”](#) on page 1320.

To grant permissions at the server level for Exchange Server 2003 for the Archiving Option

- 1 In the left pane of the **Microsoft Exchange System Manager**, expand the **Servers** container.
- 2 Right-click the Exchange Server, and then click **Properties**.
- 3 On the **Security** tab, click **Add**.
- 4 Select the Backup Exec service account to add it to the list.
- 5 Click **OK**.
- 6 In the **Name** list, click the Backup Exec service account.
- 7 In the **Permissions** list, ensure that all of the check boxes in the **Allow** column are selected.
- 8 Select any check boxes that are not selected.
- 9 Click **OK**.

How to calculate disk space requirements for the Exchange Mailbox Archiving Option

Backup Exec requires permanent disk space for the following components in the Exchange Mailbox Archiving Option:

- The vault store partitions.
- The index locations.
- The SQL Server database, which contains the Directory, vault store, and fingerprint databases.

You should also consider the size of the temporary cache that is required for the Virtual Vault feature. The default maximum size of the temporary cache is 10 GB, but it can be increased.

See [“About the temporary cache location on the Backup Exec server”](#) on page 1400.

[Table S-3](#) describes the formulas that you can use to estimate the disk space requirements for these components for the Exchange Mailbox Archiving Option.

The following values and variables are used in the formulas:

- N is the number of emails.
- m is the average number of identical copies of attachments across user mailboxes.
- The compression factor for attachments is estimated as 60%.
If the attachments are mostly Office 2007 files, the compression factor to use is 90%.
- The average number of emails that have attachments is estimated at 20%.
- The average size of an email attachment is estimated at 250 KB.

Table S-3

Calculations for disk space requirements for the Exchange Mailbox Archiving Option

Component	Requirements
Vault store partitions	<p>The size of a vault store partition depends on the following items:</p> <ul style="list-style-type: none"> ■ The size of the emails. ■ The type of attachments. ■ The number and size of the attachments. ■ The number of emails with attachments. <p>Note: If single instance storage is enabled, items are shared within and across vault stores and vault store partitions. Shareable parts of a message that exceed the single-instance threshold of 20 KB are shared. These parts include attachments and message bodies. User information and shareable parts under the single instance threshold are not shared.</p> <p>See “About single instance storage of archived items” on page 1349.</p> <p>You can use the following calculations to approximate the disk space requirements of a vault store partition:</p> <ul style="list-style-type: none"> ■ Approximate vault store partition size for which single instance storage is not enabled: $(N \times 16) + (N \times 0.2 \times 0.6 \times 250)$ kilobytes ■ Approximate vault store partition size for which single instance storage is enabled: $(N \times 16) + ((1/m) \times (N \times 0.2 \times 0.6 \times 250))$ kilobytes <p>For example, you want to know the disk space requirements for a vault store partition for 100,000 emails. You estimate that each email attachment is shared across three people on average.</p> <p>If single instance storage of archived items is not enabled, the calculation for the disk space requirements is as follows:</p> $(100000 \times 16) + (100000 \times 0.2 \times 0.6 \times 250) \text{ kilobytes} = 4.6 \text{ GB approximately}$ <p>If single instance storage is enabled, the calculation for the disk space requirements is as follows:</p> $(100000 \times 16) + ((1/3) \times 100000 \times 0.2 \times 0.6 \times 250) \text{ kilobytes} = 2.6 \text{ GB approximately}$

Table S-3 Calculations for disk space requirements for the Exchange Mailbox Archiving Option (*continued*)

Component	Requirements
Indexes	<p>The size of an index is approximately 8% of the total size of the items that are archived. The percentage may be less if there is less content to index. For example, there is less content to index when there are large attachments such as MP3 or .jpeg files.</p> <p>For example, you have 100,000 emails that each have a body size of 8 KB. About 20% of the emails have attachments, each with an average total size of 250 KB. The index size is approximately 450 MB.</p>
Directory database	<p>The Directory database only grows when a new mailbox or share is archived for the first time.</p> <p>The recommended disk space is 500 MB.</p>
Vault store database	<p>The size of a vault store database is approximately $N \times 500$ bytes. The vault store database grows with every item that is archived. Temporary space is used to hold information on the items that have not been backed up or indexed.</p>
Fingerprint database	<p>The fingerprint database is created only if you enable single instance storage of archived items. Backup Exec initially allocates 212 MB for the fingerprint database. The fingerprint database grows with every item that is archived.</p> <p>If the database grows to more than 212 MB, use the following calculation to estimate the disk space that it requires:</p> $1/m \times N \times 0.2 \times 500 \text{ bytes}$ <p>See "About single instance storage of archived items" on page 1349.</p>

How to calculate disk space requirements for the File System Archiving Option

Backup Exec requires permanent disk space for the following File System Archiving Option components:

- Vault store partitions.
- Indexes.
- SQL Server database.

[Table S-4](#) describes the formulas that you can use to estimate the disk space requirements for these components for the File System Archiving Option.

The following values and variables are used in the formulas:

- *N* is the number of files.
- *m* is the average number of identical copies per file.
If *m* is unknown, use the estimate of 1.2.
- The compression factor for files is estimated as 50%.
This estimate applies to a mix of files that contain mostly Office 2003 documents. Office 2007 documents do not compress but when mixed with non-Office files, the compression average is 80% of the original size. Pure image files are not compressed.

Table S-4 Calculating disk space requirements for the File System Archiving Option components

Component	Disk space requirements
Vault store partition	<p>You can use the following calculations to approximate the disk space requirements of a vault store partition:</p> <ul style="list-style-type: none"> ■ Approximate vault store partition size for which single instance storage of archived items is not enabled: $(N \times 4) + (N \times \text{average file size in kilobytes} \times 0.5)$ kilobytes ■ Approximate vault store partition size for which single instance storage is enabled: $(N \times 4) + ((1/m) \times N \times \text{average file size} \times 0.5)$ kilobytes <p>For example, you want to know the disk space requirement for a vault store partition for 10,000 files. The average size of each file is 250 KB, and the average number of identical copies per file is 1.2.</p> <p>If single instance storage of archived items is not enabled, then the calculation for the disk space requirement is as follows:</p> $(10000 \times 4) + (100000 \times 250 \times 0.5) \text{ kilobytes} = 1.3 \text{ GB approximately}$ <p>If single instance storage is enabled, then the calculation for the disk space requirement is as follows:</p> $(10000 \times 4) = ((1/1.2) \times 10000 \times 250 \times 0.5) \text{ kilobytes} = 1.08 \text{ GB approximately}$

Table S-4 Calculating disk space requirements for the File System Archiving Option components (*continued*)

Component	Disk space requirements
Indexes	<p>You can estimate that the index files require approximately 2% of the total size of the files that are archived. The percentage may be less if there is less content to index. If the files are all compressed image file, indexing is less than if the files are mostly small text messages. Numerous small text messages require disk space requirements similar to those of the Archive Option for Exchange for indexing.</p> <p>For example, to archive 10 GB of data, at least 200 MB of available disk space is required to store the index files.</p>
Directory database	<p>The Directory database only grows when a new mailbox or share is archived for the first time.</p> <p>The recommended disk space is 1 GB.</p>
Vault store database	<p>The vault store database grows with every item that is archived. Temporary space is used to hold information on the items that have not been backed up or indexed.</p> <p>The size of a vault store database is approximately $N \times 3000$ bytes.</p>
Fingerprint database	<p>The fingerprint database is created only if you enable single instance storage of archived items. The fingerprint database holds the shareable parts of archived items. Shareable parts of an item that exceed the single-instance threshold of 20 KB are shared. For the File System Archiving Option, it is expected that all files are larger than the 20-KB threshold.</p> <p>Backup Exec initially allocates 212 MB for the fingerprint database. The fingerprint database grows with every item that is archived.</p> <p>If the database grows to more than 212 MB, use the following calculation to estimate the disk space that it requires:</p> <p>$1/m \times N \times 500$ bytes</p> <p>See “About single instance storage of archived items” on page 1349.</p>

Installing the Backup Exec Archiving Option

You can install one or both of the following options locally as a separate, add-on component of Backup Exec:

- Exchange Mailbox Archiving Option
- File System Archiving Option

When you install or upgrade the Backup Exec Exchange Mailbox Archiving Option, the Virtual Vault feature is automatically installed and enabled on the Backup Exec server.

During installation of the Exchange Mailbox Archiving Option, the installation wizard prompts you to specify the following information:

- A path to a location on the Backup Exec server where archived emails are temporarily cached for synchronization with Virtual Vault clients.
The default location is `C:\Program Files\Symantec\Backup Exec\AOCache`.
- The maximum size of the temporary cache on the Backup Exec server.
The default is 10 GB.
- A fully qualified domain name to create a DNS alias for the Backup Exec Exchange Mailbox Archiving Option.
An example of a fully qualified domain name is `serveralias.domain.com`. The use of a DNS alias simplifies operations if you must move the Backup Exec Archiving Option to another computer. The DNS alias is also used for the connection to the Backup Exec server for virtual vault synchronization.
- Ensure that the server on which you are installing the Archiving Option does not have any unformatted hard drives. The Archiving Option cannot be installed on an unformatted drive. If your server has an unformatted drive, you should either format the drive or remove it before you attempt to install the Archiving Option.

Command line switches are also available for silent mode installation of these options.

See [“About using the command line to install the Backup Exec Exchange Mailbox Archiving Option”](#) on page 1333.

Before you attempt to install the Archiving Option, verify that all requirements are met.

See [“Installing additional Backup Exec options to the local Backup Exec server”](#) on page 76.

See [“Repairing the Backup Exec Archiving Option”](#) on page 1337.

See [“Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option”](#) on page 1314.

See [“About Enterprise Vault services for the Archiving Option”](#) on page 1336.

See [“About installing Enterprise Vault on a Backup Exec server on which the Archiving Option is installed”](#) on page 1336.

See [“About Backup Exec Virtual Vault ”](#) on page 1389.

See [“About the Internet Information Services \(IIS\) 7.0 role services installed by Backup Exec”](#) on page 1333.

See [“About Vault Cache synchronization”](#) on page 1398.

About using the command line to install the Backup Exec Exchange Mailbox Archiving Option

If you use the command line to install or upgrade Backup Exec, an additional switch is required to install the Exchange Mailbox Archiving Option. Use this switch to specify a fully-qualified DNS alias.

Use the following command line switch as part of the silent mode installation:

`/BEAODNSALIAS: <fully qualified DNS alias >`

See [“Command line switches for silent mode installation of Backup Exec”](#) on page 109.

See [“Requirements for the Exchange Mailbox Archiving Option”](#) on page 1316.

About the Internet Information Services (IIS) 7.0 role services installed by Backup Exec

On Windows Server 2008, Backup Exec automatically installs Internet Information Services (IIS) 7.0 and some IIS role services. These role services are described in the following table.

Table S-5 Installed IIS 7.0 role services

IIS 7.0 features	IIS 7.0 role services
Common HTTP Features	<p>The following role services are installed on the Backup Exec server:</p> <ul style="list-style-type: none">■ Static Content■ Default Document■ Directory Browsing■ HTTP Errors■ HTTP Redirection
Application Development Features	<p>The following role services are installed on the Backup Exec server:</p> <ul style="list-style-type: none">■ ASP.NET■ .NET Extensibility■ Active Server Pages (ASP)■ Internet Server Application Programming Interface (ISAPI) Extensions■ ISAPI Filters
Health and Diagnostics Features	<p>The following role services are installed on the Backup Exec server:</p> <ul style="list-style-type: none">■ HTTP Logging■ Logging Tools■ Request Monitor■ Tracing
Security Features	<p>The following role services are installed on the Backup Exec server:</p> <ul style="list-style-type: none">■ Basic Authentication■ Windows Authentication■ Request Filtering■ IP and Domain Restrictions
Performance Features	<p>The Static Content Compression role service is installed on the Backup Exec server.</p>

Table S-5 Installed IIS 7.0 role services (*continued*)

IIS 7.0 features	IIS 7.0 role services
Management Tools	<p>The following role services are installed on the Backup Exec server:</p> <ul style="list-style-type: none">■ IIS Management Console■ IIS Management Scripts and Tools■ Management Service■ IIS 6.0 Management Compatibility■ IIS Metabase Compatibility■ IIS 6 WMI Compatibility■ IIS 6 Scripting Tools■ IIS 6 Management Console

See [“Requirements for the Exchange Mailbox Archiving Option”](#) on page 1316.

About uninstalling or reinstalling the Archiving Option

If you uninstall both the Exchange Mailbox Archiving Option and the File System Archiving Option, the Enterprise Vault files and the Enterprise Vault services that are included in the Archiving Option are removed. All archive-related jobs display the status as **Disabled**. You cannot run, edit, or save a disabled job, but you can delete a disabled job.

If you uninstall only one option, then no changes occur for the existing archive jobs. You can continue to edit and run the archive jobs as usual.

If you reinstall one or both options, all previously archived data is available if you specify the same Backup Exec installation folder path that you used in the initial installation. You can rerun any disabled jobs if no changes were made to the Backup Exec Database. Otherwise, if you try to run the jobs again, the jobs fail.

If you want to reinstall the Archiving Option and use new databases, you must reinstall both options. However, all previously archived data becomes unavailable. Archiving selections, rules, and archiving configurations are also unavailable.

You must remove the following items before you can reinstall the Archiving Option with new databases:

- The Enterprise Vault databases that are included in the Archiving Option. If SQL Express is used, these databases are located by default at `\Program Files\Symantec\Backup Exec\Data`. If the full version of SQL Server is used, the archive databases are in the path that the administrator defined.

- The BEAODatabase.xml file, which is located at \Program Files\Symantec\Backup Exec.
- The index files, which are located by default at \Program Files\Symantec\Backup Exec\ArchiveIndex.
- The temporary cache, which by default is located at \Program Files\Symantec\Backup Exec\AOCache.

If you reinstall one or both options and you use a different Backup Exec service account than you used for the initial installation, then you must use the Backup Exec Services Manager to update the service account after the reinstallation is complete.

See [“Uninstalling Backup Exec options from the local Backup Exec server”](#) on page 131.

See [“About the temporary cache location on the Backup Exec server”](#) on page 1400.

About installing Enterprise Vault on a Backup Exec server on which the Archiving Option is installed

If you install Enterprise Vault on a Backup Exec server on which the Archiving Option is installed, all archiving functionality is unavailable. Archive jobs that are active when Enterprise Vault is installed run to completion, but scheduled archive jobs do not run.

All archive-related jobs display the status as **Disabled**. You cannot run, edit, or save a disabled job, but you can delete a disabled job.

If you subsequently uninstall Enterprise Vault, archiving functionality remains unavailable.

About Enterprise Vault services for the Archiving Option

Symantec Enterprise Vault technology is the foundation for the Archiving Option. When you install the Archiving Option, some Enterprise Vault services are also installed. The Enterprise Vault services that run on the Backup Exec server use the same credentials as the Backup Exec service account.

The following Enterprise Vault services are installed on the Backup Exec server:

- Enterprise Vault Admin Service
- Enterprise Vault Directory Service
- Enterprise Vault Indexing Service
- Enterprise Vault Storage Service

- Enterprise Vault Task Controller Service

You must always use the Backup Exec Services Manager on the Backup Exec server to update your Backup Exec credentials. The Backup Exec Service Manager automatically updates the Enterprise Vault service credentials with the same credentials.

Note: Use of the Windows Services applet to edit the credentials of an Enterprise Vault service or a Backup Exec service is not supported. The use of this applet can leave the Archiving Option unsynchronized with the Backup Exec service account credentials. Errors can occur during the archiving operations.

See [“Changing the credentials for a service account”](#) on page 621.

Repairing the Backup Exec Archiving Option

If you have missing or corrupted Backup Exec Archiving Option files in an existing setup, or if there are specific installation failures during an upgrade, you can run the Repair option. The Repair option lets you provide service account credentials. If the Exchange Mailbox Archiving Option is configured, the Repair option lets you provide the DNS alias.

See [“Repairing Backup Exec”](#) on page 120.

How the Archiving Option works

To process an archive job, Backup Exec performs the following actions:

- Reads the associated backup of the file system and Exchange server from which you make archive selections.
- Applies the archive rules that you specify to identify the files and mail messages that are eligible for archiving.
- Checks if eligible files already exist in the archives.
If a file already exists in the archives, it is not archived again.
- Adds the data to the archives.
All of the archived data content is indexed to enable fast searching and retrieval of archived items.
- Deletes the archived files and mail messages from their original location.
Depending on an option that you specify, deletion occurs immediately after the archive job completes or after the vault store is backed up.

The Archiving Option operations that you can perform are described in the following table:

Table S-6 Operations that you can run

Operation	More information
Create an archive job to archive file system data and Exchange mail messages to vault stores.	See “About creating an archive job” on page 1340.
Create disk-based vault stores to use as storage for the archived data.	See “About vault stores in the Archiving Option” on page 1351.
Restore individual items from the archives.	See “About restoring items from the archives” on page 1374.
Delete individual items from the archives.	See “About deleting data from the archives” on page 1375.
Delete expired archive data from an archive automatically to free disk space, or ensure that archive data is never deleted from an archive.	See “Preventing the deletion of expired archived items from an archive” on page 1377.
Back up the Archiving Option components. These components include vault stores, vault store partitions, archives, databases, and index locations.	See “About backing up Archiving Option components” on page 1377.
Restore the Archiving Option components. These components include vault stores, vault store partitions, archives, databases, and index locations.	See “About restoring an Archiving Option component” on page 1382.
Synchronize the archive permissions with mailbox permissions, and share and folder permissions.	See “About synchronizing archive permissions and settings” on page 1350.

See [“Best practices for the Archiving Option”](#) on page 1339.

See [“Editing default settings for the Archiving Option”](#) on page 1346.

Types of data not included in archive jobs

The Archiving Option does not include some types of data in archive jobs.

Table S-7 The types of data that are not included in archive jobs

Archive Option	Types of data
File System Archiving Option	<p>The following types of data are not included in a file system archive job:</p> <ul style="list-style-type: none">■ Hard links■ Files with alternate streams■ Reparse points■ Sparse files■ Files in Microsoft Distributed File System Replication (DFSR) shares, system folders, or recycle bins■ Files that have an encrypted, hidden, or system attribute■ Files that are in mount point directories You can share the root of the mount point target, and then select it for archiving.
Exchange Mailbox Archiving Option	<p>The following types of data are not included in an Exchange mailbox archive job:</p> <ul style="list-style-type: none">■ Mail messages that have pending reminders.■ Any Exchange items other than mail messages, such as address book entries and calendar items.■ Mail messages in Exchange managed folders, journal mailboxes, or in public folders.

Best practices for the Archiving Option

Following are best practices when you use the Archiving Option:

- Use the default full recovery model for the SQL Server instance that hosts the Backup Exec Database and the Archiving Option databases. All Archiving Option databases that are created on the SQL Server are then also created with the full recovery model.
- Do not use different Backup Exec servers to archive files or mailboxes from the same file server or Exchange server.
- Consider archiving a smaller amount of data at first, such as a mailbox or a folder. All backup data may be eligible when you run the first archive job. Over

a period of time, the amount of eligible archive data lessens, and it becomes a predictable amount.

- Ensure that an item is included in only one archive job. Unlike backup jobs, archive jobs cannot share the same items.
- Ensure that all subdirectories are included in only one archive job.
- Do not archive the system drive. The Archiving Option does not archive system files.
- If you redirect the restore of Archiving Option components to a new server because of hardware failure, redirect the restore of the Directory database first. After the redirected restore of the Directory database is complete, you must run some additional tasks in a separate program called Backup Exec Utility. The tasks in the Backup Exec Utility update the Directory database with the new locations of the components. You should run the Backup Exec Utility tasks before you redirect the restore of any other Archiving Option components.

About creating an archive job

You create an archive job by adding an archive stage to a backup job. You can add the archive stage when you create the backup job, or you can edit an existing backup job and add an archive stage. When you add the archive stage, you choose the file system shares and folders and Exchange mailboxes that you want to archive. You can set other archive job options when you create the archive stage. You can also edit the default settings that apply to all archive jobs.

You can create an archive job for multiple servers using one of the following options:

- On the Backup and Restore page, select multiple servers and select **Backup to Disk and then Archive**. While editing the archive stage, user can make archive selections from all the selected servers.
- Select the **Edit Backups** option to add additional servers to an existing backup definition. Selection from the additional servers becomes available while editing the archive stage. User can make archive selections from the newly added servers.

Note: Archive is not available for server groups.

See [“Setting archive job options”](#) on page 1341.

See [“Viewing the servers that have archiving jobs”](#) on page 1341.

See [“Editing default settings for the Archiving Option”](#) on page 1346.

See [“About the Archiving Option”](#) on page 1313.

Viewing the servers that have archiving jobs

You can view a list of all of the servers for which you have created archiving jobs.

See [“About creating an archive job”](#) on page 1340.

To view the servers that have archiving jobs

- ◆ Do either of the following:

To view archiving jobs for servers from the **Backup and Restore** tab Do the following in the order listed:

Backup and Restore tab

- On the **Backup and Restore** tab, multi-select any number of servers.
- Right-click the servers, and then click **Details**.
- In the **Views** group, click **List**, and then click **Sort and Filter**.
- Click **Filter**, and then select **Job Type**.
- Click **Enable this filter**, and then click **Archive**, and then click **OK**.

To view archiving jobs for servers from the **Storage** tab Do the following in the order listed:

Storage tab

- On the **Storage** tab, double click a vault store.
- On the left, click **Archives**.
- View the list of servers in the **Server** column.

Setting archive job options

You can choose the file system shares and folders and Exchange mailboxes that you want to archive, and set other options for the archive job.

See [“About creating an archive job”](#) on page 1340.

Setting archive job options

- 1 Create a backup with an archive stage, or edit a stage that is part of a backup job.

See [“Backing up data”](#) on page 158.

See [“Editing a stage”](#) on page 221.

- 2 On the Archive box in the backup definition, click **Edit**.

3 On the **Archive Options** dialog box, do any of the following:

- To set a time and frequency for when to run the archive job

Click **Schedule**.
- To configure Backup Exec to notify specified recipients when the archive job completes

Click **Notification**.

See ["Notification options for jobs"](#) on page 305.
- To configure file system shares and folders to archive

Click **File System Selections**.

See ["File System Selections options for archive jobs"](#) on page 1342.
- To select Exchange servers or mailbox databases for archival

Click **Exchange Selections**.

See ["Exchange Selections options for archive jobs"](#) on page 1343.
- To create a new mailbox group

Click **Create Mailbox Group**.

See ["Mailbox Group options for archive jobs "](#) on page 1344.
- To specify the storage that you want to use for the archive job.

Click **Storage and Settings**.

See ["Storage and Settings options for archive jobs"](#) on page 1345.
- You can also specify the system mailbox to use if you archive Exchange mailboxes.

4 Click **OK**.

File System Selections options for archive jobs

You can select the folders or shares where you want Backup Exec to find data to archive. You can apply the same archive settings to all of the selections, or apply different archive settings to different selections.

See ["Setting archive job options"](#) on page 1341.

Table S-8 File System Selections options for archive jobs

Item	Description
Show administrative shares	Displays the administrative shares from which you can select the files and folders where you want Backup Exec to find data to archive.

Table S-8 File System Selections options for archive jobs (*continued*)

Item	Description
Add to Include List	Lets Backup Exec search the share selection or the folder selection for the eligible data to include in the archive stage.
Add to Exclude List	Lets Backup Exec search the share selection or the folder selection for the eligible data to exclude from the archive stage.
Remove from List	Lets you remove the share or folder from the selections list.
Share or Folder Path	Displays the location of the share selection or the folder selection.
Include/Exclude	Lets you specify if you want to include in or exclude the shares and folders from the archive stage.
Archive Setting	Lets you select the retention category and archive rules to apply to specific share and folder selections.
New Archive Settings	Lets you create a retention category and archive rules to apply to the selections. See “About archive settings in the Archiving Option” on page 1363.

Exchange Selections options for archive jobs

You can select Exchange servers or mailbox databases for archival.

See [“Setting archive job options”](#) on page 1341.

Table S-9 Exchange Selections options for archive jobs

Item	Description
Select the Exchange servers or mailbox databases that you want to archive	Lists the configured Exchanger servers and mailbox databases.
Selected Exchange Items	Lists the selected Exchanger servers and mailbox databases.

Table S-9 Exchange Selections options for archive jobs (*continued*)

Item	Description
Add to list	Adds the selected Exchanger servers and mailbox databases to the items that are selected for archival.
Remove from list	Removes the selected Exchanger servers and mailbox databases from the items that are selected for archival.

Mailbox Group options for archive jobs

You can create mailbox groups for users in an Exchange server domain. You can then select these mailbox groups for archival. While creating the mailbox group, you can specify the retention rules and archive rules.

See [“Setting archive job options”](#) on page 1341.

Table S-10 Mailbox Group options for archive jobs

Item	Description
Select domain to create new mailbox group	Lists the domain of the available Exchange servers.
New Mailbox Group	Opens a new window to create a new mailbox group to archive. See “Mailbox group options for mailbox group details” on page 1361.
Mailbox Group	Displays the names of the mailbox groups that you can select for archiving. See “About Exchange mailbox groups in archive jobs” on page 1358.
Users	Lists the names of the users whose mailboxes are included in a mailbox group for archiving.
Virtual Vault	Indicates if Virtual Vault is enabled for this mailbox group.
Retention Category	Indicates the retention category that is applied to the mailbox group. A retention category specifies the period of time for which you want to keep items in the archives.

Table S-10 Mailbox Group options for archive jobs (*continued*)

Item	Description
Retention Period	Indicates the amount of time that the item stays in the archives.
Items Older Than	Indicates if the items that are older than a specified period are included or excluded from the archive job.
File Size Greater Than	Indicates if the items have a size that is greater than a specified amount are included or excluded from the archive job.
Messages With Attachments	Indicates if messages with attachments are included or excluded from the archive job.
Unread Messages	Indicates if the messages that have not been read are included or excluded from the archive job.

Storage and Settings options for archive jobs

When you add an archive stage to a backup job, you must specify a vault store where Backup Exec stores the archived data. If you archive Exchange mailboxes, you must also specify a system mailbox.

See [“Setting archive job options”](#) on page 1341.

Table S-11 Storage and Settings options for archive jobs

Item	Description
Server	Displays the name of the server that is selected for the archive job.
Vault Store Name	<p>Displays the vault store where Backup Exec stores the archived data.</p> <p>You must specify a vault store to the server if one is not already specified. If you change the specified vault store, the change only affects the mailboxes or shares that you archive after you reassign the vault store. The shares and mailboxes that already have an archive in the previously specified vault store continue to be archived to that same archive.</p> <p>If there is no vault store, you must create one.</p>

Table S-11 Storage and Settings options for archive jobs (*continued*)

Item	Description
System Mailbox	<p>Displays the name of the system mailbox on the Exchange Server for Backup Exec to log on to. If a system mailbox is not assigned, you must assign one.</p> <p>This option only appears if you choose to archive Exchange mailboxes.</p> <p>See “Requirements for the Exchange Mailbox Archiving Option” on page 1316.</p>
New Vault Store	<p>Lets you create a new vault store.</p> <p>See “About vault stores in the Archiving Option” on page 1351.</p>

Editing default settings for the Archiving Option

You can edit the default settings for the Archiving Option.

See [“About the Archiving Option”](#) on page 1313.

To edit default settings for the Archiving Option

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, click **Archive**.
- 3 Select the appropriate options.
See [“Default settings for the Archiving Option ”](#) on page 1346.
- 4 Click **OK**.

Default settings for the Archiving Option

You can use the default settings that Backup Exec creates during installation for the Archiving Option. Or, you can choose your own default settings.

See [“Editing default settings for the Archiving Option”](#) on page 1346.

Table S-12 Default settings for the Archiving Option

Item	Description
Synchronize archive permissions and mailbox group members every day at	<p>Indicates the time at which Backup Exec synchronizes the correct archive settings and archive permissions to each archived mailbox in all mailbox groups.</p> <p>At this time, the archive permissions are also synced for each archived file share and folder.</p> <p>The default time is 3:00 A.M.</p> <p>See “About synchronizing archive permissions and settings” on page 1350.</p>
To run synchronization now, click Run Now	<p>Starts the synchronization task immediately.</p>
Delete archived items that have expired retention periods every day at	<p>Indicates the time at which Backup Exec searches the vault store partitions to delete the archived items that have expired retention periods.</p> <p>The default time is 4:00 A.M.</p> <p>For individual archives, you can prevent Backup Exec from automatically deleting the expired archived items.</p> <p>See “Preventing the deletion of expired archived items from an archive” on page 1377.</p>

Table S-12 Default settings for the Archiving Option *(continued)*

Item	Description
Enable single instance storage of archived items	<p>Lets Backup Exec identify the shareable parts of an item, such as a message attachment or the contents of a document. Backup Exec then stores the parts separately, and only once. When Backup Exec identifies a shareable part that is already stored in a vault store, it references the stored shareable part instead of archiving it again.</p> <p>Enabling this option can provide a significant reduction in the storage space that is required for archived items.</p> <p>See “About single instance storage of archived items” on page 1349.</p> <p>If you enable this option, you should back up the fingerprint databases. Single instance storage-related information is contained in the fingerprint databases for all of the vault stores.</p>
Default retention category	<p>Displays the retention category that applies to the Backup Exec archive jobs by default. A retention category specifies the period of time for which you want to keep items in the archives.</p> <p>You can edit a retention category to change the retention period.</p> <p>The default retention category specifies a retention period of infinite.</p> <p>See “About retention categories for archived items” on page 1366.</p>
Temporary cache location for synchronization of archived emails with Virtual Vault clients	<p>Displays the location where copies of new items that are added to the archive on the network are held.</p> <p>See “About the temporary cache location on the Backup Exec server” on page 1400.</p>

Table S-12 Default settings for the Archiving Option *(continued)*

Item	Description
Maximum cache size	Displays the maximum size of the temporary cache location. The default is 10 GB. See “About the temporary cache location on the Backup Exec server” on page 1400.
Retention categories	Lets you manage retention categories. See “About retention categories for archived items” on page 1366.
Archive Settings	Lets you manage archive settings. See “About archive settings in the Archiving Option” on page 1363.
Mailbox Groups	Lets you manage mailbox groups. See “About Exchange mailbox groups in archive jobs” on page 1358.
Index Locations	Lets you manage index locations. See “About managing index locations in the Backup Exec Archiving Option” on page 1370.

About single instance storage of archived items

Single instance storage of the archived items lets Backup Exec identify the shareable parts of an item. An example of a shareable part is a message attachment or the contents of a document. Backup Exec then stores the parts separately, and only once. When Backup Exec identifies a shareable part that is already stored in a vault store, it references the stored shareable part instead of archiving it again.

If single instance storage is enabled, items are shared within and across vault stores and vault store partitions. The vault store partitions may be on different storage. Shareable parts of a message that exceed the single-instance threshold of 20 KB are shared. These parts include attachments and message bodies. The user information and the shareable parts that are under the single instance storage threshold are not shared.

Enabling this option can provide a significant reduction in the storage space that is required for archived items. If you enable single instance storage, you should back up the fingerprint databases. Single instance storage-related information is contained in the fingerprint databases for all of the vault stores.

See [“Enabling single instance storage of archived items”](#) on page 1350.

Enabling single instance storage of archived items

You can enable single instance storage of archived items.

See [“About single instance storage of archived items”](#) on page 1349.

To enable single instance storage of archived items

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, click **Archive**.
- 3 Ensure that **Enable single instance storage of archived items** is selected.

See [“Default settings for the Archiving Option ”](#) on page 1346.

About synchronizing archive permissions and settings

Backup Exec runs a daily synchronization task for the Archiving Option.

Synchronization associates the correct archive settings to each mailbox in all of the mailbox groups. This task also ensures that archive permissions are synchronized with the mailbox permissions for each mailbox that is archived. For the File Systems Archiving Option, archive permissions are synchronized with share and folder permissions for each file that is archived.

For Exchange 2010 support, the synchronization task synchronizes Exchange server membership of any Database Availability Groups (DAG) that you select for archiving. The synchronization task detects if Exchange servers are added to a selected DAG. If so, then the vault store and the system mailbox that are associated with an existing DAG member server are automatically associated to the added server. If necessary, you can modify the system mailbox association by editing the archive stage for which this DAG is selected for backup.

The job log for the synchronization task displays the following information:

- All of the archives and folders for which permissions are synchronized.
- Any Exchange servers that are added to a DAG.
- The vault stores and system mailboxes that are automatically associated with any new Exchange servers that are added to a DAG.
- Any Exchange servers that are removed from a DAG.

You can specify the time of day to run this operation. Symantec recommends that you schedule the archive jobs to run at a different time than the synchronization operation.

An alert is sent to the administration console when the synchronization operation completes. The alert displays the summary statistics of the operation, and contains a link to the operation's job log.

Note: The Exchange servers and file system servers must be online and accessible by the Backup Exec server for synchronization to occur.

See [“Editing default settings for the Archiving Option”](#) on page 1346.

About vault stores in the Archiving Option

A vault store is a disk-based container for the archived data that Backup Exec archives from one server. When you add an archive stage to a backup job, you select a vault store as the storage device to which you want to send the archived data. A vault store contains at least one vault store partition that is the physical location where the archived items are stored. You can create additional vault store partitions for a vault store when it requires more disk space.

See [“About vault store partitions in the Archiving Option”](#) on page 1354.

Each vault store has an associated database. The database holds information about the archives in the vault store and all the items that are stored in each archive. For example, when an item is archived, the vault store's database is updated with this information. Single instance storage-related information is contained in the fingerprint databases for all of the vault stores.

The following vault store properties let you manage the deletion of archived items:

- Delete an item from its original location on the server immediately after it is archived, or after the vault store is backed up.
See [“About deleting archived data from its original location”](#) on page 1363.
- Delete the archived items that have expired retention periods from specific archives in vault stores.
See [“Preventing the deletion of expired archived items from an archive”](#) on page 1377.

You can back up vault stores and their associated databases, along with other Archiving Option components.

You must use the Configure Storage wizard to create a vault store.

See [“Editing or viewing vault store properties”](#) on page 1352.

See [“About deleting an Archiving Option vault store”](#) on page 1353.

Editing or viewing vault store properties

You can edit or view vault store properties.

See [“About vault stores in the Archiving Option”](#) on page 1351.

To edit or view vault store properties

- 1 On the **Storage** tab, double-click the vault store for which you want to view properties.
- 2 On the left, click **Properties**.
- 3 Edit the information as appropriate.

See [“Vault store properties”](#) on page 1352.

Vault store properties

You can edit the vault store properties.

See [“Editing or viewing vault store properties”](#) on page 1352.

Table S-13 Vault store properties

Item	Description
Name	<p>Displays the name of the vault store where Backup Exec stores archived data. You can edit this field to change the name.</p> <p>Note: Do not name a vault store with the same name that is already in use by an Archiving Option component, such as Fingerprint Databases or All Partitions. A vault store that has the same name as another Archiving Option component can cause errors when you make backup selections. Backup job or restore job failures can also occur.</p>
Description	<p>Displays a description of the vault store. You can edit this field to change the description.</p>

Table S-13 Vault store properties (*continued*)

Item	Description
State	<p>Displays the following states:</p> <ul style="list-style-type: none">■ Available Archive jobs can send data to this vault store.■ Being deleted The vault store is in the process of being deleted. Archive jobs cannot send data to this vault store.■ In backup mode A backup or restore job is running for the vault store.
Database name	<p>Displays the name of the database that is associated with this vault store.</p>
Archive count	<p>Displays the number of archives that the vault store contains.</p>
Total size	<p>Displays the total size of all of the items that are archived in the vault.</p>
Item deletion method	<p>Designates when to delete archived items from their original locations.</p> <p>You can choose to delete the item immediately after it is archived or delete it after the vault store is backed up. If you delete an item immediately after it is archived, the item is deleted from its original location after the archive job has successfully completed.</p> <p>See “About deleting archived data from its original location” on page 1363.</p>

About deleting an Archiving Option vault store

You can delete a vault store if one of the following conditions apply:

- It is not assigned to any servers on which you have made archive selections.
- It is the only vault store and you have deleted all of the archive stages that send archive data to it.

When you delete a vault store, you cannot cancel or undo the operation.

When you delete a vault store, all partitions, archives, and archived items in that vault store are also deleted. You must reassign another vault store to all of the servers that were assigned to the vault store that you want to delete.

See [“Deleting a vault store”](#) on page 1354.

Deleting a vault store

You can delete a vault store from Backup Exec.

See [“About deleting an Archiving Option vault store”](#) on page 1353.

To delete a vault store

- 1 On the **Storage** tab, right-click the vault store that you want to delete.
- 2 Click **Delete**.
- 3 If no other vault stores exist, you must do one of the following:
 - Delete all of the existing archive stages before you can delete this vault store.
See [“Deleting scheduled jobs”](#) on page 263.
 - Create a new vault store, assign it to all of the affected archived servers, and then delete the selected vault store.

About vault store partitions in the Archiving Option

A vault store partition represents the physical location where the archived items are stored. A vault store can contain one or more vault store partitions. Backup Exec creates one vault store partition in each vault store by default.

As the data in a vault store grows, you can create more vault store partitions to provide additional capacity. You can specify a local drive or a network share as a location for a vault store partition. You cannot specify a path that is a subdirectory in the path for another vault store partition. You must use the Configure Storage wizard to create a vault store partition. The Configure Storage wizard is on the **Storage** tab, in the **Configure** group.

A vault store can contain many vault store partitions, but only one partition is open at a time. As data is archived, it is stored in the open partition. You can specify a vault store partition as open or closed by editing the partition properties.

You can restore archived items from closed partitions, as well as delete the archived items that are in closed partitions.

Backup Exec searches the vault store partitions daily to delete the archived items that have expired retention periods. You can specify the time at which this daily operation runs.

See [“Preventing the deletion of expired archived items from an archive”](#) on page 1377.

See [“Editing vault store partition properties”](#) on page 1355.

Editing vault store partition properties

You can change the state of a vault store partition to open or closed. You can also edit the name and description of a vault store partition.

See [“About vault store partitions in the Archiving Option”](#) on page 1354.

To edit vault store partition properties

- 1 On the **Storage** tab, expand the vault store that contains the vault store partition that you want to edit.
- 2 Right-click the vault store partition, and then click **Details**.
- 3 Edit the appropriate information.

See [“Vault store properties”](#) on page 1352.

Vault store partition properties

A vault store partition represents the physical location where the archived items are stored. You can create a new vault store partition or change the state of an existing vault store partition.

See [“Editing or viewing vault store properties”](#) on page 1352.

Table S-14 Vault store partition properties

Item	Description
Name	Displays the name of the vault store partition.
Description	Displays a description of the vault store partition.

Table S-14 Vault store partition properties (*continued*)

Item	Description
Location	<p>Displays the path name where the vault store partition is located.</p> <p>The path can be on a local drive or on a network share. You cannot specify a path that is a subdirectory in the path for another vault store partition.</p> <p>For example, you can create a vault store partition on C:\vault store 1. However, you cannot create another vault store partition on C:\vault store 1\vault store 2.</p> <p>Ensure that the Backup Exec service account has full permissions for the path.</p> <p>See “Changing the credentials for a service account” on page 621.</p>
State	<p>Displays one of the following states:</p> <ul style="list-style-type: none">■ Open New archived data is stored in this vault store partition.■ Closed New archived data cannot be stored in this vault store partition.

About archives in the Archiving Option

An archive is a logical group of archived items. Items in an archive are stored in different vault store partitions depending on which partition is open at the time that the item is archived. Each archived file system share has its own archive, and each archived Exchange mailbox has its own archive. Backup Exec creates the archives when it runs an archive stage.

You cannot back up archives. You can only back up the vault store partitions.

See [“Editing archive properties”](#) on page 1356.

See [“Deleting an archive”](#) on page 1357.

Editing archive properties

You can edit archive properties.

See [“About archives in the Archiving Option”](#) on page 1356.

To edit archive properties

- 1 On the **Storage** tab, expand the vault store that contains the archive for which you want to edit properties.
- 2 Double-click the archive.
- 3 In the left pane, click **Properties**.

See [“Archive properties”](#) on page 1357.

Deleting an archive

You can delete an archive. However, if you delete an archive from Backup Exec, all of the archived data in the archive is also deleted.

See [“About archives in the Archiving Option”](#) on page 1356.

To delete an archive

- 1 On the **Storage** tab, double-click the vault store that contains the archive that you want to delete.
- 2 Right-click the archive, and then click **Delete**.
- 3 When you are prompted if you want to delete the archive, click **Yes**.

Archive properties

You can view archive properties. You can also edit the setting to let Backup Exec automatically delete the archived items that have expired retention periods.

See [“Editing archive properties”](#) on page 1356.

Table S-15 Archive properties

Item	Description
Resource name	Displays the name of the file share or Exchange mailbox that is archived.
Resource type	Displays one of the following types of archive: <ul style="list-style-type: none">■ File share■ Exchange mailbox

Table S-15 Archive properties (*continued*)

Item	Description
Status	Displays one of the following statuses as appropriate: <ul style="list-style-type: none">■ Available■ Being created■ Being deleted
Server	Displays the name of the server on which the archive is stored.
Automatically delete archived items that have expired retention periods	Lets Backup Exec delete the archived items that have expired retention periods from the archives. You can set a time at which Backup Exec deletes these items daily. Disable this option if you do not want archived items to be automatically deleted from this archive. This option is enabled by default.

About Archiving Option operation entries in the audit log

Audit logs provide information about the operations that have been performed in Backup Exec.

You can view information about archiving operations for the following:

- Vault stores
- Vault store partitions
- Archive settings
- Retention categories

See [“Configuring audit logs”](#) on page 624.

About Exchange mailbox groups in archive jobs

A mailbox group contains the selections on the Exchange Server that you want to archive. A mailbox group consists of user mailboxes to which you want to assign

the same archive settings. For example, you can add a single user, or a user group, or a deliverable list (DL) to a mailbox group.

Backup Exec applies the archive settings sequentially to each mailbox group in the list. The archive settings of the first mailbox group that a mailbox is found in are applied to that mailbox.

The order of the mailbox groups is important. You should arrange the mailbox groups that have specific selections of users, groups, and distribution lists at the top of the list. Arrange the mailbox groups that contain the least specific selections at the bottom of the list. For example, a mailbox group that contains specific users should be listed before a mailbox group that contains a user group. For example, you want to ensure that the correct archive settings are applied to the users that are in multiple groups.

You would arrange the following example mailbox groups in the order listed:

- The Managers group contains individual user accounts and requires all messages to be archived.
- The Some Users group contains some users in an organizational unit and requires messages to be archived from the last two months.
- The All Users group contains the entire Exchange organizational unit and requires messages to be archived from the last six months.

You can select the following items to archive in a mailbox group:

- Distribution lists
- User groups
- Users

You can create mailbox groups when you create an archive stage for Exchange Server mailboxes, or at any time from the global defaults.

See [“Managing Exchange mailbox groups”](#) on page 1359.

Managing Exchange mailbox groups

You can configure and manage mailbox groups for archive jobs for the Exchange Mailbox Archiving Option.

See [“About Exchange mailbox groups in archive jobs”](#) on page 1358.

To manage Exchange mailbox groups

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, click **Archive**.

- 3 Click **Manage Mailbox Groups**.
- 4 Enter the necessary information as appropriate.
See [“Manage mailbox groups options”](#) on page 1360.

Manage mailbox groups options

You can configure or edit the mailbox groups that contain the selections for an archive job.

See [“Managing Exchange mailbox groups”](#) on page 1359.

Table S-16 Manage mailbox group options

Item	Description
Domain	Displays the domains that you can select.
Mailbox Group	Displays the names of the mailbox groups in the selected domain that this stage archives. If there are no mailbox groups, you must create one before you can run an archive job.
Retention Category	Displays the retention category that the mailbox group is assigned to.
Virtual Vault Status	Indicates if Virtual Vault is enabled for this mailbox group. See “About enabling and disabling Virtual Vault” on page 1401.
Add	Lets you create a mailbox group to add to an archive job. See “Mailbox group options for mailbox group details” on page 1361.
Edit	Lets you edit the selected mailbox group. See “Mailbox group options for mailbox group details” on page 1361.
Delete	Deletes the selected mailbox group from the list of mailbox groups to archive.

Table S-16 Manage mailbox group options (*continued*)

Item	Description
Move Up	Moves the selected mailbox group up in the list. The retention category and the archive rules apply to the mailbox groups in the order in which the mailbox groups are listed. A mailbox that belongs to multiple groups is archived according to the archive settings of the highest-level group that it is in. See “About Exchange mailbox groups in archive jobs” on page 1358.
Move Down	Moves the selected mailbox group lower in the list.

Mailbox group options for mailbox group details

For an archive job, create a mailbox group that contains the selections on the Exchange servers that you want to archive. You can also specify the retention category and archive rules for each group.

See [“Managing Exchange mailbox groups”](#) on page 1359.

Table S-17 Mailbox group options for mailbox group details

Item	Description
Name	Designates the name of the mailbox group.
Users	Lets you specify the users that you want to include in this mailbox group.
Enable Virtual Vault for this mailbox group	Enables the Virtual Vault for the specified mailbox group. Note: If you make changes to Virtual Vault, you must notify your end users that they must restart Outlook before changes can take effect.

Table S-17 Mailbox group options for mailbox group details (*continued*)

Item	Description
Retention category	<p>Lets you specify the retention category for the mailbox group.</p> <p>The default setting is the default retention category, which has a retention period of infinite.</p> <p>See “About retention categories for archived items” on page 1366.</p>
New Retention Category	<p>Lets you create a new retention category.</p> <p>See “Retention category properties” on page 1368.</p>
Retention period	<p>Indicates how long to retain the items in the archives.</p>
Description	<p>Displays a description of the retention category.</p>
Archive items that are older than	<p>Indicates that the items that are older than the specified time are archived.</p> <p>The default setting is six months.</p>
If the items are larger than	<p>Indicates if the items that are larger than the specified size and are older than the specified time are archived, then archive those items after a specified number of weeks.</p> <p>You should archive larger mail messages more often than other messages. Specify a lesser amount of time in this option than in the previous Are older than option.</p> <p>The default setting is 10 MB and two weeks.</p>
Archive only messages with attachments	<p>Indicates that the messages with attachments are archived.</p> <p>The option is enabled by default.</p>
Archive unread messages	<p>Indicates that the messages that have not been read are archived.</p>

About deleting archived data from its original location

When you create a vault store, you can specify when to delete the archived data from its original location.

You can let Backup Exec do one of the following:

- Delete the item from its original location immediately after it is archived.
If the data is lost before the vault store is backed up, the only version of the data is on the backup set.
- Delete it after the vault store is backed up and the next archive job runs.

If Backup Exec deletes an item immediately after it is archived, the item is deleted from its original location after the archive job has successfully completed. If the item is modified after it is archived but before it is backed up, then it is not deleted from its original location.

See [“Editing or viewing vault store properties”](#) on page 1352.

About archive settings in the Archiving Option

Archive settings let you apply the following criteria:

- The retention category that specifies how long to keep the data in the archives.
- The rules that determine if data is eligible for archiving.

You can create archive settings for the following selections:

- File system shares
- File system folders within the shares

Note: You can name each group of archive settings that you create.

See [“Editing default settings for the Archiving Option”](#) on page 1346.

Manage Archive Settings options

You can add an archive setting or edit an archive setting to apply a different retention category and different rules to the archived data. Changes that you make here affect all of the selections to which the archive setting applies.

See [“About archive settings in the Archiving Option”](#) on page 1363.

See [“Editing default settings for the Archiving Option”](#) on page 1346.

Table S-18 Manage Archive Settings options

Item	Description
Archive Settings	Displays the name of the archive settings to apply to the file system selections.
Retention Category	Displays the name of the retention category.
Retention Period	Displays how long the items are retained in the archives for this retention category.
Description	Displays a description of the retention category.
Add	Lets you create an archive setting. See “Archive Settings Details” on page 1364.
Edit	Lets you edit an existing archive setting. See “Archive Settings Details” on page 1364.
Delete	Lets you delete an existing archive setting.
Default Folder Archive Settings	Displays the current settings for the default archive settings. See “Archive Settings Details” on page 1364.

Archive Settings Details

You can specify a retention category and archive rules to apply to data.

See [“Manage Archive Settings options”](#) on page 1363.

Table S-19 Archive settings options

Item	Description
Archive settings	Specifies the name of the archive settings to apply to file system selections. You can apply these same archive settings to other selections.
Name	Specifies the name of the retention category to apply to the selections.

Table S-19 Archive settings options (*continued*)

Item	Description
New Retention Category	Lets you create a new retention category. See “About retention categories for archived items” on page 1366.
Retention period	Displays how long the items are retained in the archives for this retention category.
Description	Displays any description that was entered for this retention category.
Name	Displays the name for this archive rule.
Action	Indicates if the filter for this archive rule is set to include or exclude files when archiving.
File Type	Indicates the type of file that this archive rule applies to.
Not accessed in	Indicates if the files that have not been accessed in a specified number of days are included or excluded from the archive stage. The default is to include the files that have not been accessed in 30 days in the archive job.
Not modified in	Indicates if the files that have not been modified in a specified number of days are included or excluded from the archive stage.
Not created in	Indicates if the files that have not been created in a specified number of days are included or excluded from the archive stage.
File size	Indicates if the files that are greater than or equal to or less than a specified size are included or excluded from the archive stage. The default is to include the files that are greater than or equal to 10 MB in the archive job.
Add	Lets you add a rule to the list of rules in the archive settings. This rule applies when you run the archive job for the file system selections. See “Archive Rule options” on page 1369.

Table S-19 Archive settings options (*continued*)

Item	Description
Edit	Lets you edit an existing rule in the archive settings.
Delete	Deletes a rule from the list of rules in the archive settings.
Move Up	Moves a rule up in the list of rules. An item is archived according to the first rule for which it meets the criteria. The top rule in the list is the first rule that applies.
Move Down	Moves a rule down in the list of rules.

About retention categories for archived items

Use retention categories to specify the period of time for which you want to keep items in the archives. You can give the retention categories meaningful names, such as Business or Personal. Retention categories make it easier for you to retrieve items because you can search for them by their category name. Each retention category has a retention period, which indicates how long you want to retain the items that are archived with this retention category.

For example, you can create a retention category named Finance Data Retention and set it to retain archived data for seven years.

The retention period starts on the date that the item is archived. Backup Exec runs a daily operation that deletes all items that have expired retention periods. You can prevent this operation from running on specific archives.

See [“Preventing the deletion of expired archived items from an archive”](#) on page 1377.

You cannot delete retention categories. You can edit retention categories, including the retention periods.

Changes that you make to a retention category apply to the following:

- All items to which the retention category is already applied.
- Any new items to which you apply the retention category.
See [“Adding or editing a retention category ”](#) on page 1367.

You can create retention categories as needed when you create an archive job. You can also specify a retention category to use as the default setting for all archive jobs. If you do not specify a retention category, then a default retention category with a retention period of infinite is applied to an archive job.

See [“About creating an archive job”](#) on page 1340.

See [“Adding or editing a retention category ”](#) on page 1367.

See [“Editing default settings for the Archiving Option”](#) on page 1346.

Adding or editing a retention category

You can add a retention category or edit an existing retention category. Changes apply to existing archived items as well as to new items to which you apply the retention category.

See [“About retention categories for archived items”](#) on page 1366.

To add or edit a retention category

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, click **Archive**.
- 3 Click **Retention Categories**.
- 4 Do one of the following:

To add a retention category

Click **Add**, edit the appropriate information, and then click **OK**.

See [“Retention category properties”](#) on page 1368.

To edit a retention category

Do the following in the order listed:

- Select the retention category that you want to edit, and then click **Edit**.
- Edit the appropriate information, and then click **OK**.

See [“Retention category properties”](#) on page 1368.

- 5 Click **OK** on the Manage Retention Category dialog box, and click **OK** on the Archive Settings dialog box.

Manage Retention Category options

You can add a retention category or edit a retention category to change the period of time for which you want to keep items in the archive. Changes that you make here affect all of the selections to which the retention category applies.

See [“Adding or editing a retention category ”](#) on page 1367.

Table S-20 Manage Retention Category options

Item	Description
Retention Category	Displays the name of the retention category.
Retention Period	Displays how long the items are retained in the archives for this retention category.
Description	Displays a description of the retention category.
Add	Lets you create a retention category. See “Retention category properties” on page 1368.
Edit	Lets you edit an retention category. See “Retention category properties” on page 1368.

Retention category properties

You can specify the period of time for which you want to keep items in the archives.

See [“Adding or editing a retention category”](#) on page 1367.

Table S-21 Retention category properties

Item	Description
Retention category	Displays the name of the retention category.
Infinite	Retains the item in the archives for an infinite amount of time. The retention period starts from the date that the item is archived.
For a period of	Retains the item in the archives for a specified period of time. The retention period starts from the date that the item is archived.
Description	Displays a description of the retention category.

Adding or editing archive rules

See [“About retention categories for archived items”](#) on page 1366.

You can add or edit archive rules in the archive settings that apply to the archive jobs.

To add or edit archive rules

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 On the left, click **Archive**.
- 3 Click **Archive Settings**.
- 4 Select an archive setting, and then click **Add** or **Edit**.
- 5 Under Archive Rules, do one of the following:

To add new archive rules Click **Add**.

To edit archive rules Select a rule, and then click **Edit**.

- 6 Edit the appropriate options.
See [“Archive Rule options”](#) on page 1369.
- 7 Click **OK**.

Archive Rule options

You can configure the rules that specify the characteristics of the data to include or exclude in the archive job.

See [“Adding or editing archive rules”](#) on page 1368.

Table S-22 Options for archive rules

Item	Description
Name	Displays the name of the archive rule.
Include when archiving	Specifies that the files that meet the requirements that you select are included in the archive job. This option is enabled by default.
Exclude from archiving	Specifies that the files that meet the requirements that you select are excluded from the archive job.
Type of files	Specifies the types of file to include or exclude from the archive job. You can type your own rule or use a predefined rule.

Table S-22 Options for archive rules (*continued*)

Item	Description
Files not accessed in	Includes or excludes the files that have not been accessed in a specified number of days. The default is to include the files that have not been accessed in 6 months in the archive job.
Files not modified in	Includes or excludes the files that have not been modified in a specified number of days.
Files not created in	Includes or excludes the files that have not been created in a specified number of days.
File size	Includes or excludes the files that are greater than or equal to, or less than or equal to a specified size. The default is to include the files that are greater than or equal to 10 MB in the archive job.

About managing index locations in the Backup Exec Archiving Option

Index locations store all of the archived data content that is indexed to enable fast searching and retrieval of archived items.

You can manage index locations by doing the following:

- **Add new index locations.**
Add index locations to another disk or to a remote computer that has more space rather than moving the indexes to a new location. If you want a particular location to be used for indexing, close all of the other index locations and leave one location open.
- **Open and close index locations.**
Close an index location to prevent new indexes from being created at that location. You can re-open index locations when appropriate. When you close a location, new items for existing indexes continue to use the closed location. New indexes, such as those for new archives, are created in the open index location. They do not use the closed location.
- **Delete index locations.**
Delete empty index locations. An index location becomes empty when all of the archived items that had index information in this location have expired.

See [“Viewing index locations”](#) on page 1371.

See [“Adding a new index location”](#) on page 1371.

See [“Opening or closing an index location”](#) on page 1372.

See [“Deleting an index location”](#) on page 1372.

Viewing index locations

You can view index locations.

See [“About managing index locations in the Backup Exec Archiving Option”](#) on page 1370.

To view index locations

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 On the left, click **Archive**.
- 3 Click **Index Locations**.
- 4 View the existing index locations.
See [“Manage Index Locations options”](#) on page 1373.
- 5 Click **Close**.

Adding a new index location

You can add index locations. For the best performance, do not place the index location paths on the same disks as vault partition paths.

Note: The index location must be on an NTFS drive, or on a UNC path on an NTFS network share. Mapped drives are not supported. The Backup Exec service account must be granted the NTFS Read and Write permissions for the index location path. The index location cannot be at the root of a volume. You must specify a path that has at least one folder name.

See [“About managing index locations in the Backup Exec Archiving Option”](#) on page 1370.

To add a new index location

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 On the left, click **Archive**.

- 3 Click **Index Locations**.
- 4 In the **Index location path** field, enter a new location, and then click **Add**.
- 5 Click **Close**.

Deleting an index location

An index location must be empty before you can delete it. To empty an index location, use the Backup Exec Utility to move all index files to a different index location.

See [“Running the Backup Exec Utility to complete Archiving Option operations”](#) on page 1386.

You cannot delete an index location if it is the only index location.

See [“About managing index locations in the Backup Exec Archiving Option”](#) on page 1370.

To delete an index location

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 On the left, click **Archive**.
- 3 Click **Index Locations**.
- 4 Select the index location that you want to delete.
See [“Manage Index Locations options”](#) on page 1373.
- 5 Click **Delete**.
- 6 When you are prompted to delete the index location, click **Yes**.
- 7 Click **Close**.

Opening or closing an index location

You can close an index location to prevent new indexes from being created at that location. You can still add items to closed indexes. You can also perform searches on the index locations that are closed.

Note: Ensure that at least one index location is open. Archive jobs must be able to store index information for archived items. Without an open index location, archive jobs fail.

See [“About managing index locations in the Backup Exec Archiving Option”](#) on page 1370.

To open or close an index location

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, click **Archive**.
- 3 Click **Index Locations**.
- 4 Select the index location that you want to open or close.
See [“Manage Index Locations options”](#) on page 1373.
- 5 Click **Set Open** or **Set Close**, as appropriate.
- 6 Click **Close**.

Manage Index Locations options

You can manage the Archiving Option index locations.

See [“Adding a new index location”](#) on page 1371.

Table S-23 Manage Index Locations options

Item	Description
Index location path	Indicates the path where you want to store the archive index files. Note: The index location must be on an NTFS drive, or on a UNC path on an NTFS network share. Mapped drives are not supported. The Backup Exec service account must be granted the NTFS Read and Write permissions for the index location path. The index location cannot be at the root of a volume. You must specify a path that has at least one folder name.
Browse	Lets you browse the network and select an NTFS folder where you want to store the index files.
Index Locations	Displays all of the index locations that have been created or configured.

Table S-23 Manage Index Locations options (*continued*)

Item	Description
Status	<p>Displays the state of the index location.</p> <p>Possible values include the following:</p> <ul style="list-style-type: none">■ Open■ Closed
Backup Mode	<p>Displays Yes if the index location is currently being backed up; otherwise, displays No.</p> <p>When an index location is in backup mode, no new index data is written to that location. You can continue to search index locations that are in backup mode.</p>
Add	<p>Lets you add new index locations.</p> <p>Note: For best performance, do not add an index location to the same disk on which vault partitions are located.</p> <p>See “Adding a new index location” on page 1371.</p>
Delete	<p>Lets you delete an index location.</p> <p>See “Deleting an index location” on page 1372.</p>
Set Open	<p>Lets you open an index location.</p> <p>See “Opening or closing an index location” on page 1372.</p>
Set Close	<p>Lets you close an index location.</p> <p>See “Opening or closing an index location” on page 1372.</p>

About restoring items from the archives

You can perform the following restore operations for archived items:

- Restore files to their original locations or to another file server.
- Restore mail messages to the original mailbox or to another mailbox on the Exchange server.

Note: In the Restore Wizard, if you redirect the restore to another Exchange server, you must add two backslash characters (\\) in front of the Exchange server name. For example, if the Exchange server name is 'ExchangeServer01', you must enter '\\ExchangeServer01'. Otherwise, you cannot continue to the next page in the Restore Wizard.

The mailbox must already exist on the server to which you want to restore the mail messages.

The archives can contain multiple versions of the same item. To restore a specific version of the item, you must select it individually. Otherwise, Backup Exec restores the latest version of an item. You can distinguish between versions of the same file by checking the modified time of an item.

Note: NTFS permissions for archived data are not restored.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

About searching for data in the archives

You can search the archives to find and select the data in the archives. The archives contain access control restrictions, but these restrictions are not applied when you search from the Backup Exec Administration Console. The search displays all archived versions of the data.

You can specify criteria such as content, retention categories, or retention periods. You can also restrict the search to one archive or to all archives that are associated with a server. You can use wildcard characters. Use a question mark (?) to represent any single character. Use an asterisk (*) to represent any number of characters.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

About deleting data from the archives

You can delete archived files and mail messages from the archives. If you need to free some disk space, you can delete items from the archives before their retention periods expire.

The archives can contain multiple versions of the same item. To delete a specific version of the item, you must select it individually. Otherwise, Backup Exec deletes the latest version of an item. You can distinguish between versions of the same file by checking the modified time of an item.

You can delete only files and mail messages from the archives. To delete an entire archive, you must delete it from the **Storage** tab.

Additionally, Backup Exec searches the vault store partitions daily to delete the archived items that have expired retention periods. You can specify the time at which this daily operation runs.

See [“Editing default settings for the Archiving Option”](#) on page 1346.

See [“Deleting an archive”](#) on page 1357.

See [“Deleting data from the archives”](#) on page 1376.

Deleting data from the archives

You can delete specific data from the archives.

See [“About deleting data from the archives”](#) on page 1375.

To delete data from the archives

1 On the **Backup and Restore** tab, double-click the server that you archived.

2 In the **Archive** group, click **Delete from Archive**.

You can also select multiple servers, and use the **Delete from Archive** option to delete data from the archives in the selected servers.

3 Edit the options as appropriate.

See [“Delete from Archive options ”](#) on page 1376.

4 Click **OK**.

Delete from Archive options

You can select the data that you want to delete from the archives.

See [“Deleting data from the archives”](#) on page 1376.

Table S-24 Delete from Archives options

Item	Description
Name	Specifies the name of the server that the data was archived from.
Select the data that you want to delete	Lets you check the archived data that you want to delete.
Name	Displays the name of the data that you select to delete from the archives.

About backing up and restoring the Archiving Option components from a remote Backup Exec server

You can back up and restore the Archiving Option components from a remote Backup Exec server on which licenses are not installed. You can also edit the backup job default settings for the Archiving Option components.

The remote Backup Exec server that you use to back up the Archiving Option components does not require licenses for the following options:

- File System Archiving Option
- Exchange Mailbox Archiving Option

You must provide the credentials of the Backup Exec service account on the Backup Exec server on which the Archiving Option is installed.

See [“About backing up Archiving Option components”](#) on page 1377.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

See [“Changing default backup job settings”](#) on page 515.

Preventing the deletion of expired archived items from an archive

Backup Exec deletes the archived items that have expired retention periods from a specific archive. You can clear this option to prevent Backup Exec from deleting expired archived items.

See [“About archives in the Archiving Option ”](#) on page 1356.

To prevent the deletion of expired archived items from an archive

- 1 On the **Storage** tab, expand the vault store that contains the archive.
- 2 Double-click the archive.
- 3 In the left pane, click **Properties**.
- 4 Disable the option **Automatically delete archived items that have expired retention periods**.
- 5 Click **OK**.

About backing up Archiving Option components

You can select any or all of the Archiving Option components for backup. If you select all of the components for backup in the same job, recovery time is faster.

However, if you create multiple backup jobs for the components, the backup jobs run faster.

The Archiving Option components that you can back up are described in the following table, along with recommendations for backup:

Table S-25 Backing up Archiving Option components

Component	Description
Archiving Option Components	Archiving Option Components contain all of the components that are associated with the Archiving Option. Symantec recommends that you select Archiving Option Components to back up all of the Archiving Option environment.
Backup Exec Archiving Site	The Backup Exec Archiving Site is a logical representation of an installation of the Archiving Option. A Backup Exec server can have only one Archiving Site. If you select this component for backup, the Directory database is also automatically backed up.
Directory database	<p>The Directory database is a Microsoft SQL Server database that contains configuration data and information about the archives.</p> <p>After the database is populated, the amount of data in the Directory database changes very little over time.</p> <p>You should back up the Directory database after you add or remove any Archiving Option component. You should also back up the Directory database if you change the location of any component. Configuration changes can include creating vault stores, creating vault store partitions, and changing vault store partition statuses.</p> <p>See “About disabling backup mode for Archiving Option components” on page 1381.</p>

Table S-25 Backing up Archiving Option components (*continued*)

Component	Description
Index location	<p>The index location stores all of the archived data content that is indexed to enable fast searching and retrieval of archived items. The indexing data is stored in index files in the location that is specified when the Archiving Option is installed.</p> <p>You should back up the index location on a regular basis.</p>
Vault store group	<p>The vault store group is a logical entity. If you select it for backup, all of the vault databases, vault store partitions, and the fingerprint databases are backed up. Because these components are closely related, you should consider selecting the vault store group to back up all of these components together.</p>
Fingerprint databases	<p>The fingerprint databases contains the single instance storage-related information for all of the vault stores in the vault store group.</p> <p>If you enable single instance storage of archived items, you should back up the fingerprint databases on a regular basis.</p> <p>See “About single instance storage of archived items” on page 1349.</p>
Vault store	<p>The vault store is a logical entity. If you select it for backup, all of the vault databases and the vault store partitions are backed up.</p>

Table S-25 Backing up Archiving Option components (*continued*)

Component	Description
All partitions	<p>A vault store partition represents the physical location where the archived items are stored. A vault store can contain one or more vault store partitions. If you select All Partitions for backup, then all of the vault store partitions in the vault store are selected for backup.</p> <p>Note: When you back up an open partition, the vault store database is automatically backed up.</p> <p>You should back up the vault store partitions on a regular basis.</p> <p>See “About vault store partitions in the Archiving Option” on page 1354.</p>
Vault store databases	<p>The vault store databases are the Microsoft SQL Server databases that contain configuration data and information about the archives. Each vault store has an associated database. Each vault store database contains an entry for each item that is archived in the associated vault store. If an item is deleted from the archive, then the references to it are deleted from the vault store database.</p> <p>You should back up the vault store databases on a regular basis.</p>

You can also back up and restore the Archiving Option components from a remote Backup Exec server on which licenses are not installed.

See [“Archive Option backup options”](#) on page 1380.

See [“About backing up and restoring the Archiving Option components from a remote Backup Exec server”](#) on page 1377.

See [“Backing up data”](#) on page 158.

See [“About consistency checks for Archiving Option databases”](#) on page 1381.

Archive Option backup options

The Archive Option page is available only on the Backup Exec server computer where the Backup Exec Archiving Option is installed. Use the Archive Option backup

options to back up the archiving option databases, vault stores, and other components.

The **Archive Option** is available on the **Backup Options** dialog box only when you select **Archiving Option Components** on the **Backup Definition Properties** dialog box. The following options are available for the Archive Option Component backup jobs.

Table S-26 Archive Option backup options

Component	Description
Full Backup method	For full backup, only one option is available. Full - Back up components This option is selected by default. You cannot change this selection.
Incremental Backup method	For incremental backup, following options are available: <ul style="list-style-type: none">■ Full - Back up components■ Differential - Back up component changes since the last full■ Incremental - Back up component changes since the last full or incremental

See [“About backing up Archiving Option components”](#) on page 1377.

See [“Backup methods in Backup Exec”](#) on page 188.

About consistency checks for Archiving Option databases

Backup Exec automatically checks the physical consistency of an Archiving Option database before a backup job and after a restore job. Any consistency check failures are reported in the Backup Exec job log. Backup Exec uses Microsoft SQL Server's utility Physical Check Only for consistency checks of Archiving Option databases.

For more information about the Physical Check Only utility, see the Microsoft SQL Server documentation.

About disabling backup mode for Archiving Option components

When you back up the Directory database, ensure that the Archiving Option components are not in backup mode.

See [“Editing or viewing vault store properties”](#) on page 1352.

If a component is in backup mode, you must take it out of backup mode by running the task **Disable Backup Mode on Archiving Option entities** in the Backup Exec Utility.

See [“Running the Backup Exec Utility to complete Archiving Option operations”](#) on page 1386.

About restoring an Archiving Option component

You can restore any of the following Archiving Option components:

- Directory database
- Vault store databases
- Fingerprint databases
- Vault store partition
- Index location

Note: The procedures for redirecting restore jobs include running tasks in Backup Exec Utility to update the new locations of the restored components.

See [“About running the Backup Exec Utility to complete redirected restores of Archiving Option components”](#) on page 1383.

Review the scenarios in the following table to find the best procedure to restore an Archiving Option component.

Table S-27 Methods for restoring an Archiving Option component

Method	More information
If a data loss occurs, and you want to restore an Archiving Option component to the same location	See “Methods for restoring data in Backup Exec” on page 232.
If hardware fails and a data loss occurs, and you want to restore an Archiving Option component to a different location	See “About moving Archiving Option components to a new location” on page 1387.
If you want to move components to new hardware, such as to a new SQL Server or a new disk	See “About moving Archiving Option components to a new location” on page 1387.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

See [“About consistency checks for Archiving Option databases”](#) on page 1381.

About running the Backup Exec Utility to complete redirected restores of Archiving Option components

You can run a redirected restore job if you want to restore one or more Archiving Option components to a different location. After the redirected restore job completes, you must run additional tasks in Backup Exec Utility to update the new locations of the restored components.

The following table lists the possible scenarios and the associated redirected restore solutions for Archiving Option components.

Table S-28 Redirected restore solutions for Archiving Option components

Scenario	Solution
The SQL Server that hosts the databases fails and a data loss occurs.	<p>Do the following in the order listed:</p> <ul style="list-style-type: none">■ Redirect the restore of the Archiving Option databases to a new SQL Server. See “Methods for restoring data in Backup Exec” on page 232.■ Run the Change Database Location task in the Backup Exec Utility.■ Run the Update SQL Server Name task in the Backup Exec Utility. <p>See “Running the Backup Exec Utility to complete a redirected restore of the Archiving Option databases” on page 1384.</p>
The local drive or network share that hosts the vault store partition fails and a data loss occurs.	<p>Do the following in the order listed:</p> <ul style="list-style-type: none">■ Redirect the restore of a vault store partition to a different path in a local drive or network share. See “Methods for restoring data in Backup Exec” on page 232.■ Run the Change Vault Partition Path in the Backup Exec Utility. <p>See “Running the Backup Exec Utility before redirecting the restore of an Archiving Option vault store partition” on page 1385.</p>

Table S-28 Redirected restore solutions for Archiving Option components
(continued)

Scenario	Solution
The disk that contains the index files fails and a data loss occurs.	<p>Do the following in the order listed:</p> <ul style="list-style-type: none">■ Redirect the restore of the index files to a new location. See “Methods for restoring data in Backup Exec” on page 232.■ Run the Change Index Location task in the Backup Exec Utility. See “Running the Backup Exec Utility before redirecting the restore of Archiving Option index files” on page 1385.

Running the Backup Exec Utility to complete a redirected restore of the Archiving Option databases

After you redirect the restore of the Archiving Option databases to a new SQL Server, you must run the **Change Database Location** task in Backup Exec Utility.

See [“About running the Backup Exec Utility to complete redirected restores of Archiving Option components”](#) on page 1383.

To run the Backup Exec Utility to complete a redirected restore of the Archiving Option databases

- 1 When the redirected restore job is complete, start the **Backup Exec Utility**.
See [“Running the Backup Exec Utility to complete Archiving Option operations”](#) on page 1386.
- 2 In the **Backup Exec Utility** task pane, under **Archiving Option Tasks**, click **Change Database Location**.
- 3 In **Destination SQL server instance**, type the name of the new SQL Server.
- 4 Click **OK**.
- 5 In the **Backup Exec Utility** task pane, under **Archiving Option Tasks**, click **Update SQL Server Name**.
- 6 After the operation completes, exit the Backup Exec Utility.

Running the Backup Exec Utility before redirecting the restore of an Archiving Option vault store partition

You can redirect the restore of a vault store partition to a different path in a local drive or network share.

If you restore a vault store partition that has an **Open** status, its vault store database is automatically restored.

See [“About running the Backup Exec Utility to complete redirected restores of Archiving Option components”](#) on page 1383.

If a vault store partition needs more disk space, you can create a new partition.

To run the Backup Exec Utility before redirecting the restore of an Archiving Option vault store partition

- 1 Start the **Backup Exec Utility**.

See [“Running the Backup Exec Utility to complete Archiving Option operations”](#) on page 1386.

- 2 In the **Backup Exec Utility** task pane, under **Archiving Option Tasks**, click **Change Vault Partition Path**.
- 3 Select the name of the vault store partition.
- 4 In **New Vault Store Partition Path**, type the new path to which you want to restore the vault store partition.
- 5 Ensure that **Move Vault Store Partition Files** is not selected.
- 6 Click **OK**.
- 7 On the Backup Exec Administration Console, start the Restore Wizard to restore the vault store partition.

See [“Methods for restoring data in Backup Exec”](#) on page 232.

Running the Backup Exec Utility before redirecting the restore of Archiving Option index files

You can redirect the restore of the index files to a new location.

Note: You must locate the index files on a local NTFS drive.

See [“About restoring an Archiving Option component”](#) on page 1382.

To run the Backup Exec Utility before redirecting the restore of Archiving Option index files

- 1 Start the **Backup Exec Utility**.
See [“Running the Backup Exec Utility to complete Archiving Option operations”](#) on page 1386.
- 2 In the **Backup Exec Utility** task pane, under **Archiving Option Tasks**, click **Change Index Location for the Archiving Option**.
- 3 Select the index location that you want to modify.
The application displays the selected index location path in the **New Index Location** field.
- 4 In the **New Index Location** field, type the new path to which you want to restore the index files.
- 5 Ensure that **Move Index Files** is not selected.
- 6 Click **OK**.
- 7 On the Backup Exec Administration Console, create a restore job to restore the index files.
See [“Methods for restoring data in Backup Exec”](#) on page 232.

Running the Backup Exec Utility to complete Archiving Option operations

You must run the Backup Exec Utility to complete some operations for an Archiving Option component.

See [“About running the Backup Exec Utility to complete redirected restores of Archiving Option components”](#) on page 1383.

To run the Backup Exec Utility to complete Archiving Option operations

- 1 From the Backup Exec installation directory, double-click **BEUtility.exe**.
- 2 In the **Properties** pane, under **Archiving Option Tasks**, click the appropriate task.
- 3 Click **Help** for information about a task.

About moving Archiving Option components to a new location

You can use the Backup Exec Utility to move Archiving Option components to a new location. Ensure that no other archive-related operations are running when you move a component.

If you must move a component because the hardware that hosts the component has failed, you should use a redirected restore job.

See [“About running the Backup Exec Utility to complete redirected restores of Archiving Option components”](#) on page 1383.

Table S-29 Moving Archiving Option components to a new location

Component	More information
Index location	<p>You can move an index location if the disk where the index files are stored runs out of space.</p> <p>Use the Change Index Location task in Backup Exec Utility.</p>
Databases	<p>You can move databases to a different SQL Server. For example, you can move the databases if the current SQL Server becomes overloaded.</p> <p>Use the Change Database Location task in Backup Exec Utility.</p>
Vault store partitions	<p>You can move vault store partitions if you must remove the current drive or network share that contains the partition.</p> <p>Note: If a vault store partition only requires more disk space, you can create a new partition and designate it as open.</p> <p>Use the Change Vault Store Partition Path task in Backup Exec Utility.</p>

See [“Running the Backup Exec Utility to complete Archiving Option operations”](#) on page 1386.

Troubleshooting archive jobs

If there are issues with archive jobs, you can find information in the following sources:

- Backup Exec job logs.
See [“Viewing the job log”](#) on page 267.
- Enterprise Vault event log that is located in the Windows Event Viewer.
See [“Viewing the Enterprise Vault event log for Archiving Option events”](#) on page 1388.

An Exchange Mailbox Archiving Option job may not find data to archive for the following reasons:

- Only back up sets for which the Granular Recovery Technology option was enabled and that exist on disk storage devices can be archived.
- The associated Exchange mail stores may not be backed up, or the mailbox or user may have been deleted in the last 14 days.

See [“Requirements for both the Exchange Mailbox Archiving Option and the File System Archiving Option”](#) on page 1314.

Viewing the Enterprise Vault event log for Archiving Option events

You can view the Windows Event Viewer to review the Enterprise Vault event log for information about Archiving Option events. Enterprise Vault generates many log entries. You must take some action to make sure that the log files do not grow too large. For information on how to control the log file size, see the Windows Event Viewer help.

Reports for the Archiving Option

The reports in the following table are available to help you monitor your Archiving Option environment.

See [“Reports in Backup Exec”](#) on page 629.

Table S-30 Reports for the Archiving Option

Report	Description
Vault Store Usage Summary	Displays the archived items that are in each vault store and the total size of the vault store.
Vault Store Usage Details	Displays the archives that are in each store and the size of each archive.

Table S-30 Reports for the Archiving Option (*continued*)

Report	Description
File System Archive Settings	Displays the archive settings that are applied to archive selections for each server.
Exchange Mailbox Group Archive Settings	Displays the archive settings that are applied to mailbox groups in each domain.
Archive Selections by Archive Rules and Retention Categories	Displays the archive rules and retention categories that are applied to each archive selection.
Archive Job Success Rate	Displays the number of archive jobs that ran successfully.
Failed Archive Jobs	Displays a list of archive jobs that failed recently.
Overnight Archive Summary	Displays a summary of archive jobs for the last 24 hours.

About Backup Exec Virtual Vault

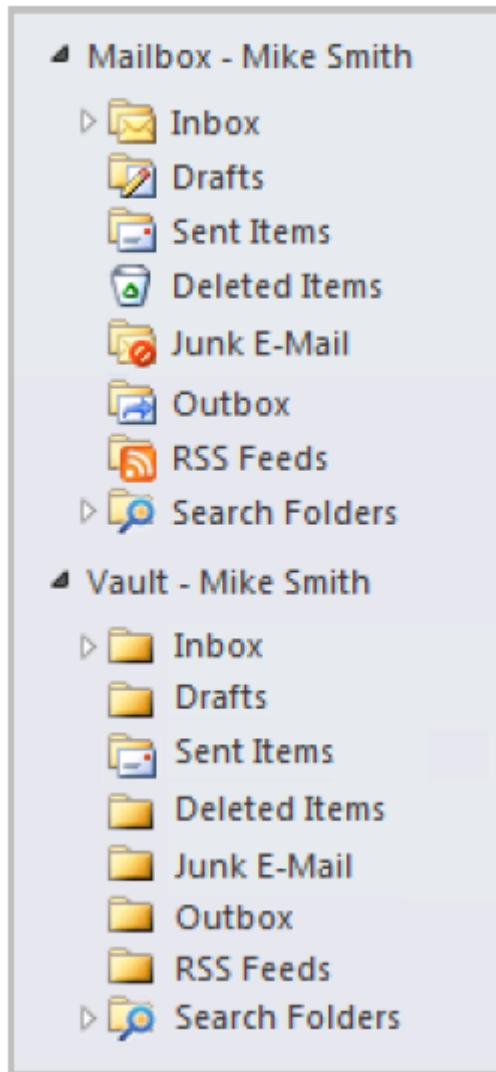
After an archive job runs on an end user's mailbox, the email messages are archived and stored in an archive on the network. Each archived Exchange mailbox has its own archive on the network. Typically, only the administrator can access the archives on the network. However, with Backup Exec's Virtual Vault feature, the end users can access their archived email messages from a vault on their Outlook navigation pane.

The Backup Exec administrator must make the Backup Exec Outlook Add-In (add-in) available to the end users to install on their computers. The add-in is a component of the Virtual Vault feature that runs on the end user's computer. The add-in maintains a cache that is called the Vault Cache on each end user's computer. The add-in regularly synchronizes the end user's Vault Cache with the mailbox archive on the network. The Vault Cache acts like a local copy of an end user's archived mailbox.

End users can view, search, read, forward, and reply to the email messages in their vaults. End users cannot move, delete, or rename a folder in their vaults. End users can also search the contents of their vaults by using Outlook Instant Search, Outlook Advanced Find, or Windows Desktop Search.

In the following example, the vault for the end user named Mark Smith appears as **Vault - Mark Smith**.

Figure S-1 An example of an end user's vault on the Outlook Navigation Pane



Backup Exec's Virtual Vault feature is automatically installed and enabled on the Backup Exec server when you install or upgrade the Backup Exec Exchange Mailbox Archiving Option. Virtual Vault is enabled for all mailbox groups by default. The Backup Exec administrator can disable the use of Virtual Vault for all end users that belong to a mailbox group.

See [“About installing the Backup Exec Outlook Add-In to end users' computers”](#) on page 1391.

See [“Best practices for Virtual Vault”](#) on page 1396.

See [“About using the command line to install the Backup Exec Exchange Mailbox Archiving Option”](#) on page 1333.

See [“About configuring Outlook in cached Exchange mode”](#) on page 1396.

See [“About enabling and disabling Virtual Vault ”](#) on page 1401.

About the Backup Exec Outlook Add-In

The Backup Exec Outlook Add-In is a component of Backup Exec's Virtual Vault. The add-in must be installed on the end users' computers.

The add-in maintains the end user's Vault Cache. The Vault Cache is a location that stores the archived email message headers and the email message contents that are downloaded from the Backup Exec server. These items are downloaded from the Backup Exec server when the Vault Cache synchronizes with the mailbox archives on the network. To synchronize the Vault Cache, the add-in connects to the Backup Exec server by using the DNS alias that you entered during the Archiving Option installation or upgrade.

End users can perform the following tasks by using the Backup Exec Outlook Add-In:

- Start synchronization between the Vault Cache and the Backup Exec server.
- Suspend or resume synchronization.
- Edit the size of the Vault Cache.
- Specify the amount of time that Backup Exec waits after Outlook starts before it starts synchronization.

The Backup Exec Outlook Add-In contains Help that the end users can access on their computers. The Help is located on the Outlook toolbar, in **Help > Backup Exec Help Topics**.

See [“About installing the Backup Exec Outlook Add-In to end users' computers”](#) on page 1391.

See [“About the Vault Cache”](#) on page 1397.

About installing the Backup Exec Outlook Add-In to end users' computers

Several methods are available to install the Backup Exec Outlook Add-In to end users' computers.

Table S-31 Installation methods for the Backup Exec Outlook Add-In

Installation method	Information
Copy the Backup Exec Outlook Add-In.msi file to a network share that end users can access. Then, send instructions to specific end users about how to install the add-in.	See “Installing the Backup Exec Outlook Add-In by copying it to a network share” on page 1395.
Use the command line to install the Backup Exec Outlook Add-In.msi file to remote computers, one at a time.	See “Installing the Backup Exec Outlook Add-In by using the command line (Silent Mode)” on page 1392.
Install the Backup Exec Outlook Add-In in an Active Directory network.	See “Installing the Backup Exec Outlook Add-In in an Active Directory network” on page 1393.
Use Microsoft's System Center Configuration Manager to deploy the Backup Exec Outlook Add-In to end users' computers.	See your Microsoft documentation for instructions on how to use System Center Configuration Manager, or search the Symantec Knowledge Base for specific instructions.

Installing the Backup Exec Outlook Add-In by using the command line (Silent Mode)

You can use the command line to install the Backup Exec Outlook Add-In to a single computer.

See [“About installing the Backup Exec Outlook Add-In to end users' computers”](#) on page 1391.

To install the Backup Exec Outlook Add-In by using the command line (Silent Mode)

- 1 Open a Windows command prompt.
- 2 Type the following command:

```
msiexec /i <path to the Backup Exec Outlook Add-In.msi package>  
/qn+ ALLUSERS=1
```

- 3 Optionally, to add logging, type the following command:

```
msiexec /i <path to the Backup Exec Outlook Add-In.msi package>  
/qn+ /lv* <Path to log file> ALLUSERS=1
```

Installing the Backup Exec Outlook Add-In in an Active Directory network

You can centrally manage the installation of the Backup Exec Outlook Add-In to computers in an Active Directory network. You configure the installation once, and then use a Group Policy Object to assign that installation to computers in an Organizational Unit. The add-in is installed automatically whenever a computer in the Organizational Unit is restarted.

See [“About installing the Backup Exec Outlook Add-In to end users' computers”](#) on page 1391.

Note: Review your organization's deployment plans before you implement a rollout of the Backup Exec Outlook Add-In to client computers. You should also review your Group Policy Desktop Management and Active Directory documentation.

Table S-32 Installing the Backup Exec Outlook Add-In in an Active Directory network

Step	Action	Description
Step 1	Create a distribution point (share) that contains the source file that you want to install.	You must copy the Backup Exec Outlook Add-In package to the distribution point. See “Creating a software distribution point to deploy the Backup Exec Outlook Add-In” on page 1394.
Step 2	Configure a Group Policy Object to assign the directory in the distribution point to computers in an Active Directory Organizational Unit.	The software is installed automatically when the computers in the Organizational Unit are restarted. Note: Computers in the Organizational Unit may need to be restarted twice if the Windows Fast Logon Optimization feature is enabled. See “Configuring a Group Policy Object for the Backup Exec Outlook Add-In” on page 1394.

Creating a software distribution point to deploy the Backup Exec Outlook Add-In

To install the Backup Exec Outlook Add-In in an Active Directory network, you must create a software distribution point (share).

See [“Installing the Backup Exec Outlook Add-In in an Active Directory network”](#) on page 1393.

Table S-33 Creating a software distribution point

Step	Description
Step 1	Create a shared folder, and then set read-only permissions for all of the client computers to which you want to deploy the package.
Step 2	Copy the Symantec Backup Exec Add-In.msi package to the shared folder.
Step 3	Configure a Group Policy Object to deploy the source files. See “Configuring a Group Policy Object for the Backup Exec Outlook Add-In” on page 1394.

Configuring a Group Policy Object for the Backup Exec Outlook Add-In

To install the Backup Exec Outlook Add-In in an Active Directory network, you must configure a Group Policy Object after you create a software distribution point.

See [“Creating a software distribution point to deploy the Backup Exec Outlook Add-In”](#) on page 1394.

Refer to your Microsoft Windows documentation for information on configuring a Group Policy Object.

Note: Symantec recommends that you create a new Group Policy Object and apply it only to the computers to which you want to deploy the Backup Exec Outlook Add-In.

To configure a Group Policy Object for the Backup Exec Outlook Add-In

- 1 From the Active Directory snap-in that manages users and groups, right-click the domain name, and then click **Properties**.
- 2 Create a new Group Policy Object.

- 3 Under **Computer Configuration**, expand **Software Settings**.
- 4 Right-click **Software Installation**, click **New**, and then click **Package**.
- 5 On the **File Open** dialog box, browse to the software distribution point by using the UNC name.
- 6 Select the package file `Symantec Backup Exec Outlook Add-In.msi`, and then click **Open**.
- 7 When you are prompted, click **Advanced**.
- 8 After Active Directory checks the MSI package, on the **General Properties** tab, ensure that the correct version of the option is selected for installation.
- 9 On the **Deployment** tab, set up the configuration for your environment.
- 10 Ensure that the option **Make this 32-bit x86 application available to WIN64 machines** is selected.
- 11 Close all of the dialog boxes.

Installing the Backup Exec Outlook Add-In by copying it to a network share

The Backup Exec Outlook Add-In is a component of the Virtual Vault feature, and must be installed on an end user's computer. You can copy the Backup Exec Outlook Add-In to a network share that is accessible to end users. Then, you can send instructions to the end users whom you want to install the Backup Exec Outlook Add-In.

You can copy and paste the instructions that are included in the following procedures in an email message to the end users.

See [“About installing the Backup Exec Outlook Add-In to end users' computers”](#) on page 1391.

To copy the Backup Exec Outlook Add-In to a network share

- 1 At the Backup Exec server, place the installation media in the appropriate drive.
- 2 Navigate to the following directory on the installation media:
`BE\WinNT\Install\EVE\Outlook Add-In`
- 3 Copy the .msi file to a network share to which end users also have access.
- 4 In the following instructions for end users, replace the text in brackets with the actual name of the network share to which you have copied the .msi file.
- 5 Send the following instructions to the end users whom you want to install the Backup Exec Outlook Add-In.

To install the Backup Exec Outlook Add-In

- 1 On your computer, close Outlook.
- 2 Navigate to the following folder:
<Enter the name of the network share to which you have copied the .msi file>
- 3 Double-click `Symantec Backup Exec Outlook Add-In.msi`.
- 4 Follow the instructions on your screen.
- 5 In Outlook, turn on cached Exchange mode.
For instructions on how to configure cached Exchange mode, refer to your Outlook Help.
- 6 In Outlook, turn off AutoArchive.
For instructions on how to turn off AutoArchive, refer to your Outlook Help.
- 7 Restart Outlook.

About configuring Outlook in cached Exchange mode

The end user must configure Outlook in cached Exchange mode to enable the vault to appear in the Outlook Navigation Pane, as well as to allow preemptive caching. Preemptive caching lets the items that are almost eligible for archiving to be archived in the Vault Cache directly from the Outlook offline folders. These items do not have to be downloaded from the Backup Exec server. The Backup Exec Outlook Add-In copies email messages from the end user's Outlook .OST file to the Vault Cache.

End users should see their Outlook Help for more information about how to set cached Exchange mode.

See [“Installing the Backup Exec Outlook Add-In by copying it to a network share”](#) on page 1395.

Best practices for Virtual Vault

After you install the Backup Exec Outlook Add-In for a few end users, monitor the Backup Exec server's performance to ensure that it performs within acceptable limits. For large organizations, it may be beneficial to phase the availability of the Virtual Vault feature to the end users.

You can limit the number of end users who can use Virtual Vault by doing one or both of the following actions:

- Enable Virtual Vault for specific mailbox groups.
See [“About enabling and disabling Virtual Vault”](#) on page 1401.

- Make the Backup Exec Outlook Add-In available only to specific end users.
See [“Installing the Backup Exec Outlook Add-In by copying it to a network share”](#) on page 1395.

About the Vault Cache

A Vault Cache is a location on the end user's computer that stores the archived email messages that are downloaded from the Backup Exec server. These items are downloaded from the Backup Exec server when the Vault Cache synchronizes with the archives on the network. The end user's vault on the Outlook Navigation Pane reflects the contents from the Vault Cache rather than the contents from the archive on the network.

By default, the Vault Cache is created in the following locations when the end user installs the Backup Exec Outlook Add-In:

- For Windows XP, in \Documents and Settings\<User>\Local Settings\Application Data\KVS\Enterprise Vault
- For Windows Vista/7, in \Users\<User>\AppData\Local\KVS\Enterprise Vault

The Outlook Reading Pane shows the selected email message header. If the email message is in the Vault Cache, the Reading Pane also shows the content. If only the email message header is shown but no content, the Reading Pane displays a link to the original email message.

The Vault Cache is limited to a maximum of 10% of free space on the end user's computer or to 1 GB, whichever is smaller. After the limit is reached, the oldest email messages are automatically deleted from the Vault Cache to make room for new email messages. Even though the archived email messages are deleted from the end user's Vault Cache, the messages are still kept in the archives on the network.

End users can synchronize another user's archive to their Vault Cache if they have the correct permissions. A separate vault is displayed in the end user's Outlook Navigation Pane for each archive on the network that is synchronized to the Vault Cache. Instructions on how to synchronize to other archives on the network is provided in the Backup Exec Outlook Add-In Help on the end users' computers.

Note: Vault Cache is not available to Microsoft Entourage users.

The end user can edit the settings in the Backup Exec Outlook Add-In to manage their Vault Cache.

See [“About the Backup Exec Outlook Add-In”](#) on page 1391.

See [“About Vault Cache synchronization”](#) on page 1398.

About Vault Cache synchronization

Vault Cache synchronization updates the end user's Vault Cache with any changes that are made to the archives on the network. These changes include the email messages that were created, updated, and deleted since the last synchronization.

The Vault Cache synchronization process consists of the following:

- **Header synchronization**
Updates the Vault Cache with information about archived email messages. The email message header contains enough information to enable the message to be represented in the end user's vault in the Outlook Navigation Pane. The email message header also contains information to associate the header with the content of the email message. Where changes have occurred in the archives on the network, Vault Cache synchronization downloads header information from the Backup Exec server. The changes are then applied to the Vault Cache. Header synchronization also synchronizes any changes that are made to the folder hierarchy in the archives on the network.
- **Content synchronization**
Downloads the archived email messages from one or more archives on the network to the Vault Cache.

After the end user installs the Backup Exec Outlook Add-In, the initial header synchronization automatically starts three minutes after Outlook starts. Content synchronization may also be performed. If the archive contains a large number of items, content synchronization takes much longer than header synchronization.

A vault is automatically added to an end user's Outlook profile when the following criteria are met:

- The Exchange archive job has processed all of the archives that the end user can access.
- The initial header synchronization has completed.
- The end user has not previously removed the vault from the profile.

After the Exchange Mailbox Archiving job runs and after the Vault Cache synchronization completes, the Vault Cache is current with the mailbox archives on the network.

Vault Cache synchronization can start in the following ways:

- The Backup Exec Outlook Add-In automatically synchronizes the Vault Cache with the archives on the network once a day. If the add-in cannot connect to Backup Exec, then it waits for five minutes before it attempts to contact the Backup Exec server again.
If a scheduled synchronization time is missed, the add-in attempts a synchronization the next time that the end user opens Outlook. By default, the

add-in waits three minutes after the end user opens Outlook before it attempts to retry a missed synchronization. This default can be changed by the end user.

- The end user can manually start the Vault Cache synchronization at any time. A manual synchronization does not affect the next scheduled time for automatic synchronization. Unlike a scheduled synchronization, a manual synchronization that fails is not retried.

End users may miss the scheduled Vault Cache synchronization on a weekend when they do not use Outlook. In this situation, a large number of header synchronization requests may occur around the same time on Monday. To avoid an excessive load on the Backup Exec server, only a limited number of header synchronization requests are accepted. Because of this limitation, scheduled synchronization may not process immediately for some end users. No error messages are displayed to end users if their synchronization is not processed immediately. Their header synchronization request is repeated every five minutes until it succeeds. When the synchronization succeeds, the daily scheduled synchronization time is reset to the time of the successful synchronization.

Backup Exec restricts the number of active content download requests at any given time. Backup Exec preconfigures this setting to control the amount of system resources that the Vault Cache content downloads use. The content download from the Backup Exec server to the Vault Cache uses Microsoft Background Intelligent Transfer Service technology.

End users can access the Backup Exec Outlook Add-In Help for information about how to view the synchronization status, and how to synchronize the Vault Cache.

See [“About Backup Exec Virtual Vault ”](#) on page 1389.

About preemptive caching

To minimize downloads to the Vault Cache, the Backup Exec Outlook Add-In regularly searches the Outlook .OST cache file for any email messages that are due to be archived soon. The Backup Exec Outlook Add-In automatically adds these email messages to the Vault Cache. This feature is called preemptive caching. Preemptive caching takes place on the end user's computer. It reduces the number of items that must be downloaded from the archive on the network to the Vault Cache when the two are synchronized.

The Backup Exec Outlook Add-In performs preemptive caching for email messages seven days before they are due for archiving as per the Exchange archiving settings for the mailbox.

The end user must configure Outlook in cached Exchange mode for preemptive caching to work. The Outlook Add-In copies email messages from the end user's

Outlook .OST file to the Vault Cache. If the end user's Outlook does not have a .OST file, then preemptive caching cannot be performed.

See [“About configuring Outlook in cached Exchange mode”](#) on page 1396.

About the temporary cache location on the Backup Exec server

If new items are added to the archive on the network, copies of these items are held temporarily in a cache on the Backup Exec server. The items are then downloaded to the end user's computer. When you install the Backup Exec Archiving Option, the cache location is set by default to a folder in the installation directory. Typically, this directory is `C:\Program Files\Symantec\Backup Exec\AOCache`.

Backup Exec notifies you if the cache location on the Backup Exec server reaches the maximum size that is specified. If you do not free some disk space, or specify another cache location, the oldest items in the cache are deleted to free some space.

You can change the location of the cache and increase the size of the cache.

See [“Changing the temporary cache location”](#) on page 1400.

See [“Changing the temporary cache size”](#) on page 1401.

Changing the temporary cache location

You can change the location of the temporary cache on the Backup Exec server. The temporary cache location is used to synchronize the archived email messages with the Vault Cache on the end users' computers. The temporary cache location must be on a local NTFS drive. Mapped drives are not supported. The temporary cache location cannot be at the root of a volume. You must specify a path that has at least one folder name. The Backup Exec service account must be granted the NTFS rights Read and Write for the temporary cache location.

See [“About the temporary cache location on the Backup Exec server”](#) on page 1400.

Note: You must restart the Backup Exec services before the changes can take effect.

To change the temporary cache location

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, click **Archive**.

- 3 Under **Virtual Vault Settings**, enter a new cache location.
- 4 Click **OK**.

Changing the temporary cache size

The default maximum size of 10 GB is used for the temporary cache on the Backup Exec server. The temporary cache location is used to synchronize the archived email messages with the Vault Cache on the end users' computers. You can increase the size of the temporary cache as needed.

See [“About the temporary cache location on the Backup Exec server”](#) on page 1400.

To change the temporary cache size

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, click **Archive**.
- 3 Under **Virtual Vault Settings**, in the **Maximum cache size** field, select a size for the cache.
- 4 Click **OK**.
- 5 Restart the Backup Exec services for the changes to take effect.

About enabling and disabling Virtual Vault

By default, Virtual Vault is enabled for all mailbox groups. The Backup Exec administrator can disable the use of Virtual Vault for all end users that belong to a mailbox group.

Enabling or disabling Virtual Vault on mailbox groups takes effect only after the Backup Exec server synchronizes the settings with the end users' mailboxes. This synchronization operation may take a few minutes. Then, end users must restart Outlook before changes to Virtual Vault can take effect.

If you disable Virtual Vault for end users, the files in the Vault Cache on the end users' computers are not removed automatically. The files can remain in the Vault Cache to be reused if you enable Virtual Vault again.

See [“Enabling and disabling Virtual Vault for mailbox groups”](#) on page 1401.

Enabling and disabling Virtual Vault for mailbox groups

You can enable or disable Virtual Vault for the end users that are in a mailbox group.

See [“About enabling and disabling Virtual Vault”](#) on page 1401.

To enable and disable Virtual Vault for mailbox groups

- 1 Click the Backup Exec button, select **Configuration and Settings**, and then select **Backup Exec Settings**.
- 2 In the left pane, click **Archive**.
- 3 Click **Mailbox groups**.
- 4 Select the domain, and then do one of the following:

To create a new mailbox group for which you want to enable or disable Virtual Vault Click **New**.

To enable or disable Virtual Vault for an existing mailbox group Select the mailbox group, and then click **Edit**.

- 5 Do one of the following:

To enable Virtual Vault for all of the end users in the mailbox group Ensure that **Enable Virtual Vault** is selected.

This option is enabled by default.

To disable Virtual Vault for all of the end users in the mailbox group Clear the **Enable Virtual Vault** check box.

- 6 Click **OK**.

Accessibility and Backup Exec

This appendix includes the following topics:

- [About accessibility and Backup Exec](#)
- [About keyboard shortcuts in Backup Exec](#)
- [List box navigation in Backup Exec](#)
- [Tabbed dialog box navigation in Backup Exec](#)
- [About setting accessibility options](#)

About accessibility and Backup Exec

Symantec products meet federal accessibility requirements for software as defined in Section 508 of the Rehabilitation Act:

<http://www.access-board.gov/508.htm>

Symantec products are compatible with operating system accessibility settings as well as a variety of assistive technologies. All manuals also are provided as accessible PDF files, and the online help is provided as HTML displayed in a compliant viewer.

Keyboard navigation is available for all graphical user interface operations and menu items. Backup Exec uses standard operating system navigation keys.

Items in the task pane that do not have keyboard shortcuts can be accessed by using the operating system's "mouse keys", which allow you to control the mouse through the numerical keyboard.

To see a table of the standard Microsoft navigation keys and keyboard shortcuts, select your version of Microsoft Windows from the table at:

<http://www.microsoft.com/enable/products/keyboard.aspx>

About keyboard shortcuts in Backup Exec

All menu items can be selected by using accelerator or mnemonic keyboard shortcuts. An accelerator is a key combination that provides shortcut access to a user interface function. A mnemonic (sometimes referred to as a "hotkey ") is a key equivalent (used in combination with the **Alt** key) for selecting user interface components such as menu items. The mnemonic "hotkey " letter is on the item in the user interface.

Select secondary menu items by opening the main menu and using the up or down arrow key until the item is highlighted. Press the right arrow key to open a submenu, and press **Enter** to select your choice.

Keyboard shortcuts are not case-sensitive. Mnemonic keystrokes may be pressed either sequentially or simultaneously. All menu items have mnemonics, but not all menu items have accelerators.

Routine functions such as opening, saving, and printing files can be performed using the standard Microsoft keyboard shortcuts. Other menu items are unique to Backup Exec.

The following table lists keyboard shortcuts to display labels and the contents of the Backup Exec button:

Table T-1 Keyboard shortcuts in Backup Exec

Accelerator	Mnemonic	Result
Alt	F10	Displays the mnemonic label for a tab that you can use in combination with the accelerator to display the tab.
Alt	A	Expands the Backup Exec button. Use the Backup Exec button to connect to the Backup Exec server, configure settings, and access installation and licensing items. You can also access Backup Exec documentation, technical support items, and various Symantec Web sites .

See "Home tab keyboard shortcuts" on page 1405.

See "Backup and Restore tab keyboard shortcuts" on page 1405.

See "Job Monitor tab keyboard shortcuts" on page 1413.

See [“Storage tab keyboard shortcuts”](#) on page 1416.

See [“Reports tab keyboard shortcuts”](#) on page 1424.

Home tab keyboard shortcuts

The following table lists the keyboard shortcuts for the **Home** tab.

See [“About keyboard shortcuts in Backup Exec”](#) on page 1404.

Table T-2 Home tab keyboard shortcuts

Accelerator	Mnemonic	Command	Result
Alt	H	Home tab	Opens the Home tab.
Alt	1	One Column	Displays the Home tab items in one column.
Alt	2	Two Columns	Displays the Home tab items in two columns.
Alt	NW	Narrow/Wide	Displays the Home tab items in two columns with a narrow panel and a wide panel.
Alt	3	Three Columns	Displays the Home tab items in three columns.
Alt	D	Reset Home Tab	Restores the contents of the Home tab to the default configuration.

Backup and Restore tab keyboard shortcuts

The following table lists the keyboard shortcuts for the **Backup and Restore** tab.

See [“About keyboard shortcuts in Backup Exec”](#) on page 1404.

See [“Backup and Restore tab keyboard shortcuts for the Jobs view”](#) on page 1408.

See [“Backup and Restore tab keyboard shortcuts for the Job History view”](#) on page 1409.

See [“Backup and Restore tab keyboard shortcuts for the Backup Sets view”](#) on page 1410.

See [“Backup and Restore tab keyboard shortcuts for the Active Alerts view”](#) on page 1412.

See [“Backup and Restore tab keyboard shortcuts for the Credentials view ”](#) on page 1412.

Table T-3 Backup and Restore tab keyboard shortcuts

Accelerator	Mnemonic	Command	Result
Alt	B	Backup and Restore tab	Opens the Backup and Restore tab.
Alt	ST	Standard	Displays Backup Exec in a view that provides descriptive text.
Alt	CO	Compact	Displays Backup Exec in a view that conserves space.
Alt	F	Sort and Filter	Displays information in a custom view or lets you create and save a custom view.
Alt	T	Tree	Displays items in a hierarchal view. This command is disabled for the list of servers on the Backup and Restore tab.
Alt	L	List	Displays items in a list that you can sort by columns. This command is disabled for the list of servers on the Backup and Restore tab.
Alt	G	Groups	Lets you view information by server group. You can add, remove, or edit a server group.
Alt	B	Backup	Defines backup jobs and settings to back up your data. You can back up data now, or schedule a time to do it.
Alt	O	One-Time Backup	Defines backup jobs and settings to run one time.

Table T-3 Backup and Restore tab keyboard shortcuts (*continued*)

Accelerator	Mnemonic	Command	Result
Alt	EB	Edit Backups	Lets you edit one or more existing backup jobs. You cannot edit synthetic or one-time backups.
Alt	CA	Backup Calendar	Lets you view all scheduled backup jobs on a calendar.
Alt	RE	Restore	Browses the backup sets from a single server, and then restores the data.
Alt	SE	Search	Searches for backup sets, and then restores the data, or copies and saves the search criteria.
Alt	DR	Create Disaster Recovery Disk	Launch the Simplified Disaster Recovery Disk Creation Wizard to help you create a recovery disk.
Alt	G	Convert to Virtual	Converts backup data to a virtual machine. You must run a full backup that includes all critical system components before you can convert backup data to a virtual machine.
Alt	AS	Add	Adds one or more servers to the list of servers. You must add servers to back them up and monitor them.
Alt	RS	Remove	Removes one or more servers from the list of servers. You can remove servers from the list if you no longer want to back them up.
Alt	US	Update	Updates the selected Backup Exec servers with the latest hot fixes and maintenance packs.

Table T-3 Backup and Restore tab keyboard shortcuts (*continued*)

Accelerator	Mnemonic	Command	Result
Alt	HA	Hold Job Queue	Pauses the job queue. Active jobs continue to run, but no new jobs run until the queue is taken off hold.
Alt	RN	Run Next Backup Now	Runs the next scheduled backup for the selected servers.

Backup and Restore tab keyboard shortcuts for the Jobs view

The following table lists the **Backup and Restore** tab keyboard shortcuts for the Jobs view.

See [“Backup and Restore tab keyboard shortcuts”](#) on page 1405.

Table T-4 Backup and Restore tab keyboard shortcuts for the Jobs view

Accelerator	Mnemonic	Command	Result
Alt	JE	Edit	Lets you edit a backup definition. You can edit a backup definition's backup selections, backup settings, and stages.
Alt	JD	Delete	Deletes a backup definition if you no longer need it.
Alt	JC	Cancel	Cancels an active job while it is running.
Alt	JP	Priority	Increases or decreases a job's priority in the job queue .
Alt	JR	Run Now	Runs a job immediately. If a job is scheduled it still runs at its scheduled time.
Alt	JH	Hold	Pauses a job or the entire job queue. The job or job queue resumes its normal schedule when you clear this option.

Table T-4 **Backup and Restore** tab keyboard shortcuts for the Jobs view
(continued)

Accelerator	Mnemonic	Command	Result
Alt	JT	Test Run	Runs a test of the selected backup job now.
Alt	JA	Job Activity	Lets you view statistical and system information about a running job. You can also cancel an active job.
Alt	HH	View Job History	Lets you view detailed information about the job history including individual job and job summary statistics.

Backup and Restore tab keyboard shortcuts for the Job History view

The following table lists the **Backup and Restore** tab keyboard shortcuts for the Job History view.

See [“Backup and Restore tab keyboard shortcuts”](#) on page 1405.

Table T-5 **Backup and Restore** tab keyboard shortcuts for the Job History view

Accelerator	Mnemonic	Command	Result
Alt	HR	Run Now	Runs a job immediately. If the job is scheduled, it still runs at its scheduled time.
Alt	HL	View Job Log	Lets you view the job log for the selected job history. The job log provides detailed job information, storage and media information, job options, file statistics, and job completion status.
Alt	HD	Delete	Deletes the selected job history and its associated job log, if you no longer need them.

Table T-5 **Backup and Restore** tab keyboard shortcuts for the Job History view *(continued)*

Accelerator	Mnemonic	Command	Result
Alt	HH	View Job History	Lets you view detailed information about the job history including individual job and job summary statistics.
Alt	EH	Error Handling	Defines the rules for handling failed jobs.
Alt	HC	Duplicate	Creates a duplicate copy of the job history, which includes all of a job's dependent backup sets. You can duplicate the job history now, or schedule a time to do it.
Alt	HV	Verify	Verifies the integrity of the collection of data and the media on which it resides for this job history. When you verify a job history, you verify all of a job's dependent backup sets. You can verify the job history now, or schedule a time to do it.
Alt	EB	Edit Backup	Lets you edit the backup for the selected job history. You cannot edit synthetic or one-time backups jobs.

Backup and Restore tab keyboard shortcuts for the Backup Sets view

The following table lists the **Backup and Restore** tab keyboard shortcuts for the Backup Sets view.

See [“Backup and Restore tab keyboard shortcuts”](#) on page 1405.

Table T-6 **Backup and Restore** tab keyboard shortcuts for the Backup Sets view

Accelerator	Mnemonic	Command	Result
Alt	SQ	Expire	Expires the selected backup set, if you no longer need it.
Alt	SH	Retain	Retains the selected backup set. You can prevent backup sets from expiring by retaining them.
Alt	SC	Catalog	Catalogs the selected backup set. Cataloging backup sets lets you view the data that is contained in them and search for files to restore.
Alt	SU	Duplicate	Creates a duplicate copy of the selected backup set. You can duplicate the backup set now, or schedule a time to do it.
Alt	SV	Verify	Verifies the integrity of the collection of data and the media on which it resides for this backup set. You can verify the backup set now, or schedule a time to do it.
Alt	SL	Expiration	Lets you to edit the expiration date and time.
Alt	SK	Show Dependent Backup Sets	Shows the dependent backup sets for the selected backup set.
Alt	SW	Restore	Browses the backup sets and then restores the data.

Backup and Restore tab keyboard shortcuts for the Active Alerts view

The following table lists the **Backup and Restore** tab keyboard shortcuts for the Active Alerts view.

See [“Backup and Restore tab keyboard shortcuts”](#) on page 1405.

Table T-7 **Backup and Restore** tab keyboard shortcuts for the Active Alerts view

Accelerator	Mnemonic	Command	Result
Alt	AR	Respond	Lets you view the alert and any additional information. You must select Respond OK to clear the alert.
Alt	AO	Respond OK	Clears the alert without displaying information about it. Respond OK only if you no longer need the alert.
Alt	AL	View Job Log	Lets you view the job log that was generated for this job. The job log provides detailed job information, storage and media information, job options, file statistics, and job completion status.
Alt	AH	Show Alert History	Displays alerts that you have responded to or that you have automatically cleared.

Backup and Restore tab keyboard shortcuts for the Credentials view

The following table lists the **Backup and Restore** tab keyboard shortcuts for the Credentials view.

See [“Backup and Restore tab keyboard shortcuts”](#) on page 1405.

Table T-8 **Backup and Restore** tab keyboard shortcuts for the Credentials view

Accelerator	Mnemonic	Command	Result
Alt	CT	Test Credentials	Lets you test the account credentials for the selected backup source.
Alt	CD	Delete Selected Resource	Deletes the selected resource.
Alt	CC	Cancel Active Text	Cancels an active test of credentials.

Job Monitor tab keyboard shortcuts

The following table lists the keyboard shortcuts for the **Job Monitor** tab.

See [“About keyboard shortcuts in Backup Exec”](#) on page 1404.

See [“Storage tab keyboard shortcuts for the Jobs view ”](#) on page 1420.

See [“Storage tab keyboard shortcuts for the Job History view ”](#) on page 1421.

See [“Storage tab keyboard shortcuts for the Backup Sets view ”](#) on page 1422.

See [“Storage tab keyboard shortcuts for the Active Alerts view ”](#) on page 1423.

Table T-9 **Job Monitor** tab keyboard shortcuts

Accelerator	Mnemonic	Command	Result
Alt	J	Job Monitor tab	Opens the Job Monitor tab.
Alt	ST	Standard	Displays Backup Exec in a view that provides descriptive text.
Alt	CO	Compact	Displays Backup Exec in a view that conserves space.
Alt	F	Sort and Filter	Displays information in a custom view or lets you create and save a custom view.
Alt	T	Tree	Displays items in a hierarchal view.
Alt	L	List	Displays items in a list that you can sort by columns.

Table T-9 **Job Monitor** tab keyboard shortcuts (*continued*)

Accelerator	Mnemonic	Command	Result
Alt	CA	Backup Calendar	Lets you view all scheduled backup jobs on a calendar.
Alt	RE	Restore	Browses the backup sets from a single server, and then restores the data.
Alt	JE	Edit	Lets you edit a backup definition. You can edit a backup definition's backup selections, backup settings, and stages.
Alt	JD	Delete	Deletes a backup definition if you no longer need it.
Alt	JC	Cancel	Cancels an active job while it is running.
Alt	JP	Priority	Increases or decreases a job's priority in the job queue .
Alt	JR	Run Now	Runs a job immediately. If a job is scheduled it still runs at its scheduled time.
Alt	JH	Hold	Pauses a job or the entire job queue. The job or job queue resumes its normal schedule when you clear this option.
Alt	JT	Test Run	Runs a test of the selected backup job now.
Alt	JA	Job Activity	Lets you view statistical and system information about a running job. You can also cancel an active job.
Alt	JJ	View Job History	Lets you view detailed information about the job history including individual job and job summary statistics.

Table T-9 **Job Monitor** tab keyboard shortcuts (*continued*)

Accelerator	Mnemonic	Command	Result
Alt	HR	Run Now	Runs a job immediately. If the job is scheduled, it still runs at its scheduled time.
Alt	HL	View Job Log	Lets you view the job log for the selected job history. The job log provides detailed job information, storage and media information, job options, file statistics, and job completion status.
Alt	HD	Delete	Deletes the selected job history and its associated job log, if you no longer need them.
Alt	HH	View Job History	Lets you view detailed information about the job history including individual job and job summary statistics.
Alt	EH	Error Handling	Defines the rules for handling failed jobs.
Alt	HC	Duplicate	Creates a duplicate copy of the job history, which includes all of a job's dependent backup sets. You can duplicate the job history now, or schedule a time to do it.
Alt	HV	Verify	Verifies the integrity of the collection of data and the media on which it resides for this job history. When you verify a job history, you verify all of a job's dependent backup sets. You can verify the job history now, or schedule a time to do it.

Table T-9 Job Monitor tab keyboard shortcuts (*continued*)

Accelerator	Mnemonic	Command	Result
Alt	EB	Edit Backup	Lets you edit the backup for the selected job history. You can edit a backup definition's backup selections, backup settings, and stages.

Storage tab keyboard shortcuts

The following table lists the keyboard shortcuts for the **Storage** tab.

See [“About keyboard shortcuts in Backup Exec”](#) on page 1404.

See [“Storage tab keyboard shortcuts for the Jobs view ”](#) on page 1420.

See [“Storage tab keyboard shortcuts for the Job History view ”](#) on page 1421.

See [“Storage tab keyboard shortcuts for the Backup Sets view ”](#) on page 1422.

See [“Storage tab keyboard shortcuts for the Active Alerts view ”](#) on page 1423.

Table T-10 Storage tab keyboard shortcuts

Accelerator	Mnemonic	Command	Result
Alt	S	Storage tab	Opens the Storage tab.
Alt	ST	Standard	Displays Backup Exec in a view that provides descriptive text.
Alt	CO	Compact	Displays Backup Exec in a view that conserves space.
Alt	F	Sort and Filter	Displays information in a custom view or lets you create and save a custom view.
Alt	T	Tree	Displays items in a hierarchal view.
Alt	L	List	Displays items in a list that you can sort by columns.
Alt	SP	Pause	Pauses a device to prevent scheduled and new jobs from running on it.

Table T-10 **Storage** tab keyboard shortcuts (*continued*)

Accelerator	Mnemonic	Command	Result
Alt	SD	Disable	Disables a device so that it is available for other applications.
Alt	SO	Offline	Troubleshoots the device to bring it online. No operations are allowed on this device until it is online again.
Alt	CC	Configure Storage	Launches the Configure Storage wizard to set up different types of storage to which you can back up data.
Alt	CT	Troubleshoot	Lets Backup Exec troubleshoot the device and provide possible solutions.
Alt	CD	Delete	Removes an item from the Backup Exec Database.
Alt	CS	Share	Shares a device between Backup Exec servers.
Alt	SS	Scan	Let you get information about the media that are in the slots, and then update the Backup Exec Database.
Alt	SI	Inventory	Mounts the media in tape drives, reads the media label, and updates the Backup Exec Database. You can inventory media now, or schedule when to do it.
Alt	SC	Catalog	Logs information about the backup sets and the storage device on which the backup sets are stored.

Table T-10 Storage tab keyboard shortcuts (*continued*)

Accelerator	Mnemonic	Command	Result
Alt	SG	Inventory and Catalog now	Mounts the media in tape drives, reads the media label, and updates the Backup Exec Database. Also, logs information about the backup sets and the storage device on which the backup sets are stored.
Alt	SZ	Initialize now	Sends a startup command to the robotic library.
Alt	SB	Label	Writes a new media label on the media in the drive. Labeling destroys any data that exists on the media.
Alt	SE	Erase	Writes an indicator at the beginning of the media that makes the data on the media inaccessible. Long erase physically erases the entire media.
Alt	VB	Blink	Blinks the status lights on the physical disk to help identify it in a virtual disk.
Alt	VU	Unblink	Turns off the blinking status lights on the physical disk.
Alt	VC	Configure Virtual Disk	Configures a virtual disk on a storage array for use with Backup Exec.
Alt	JR	Run Now	Runs a job immediately. If a job is scheduled it still runs at its scheduled time.

Table T-10 Storage tab keyboard shortcuts (*continued*)

Accelerator	Mnemonic	Command	Result
Alt	MA	Associate with media set	<p>Specifies the append period and the overwrite protection period of a media set to apply to a media.</p> <p>You can select storage and press Enter to view more detailed information about media set operations and media vault operations.</p>
Alt	MS	Scratch	<p>Associates the media with the Scratch media set so that Backup Exec can use it in an overwrite backup job.</p> <p>You can select storage and press Enter to view more detailed information about media set operations and media vault operations.</p>
Alt	MT	Retire	<p>Associates the media with the Retired media set so that Backup Exec cannot use it for backup jobs.</p> <p>You can select storage and press Enter to view more detailed information about media set operations and media vault operations.</p>
Alt	MV	Move media to vault	<p>Lets you type the labels or scan the barcode labels to move the media to a media vault.</p> <p>You can select storage and press Enter to view more detailed information about media set operations and media vault operations.</p>
Alt	RE	Restore	<p>Browses the backup sets from a single server, and then restores the data.</p>

Storage tab keyboard shortcuts for the Jobs view

The following table lists the **Storage** tab keyboard shortcuts for the Jobs view.

See [“Storage tab keyboard shortcuts”](#) on page 1416.

Table T-11 Storage tab keyboard shortcuts for the Jobs view

Accelerator	Mnemonic	Command	Result
Alt	JE	Edit Backups	Edits one or more existing backup jobs. You cannot edit synthetic or one-time backups.
Alt	JE	Edit	Lets you edit a backup definition. You can edit a backup definition's backup selections, backup settings, and stages.
Alt	JD	Delete	Deletes a backup definition if you no longer need it.
Alt	JC	Cancel	Cancels an active job while it is running.
Alt	JP	Priority	Increases or decreases a job's priority in the job queue.
Alt	JR	Run Now	Runs a job immediately. If a job is scheduled it still runs at its scheduled time.
Alt	JH	Hold	Pauses a job or the entire job queue. The job or job queue resumes its normal schedule when you clear this option.
Alt	JT	Test Run	Runs a test of the selected backup job now.
Alt	JA	Job Activity	Lets you view statistical and system information about a running job. You can also cancel an active job.

Table T-11 Storage tab keyboard shortcuts for the Jobs view (*continued*)

Accelerator	Mnemonic	Command	Result
Alt	HH	View Job History	Lets you view detailed information about the job history including individual job and job summary statistics.

Storage tab keyboard shortcuts for the Job History view

The following table lists the **Storage** tab keyboard shortcuts for the Job History view.

See [“Storage tab keyboard shortcuts”](#) on page 1416.

Table T-12 Storage tab keyboard shortcuts for the Job History view

Accelerator	Mnemonic	Command	Result
Alt	HR	Run Now	Runs a job immediately. If the job is scheduled, it still runs at its scheduled time.
Alt	HL	View Job Log	Lets you view the job log for the selected job history. The job log provides detailed job information, storage and media information, job options, file statistics, and job completion status.
Alt	HD	Delete	Deletes the selected job history and its associated job log, if you no longer need them.
Alt	HH	View Job History	Lets you view detailed information about the job history including individual job and job summary statistics.
Alt	EH	Error Handling	Defines the rules for handling failed jobs.

Table T-12 Storage tab keyboard shortcuts for the Job History view *(continued)*

Accelerator	Mnemonic	Command	Result
Alt	HC	Duplicate	Creates a duplicate copy of the job history, which includes all of a job's dependent backup sets. You can duplicate the job history now, or schedule a time to do it.
Alt	HV	Verify	Verifies the integrity of the collection of data and the media on which it resides for this job history. When you verify a job history, you verify all of a job's dependent backup sets. You can verify the job history now, or schedule a time to do it.
Alt	EB	Edit Backup	Lets you edit the backup for the selected job history. You cannot edit synthetic or one-time backups jobs.

Storage tab keyboard shortcuts for the Backup Sets view

The following table lists the **Storage** tab keyboard shortcuts for the Backup Sets view.

See [“Storage tab keyboard shortcuts”](#) on page 1416.

Table T-13 Storage tab keyboard shortcuts for the Backup Sets view

Accelerator	Mnemonic	Command	Result
Alt	SQ	Expire	Expires the selected backup set, if you no longer need it.
Alt	SH	Retain	Retains the selected backup set. You can prevent backup sets from expiring by retaining them.

Table T-13 **Storage** tab keyboard shortcuts for the Backup Sets view
(continued)

Accelerator	Mnemonic	Command	Result
Alt	SC	Catalog	Catalogs the selected backup set. Cataloging backup sets lets you view the data that is contained in them and search for files to restore.
Alt	SU	Duplicate	Creates a duplicate copy of the selected backup set. You can duplicate the backup set now, or schedule a time to do it.
Alt	SV	Verify	Verifies the integrity of the collection of data and the media on which it resides for this backup set. You can verify the backup set now, or schedule a time to do it.
Alt	SL	Expiration	Lets you to edit the expiration date and time.
Alt	SK	Show Dependent Backup Sets	Shows the dependent backup sets for the selected backup set.
Alt	SW	Restore	Browses the backup sets and then restores the data.

Storage tab keyboard shortcuts for the Active Alerts view

The following table lists the **Storage** tab keyboard shortcuts for the Active Alerts view.

See [“Storage tab keyboard shortcuts”](#) on page 1416.

Table T-14 Storage tab keyboard shortcuts for the Active Alerts view

Accelerator	Mnemonic	Command	Result
Alt	AR	Respond	Lets you view the alert and any additional information. You must select Respond OK to clear the alert.
Alt	AO	Respond OK	Clears the alert without displaying information about it. Respond OK only if you no longer need the alert.
Alt	AL	View Job Log	Lets you view the job log that was generated for this job. The job log provides detailed job information, storage and media information, job options, file statistics, and job completion status.
Alt	AH	Show Alert History	Displays alerts that you have responded to or that you have automatically cleared.

Reports tab keyboard shortcuts

The following table lists the keyboard shortcuts for the **Reports** tab.

See [“About keyboard shortcuts in Backup Exec”](#) on page 1404.

Table T-15 Reports tab keyboard shortcuts

Accelerator	Mnemonic	Command	Result
Alt	R	Run Report Now	Runs a selected report immediately.
Alt	C	New Custom Report	Creates a new report that uses the report options that you select.
Alt	N	New Scheduled Report	Schedules a report to run on a specific date and time.

Table T-15 Reports tab keyboard shortcuts (*continued*)

Accelerator	Mnemonic	Command	Result
Alt	E	Edit	Edits a scheduled report or report options for a custom report.
Alt	P	Copy	Lets you make a copy of a custom report. Backup Exec saves a copy of the report in the Custom report group, along with the original report.
Alt	D	Delete	Deletes a custom report, a scheduled report, or a completed report.

General keyboard navigation within the Backup Exec user interface

You can navigate and use Backup Exec with only the keyboard. In the user interface, the current active tree or table has a dark blue highlight, and the current active tab, radio button, or check box is enclosed within a rectangle formed by dotted lines. These areas are said to have focus and will respond to commands.

All Symantec user interfaces use the following keyboard navigation standards:

- The **Tab** key moves the focus to the next active area, field, or control, following a preset sequence. **Shift+Tab** moves the focus in the reverse direction through the sequence.
- **Ctrl+Tab** exits any console area that you internally navigate with the **Tab** key.
- Up and down arrow keys move focus up and down the items of a list.
- The **Alt** key in combination with the underlined mnemonic letter for a field or command button shifts the focus to that field or button.
- **Enter** activates your selection. For example, after pressing the **Tab** key to select **Next** in a wizard panel, press **Enter** to display the next screen.
- **Shift+F10** provides access to context menus.

Keyboard navigation within dialog boxes in Backup Exec

Dialog boxes contain groups of controls that are necessary to set options or settings for programs.

The following list contains some general rules about dialog box navigation:

- The **Tab** key moves the focus between controls within the dialog box along a preset sequence.
- A dark border indicates the default command button. Press **Enter** at any time to choose the button with a dark border.
- **Esc** chooses the Cancel button if one exists.
- The spacebar chooses a control you select with the **Tab** key.
- The spacebar changes the state of a check box that has focus. Typing a mnemonic (if one is available) will move the focus to the check box and change its state.
- Arrow keys move the focus in radio buttons, list boxes, sliders, groups of option controls, or groups of page tabs.
- Items that cannot be changed are not visited by the **Tab** key sequence. Options that are unavailable are grayed-out and can neither be selected nor given focus.

While the controls described here are typically found in dialog boxes, they also can occur in other contexts. The same navigation standards will apply.

List box navigation in Backup Exec

List boxes display a column of available choices.

There are different kinds of list boxes with the following additional navigation conventions:

- Drop-down list boxes by default show only the selected item. A small button to the right of the control shows a downward-pointing arrow. Select the arrow to display more items from the list box. If there are more choices than can fit in the preset list box area, a slider appears along the side of the list box. Show or hide the list using **Alt+Down** arrow, **Alt+Up** arrow, or **F4**. The **Tab** key selects an item.
- Extended selection list boxes support selecting single items, blocks of items, or combinations of the two. After selecting an item, hold down **Ctrl**+navigation keys to select or clear additional items or blocks of items.

Tabbed dialog box navigation in Backup Exec

Some dialog boxes use tabbed pages to subcategorize groups of many options. Each tabbed page contains different groups of controls. Use **Tab** to move the focus between tabbed pages within a dialog box. Typing the mnemonic for a tab also moves the focus to the tabbed page and displays its page of controls.

The following table lists keyboard navigation rules within tabbed dialog boxes:

Table T-16 Keyboard navigation within tabbed dialog boxes

Keyboard input	Result
Ctrl+Page Down or Ctrl+Tab	Switches to the next tab and displays the page.
Ctrl+ Page Up	Switches to the previous tab and displays the page.
Right arrow or Left arrow	When the focus is on a tab selector, chooses the next or previous tab in the current row and displays the page.

About setting accessibility options

Symantec software responds to operating system accessibility settings.

Symantec products are compatible with Microsoft's accessibility utilities. In Windows operating systems, accessibility options involving keyboard responsiveness, display contrast, alert sounds, and mouse operation can be set through the Control Panel.

Accessibility features are primarily for the English version. Localized versions of this product include support for keyboard (mouseless) navigation using accelerator keys and mnemonic keys.

For more information on setting accessibility options, see the Microsoft documentation.

Glossary

ADAMM (Advanced Device and Media Management)	A Backup Exec database that automates the tracking of tape cartridge media. ADAMM expires the backup sets that are stored on tape cartridge media according to the associated media set.
administration console	The user interface that allows you to run Backup Exec operations. The user interface can be run from the Backup Exec server or from a remote computer.
agent	A component that allows computers such as Microsoft SQL Server to interact with the Backup Exec server.
Agent for Windows	A Backup Exec system service that runs on Microsoft Windows computers and allows remote backup and restore of those computers.
alert	An event in Backup Exec that usually requires some form of user interaction or acknowledgment.
alert category	A group of one or more events that occur in Backup Exec and that can generate an alert. Examples of alert categories include Job Success, Install Warning, and Database Maintenance Failure.
alert source	A source that can generate an alert. Alert sources include jobs, media, storage devices, and computers.
alert type	The classification of an alert that lets you determine the severity of the alert. Alert types include Error, Warning, Information, and Attention Required.
allocated media	The tape cartridge media that are associated with a media set and that have current append and overwrite protection periods.
append period	The length of time that data can be added to tape cartridge media. The append period starts when the first backup job is written to the media.
archive	A logical group of archived items that the Backup Exec Archiving Option creates. Archives are contained in vault store partitions. Each archived file system share has its own archive. Each archived Exchange mailbox has its own archive.
audit log	A running history of all actions that are performed in Backup Exec. An entry into the log is created each time an action occurs that is configured to display in the audit log.
backup definition	A container for any backup selections, job templates, and stages that you specify. Backup Exec combines the job templates with the backup selections to create backup jobs. If you specify a stage, then that additional task is also run.

Backup Exec server	The computer on which Backup Exec is installed and where the Backup Exec services are running.
Backup Exec server pool	A feature of the Backup Exec Central Admin Server Option that lets you group managed Backup Exec servers in a pool to which you can restrict backup jobs.
Backup Exec service account	A user account that is configured for the Backup Exec system services. It contains a user name and password and provides the rights to log on as a service and to act as a Backup Exec administrator.
backup method	An option that you select when you run a backup job to specify a full, differential, or incremental backup.
backup set	A collection of data that is backed up from a single source of content. For example, a single source of content can be a server or it can be a Microsoft Exchange dataset. If you select multiple sources of contents, Backup Exec creates multiple backup sets.
backup strategy	The procedures that you implement for backing up your network. Backup strategies include what methods of backup are performed and when backups are performed.
baseline	The first backup job to run in a synthetic backup. The baseline backup runs one time only and backs up all of the files on the selected computer. A full backup is assembled, or synthesized, from a baseline backup and the subsequent incremental backups.
catalog	A database that Backup Exec creates during a backup or archive operation. When you select data to restore, Backup Exec uses the catalog information to find the restore selections and the storage devices on which they reside.
central administration server	A Backup Exec server on which the Central Admin Server Option (CASO) is installed. In a CASO environment, the central administration server provides centralized administration, delegated job processing, and load balancing functionality for managed Backup Exec servers.
centralized catalog	A catalog location in the Central Admin Server Option. All of the files in the catalog are kept on the central administration server.
common encryption key	A type of encryption key that anyone can use to back up data using encryption and to restore encrypted data.
custom error-handling rule	An error-handling rule that you can define for a specific error code in an error category. When a job fails with the error code that is associated with the custom error-handling rule, the retry options and the final job disposition are applied to the job.
data discovery	A Backup Exec feature that allows the detection of new backup content within a Windows domain.
data lifecycle management (DLM)	An automated disk reclamation process that Backup Exec uses to delete expired backup sets that are on disk-based storage. The disk space is then free for use by

	new backup sets. DLM deletes backup sets from disk-based storage after the amount of time expires that you specified when you created the backup job. By default, Backup Exec keeps the most recent backup sets that are necessary to restore any backed-component of a server, even if the backup sets are expired.
differential	A backup method that includes all files that have changed since the last full backup.
disk storage	A location on a locally attached internal hard drive, a USB device, a FireWire device, or a network-attached storage device to which you can back up data.
distributed catalog	A catalog location in the Central Admin Server Option. Image files in the catalog are distributed to the central administration server from every managed Backup Exec server. These distributed files are small because they do not contain the entire catalog. They contain only information about the backup set. The history files, which contain detailed information about the backup set, remain on the managed Backup Exec server.
error-handling rule	A default or custom rule that sets retry options and the final job disposition for failed or canceled jobs. Retry options let you specify how often to retry a job if it fails and the time to wait between retry attempts. The final job disposition lets you place the job on hold until you can fix the error.
event	An action that occurs during a Backup Exec operation, such as a job cancellation.
full	A backup method that includes all of the files that you select for backup.
granular restore	A restore of individual items from a backup for which you enable the Granular Recovery Technology option.
GRT (Granular Recovery Technology)	A backup option that is available with some Backup Exec agents. Granular Recovery Technology lets you restore individual items from database backups. A separate backup of the individual items is not required for you to recover one item.
imported media	The media that are created by a product other than this installation of Backup Exec, but are in storage devices in the Backup Exec environment.
incremental	A backup method that backs up only the files that have changed since the last full or incremental backup.
job delegation	A process by which jobs are distributed by a central administration server to available storage devices on managed Backup Exec servers. Job delegation is only available with the Central Admin Server Option.
job	An operation that has been scheduled for processing by the Backup Exec server. Jobs contain source or destination information, settings, and a schedule. Types of jobs include backup, restore, data discovery, reports, test run, and storage operations.
job history	A list of completed and failed backup, restore, and storage operation jobs.

job log	A log that contains the results of a job. It is created when the job runs. You can review the job log for job errors and job details.
job template	A collection of settings that Backup Exec uses to create jobs. For example, backup job settings can include encryption, scheduling options, or notifications. When a backup job is run, Backup Exec combines the job template with the backup selections to create a backup job.
legacy backup-to-disk folder	A storage device used in versions prior to Backup Exec 2012 that you could create to back up data to a folder on a hard disk. For Backup Exec 2012/2014, these legacy backup-to-disk folders are read-only. Symantec recommends that you use disk storage devices instead.
load balancing	<p>A feature in Backup Exec that automatically distributes jobs among any available storage devices in a storage device pool.</p> <p>Also a feature of the Backup Exec Central Admin Server Option in which jobs are automatically distributed from a central administration server to multiple managed Backup Exec servers for processing among the various storage devices.</p>
logon account	An account that stores the credentials of a Windows user account and that enables Backup Exec to manage user names and passwords. It can be used to browse data sources or to process jobs.
mailbox group	A group of user mailboxes to which you want to assign the same archive rules, retention categories, and vault stores in the Backup Exec Archiving Option. In Enterprise Vault, this is called a provisioning group.
managed Backup Exec server	A Backup Exec server that is managed by a central administration server. Managed Backup Exec servers are responsible for the actual processing of backup and restore jobs in a Central Admin Server Option environment. Managed Backup Exec servers are only available with the Backup Exec Central Admin Server Option.
media ID	A unique internal label that Backup Exec assigns to each media used in Backup Exec. The ID keeps statistics for each media. The media ID cannot be erased or changed.
media label	A label used to identify media. Backup Exec can assign the label automatically, but you can rename it. If the media was first used in a library with a barcode reader, the media label will already have a barcode label.
media overwrite protection level	A global setting in Backup Exec that lets you specify whether to overwrite scratch, imported, or allocated tape cartridge media regardless of the media's overwrite protection period.
media rotation	A strategy that determines when tape cartridge media can be reused, or rotated back into use, by Backup Exec. Common examples of a media rotation strategy are Son, Father/Son, and Grandfather/Father/Son.

media set	A set of rules that apply to tape cartridge media that are associated with a media set. These rules specify append periods, overwrite protection periods, and vaulting periods.
media vault	A logical representation of the actual physical location of tape cartridge media, such as a special media room, a scratch bin, or an offsite location.
mixed backup	A backup definition that contains more than one backup method for multiple data types.
offhost backup	A feature of the Backup Exec Advanced Disk-based Backup Option that enables the backup operation to be processed on a Backup Exec server instead of on the remote computer, or host computer. Moving the backup from the remote computer to a Backup Exec server enables better backup performance and frees the remote computer as well.
Offline Tape Cartridge Media vault	A location on the Storage tab that displays the tape cartridge media that are on-site but are not in tape drives, robotic libraries, or media vaults. Media are automatically moved to the offline vault if you use Backup Exec to remove media from a tape drive or robotic library.
Online Tape Cartridge Media vault	A location on the Storage tab that displays the tape cartridge media that are available in tape drives or robotic libraries. You cannot add or move media to the online media vault. Backup Exec does that automatically.
overwrite protection period	The length of time that data is retained on a specific tape cartridge media before being overwritten (unless the media is erased, formatted, moved to scratch media, or if the media overwrite protection level is set to None). The overwrite protection period is measured from the last time data was appended to the media.
preferred server configuration	A collection of one or more servers and sites that you select as preferred backup sources. Preferred server configurations take priority as backup sources in instances where data is replicated between multiple servers.
recyclable media	Tape cartridge media that are assigned to a media set but have expired data overwrite protection periods.
remote administrator	The Backup Exec user interface (Administration Console) that is run on remote computers.
replicated catalog	A catalog location in the Central Admin Server Option. All of the files in the catalog are replicated from the managed Backup Exec server to the central administration server.
restricted encryption key	A type of encryption key that anyone can use to back up data using encryption. Only the key owner or a user with knowledge of the pass phrase can restore data that was encrypted with a restricted encryption key.

retention category	A setting in the Backup Exec Archiving Option that lets you specify the period of time for which you want to keep items in the archives. You can name a retention category to make it easier to search for and retrieve archived items.
retired media	Tape cartridge media that has been taken out of service, usually because of an excessive number of errors. Media that is retired is available for restore jobs but not for backup jobs. Media must be retired before it can be deleted. If you want to use media that has been deleted, Backup Exec recognizes it as imported media. You must catalog retired media before you can restore from it.
scratch media	Tape cartridge media that are not associated with a media set and that can be overwritten. Scratch media includes new or blank media, erased media, and media moved from another group.
simulated tape library	A tape library that emulates an Advanced Intelligent Tape (AIT) media type and has the AIT media type label. A simulated tape library is created by the Tape Library Simulator.
stage	An additional task that you can run with a backup job, such as duplicating a copy of the backup data to disk storage.
storage device	A disk storage device, disk cartridge, robotic library drive, stand-alone drive, virtual drive, removable storage drive, or other type of data storage that is supported by Backup Exec.
storage device pool	A group of similar types of storage devices that enables load-balancing of Backup Exec jobs.
synthetic backup	A feature of the Advanced Disk-based Backup Option that enables a full backup to be assembled, or synthesized, from a baseline and subsequent incremental backups.
Tape Library Simulator	A utility that lets you create a virtual device on a hard disk or on any mounted volume on a computer on which the Backup Exec Remote Media Agent for Linux is installed. The virtual device that is created is called a simulated tape library.
true image restore	A feature of the Advanced Disk-based Backup Option that enables Backup Exec to restore the contents of directories to what they were at the time of any full or incremental backup. Restore selections are made from a view of the directories as they existed at the time of the particular backup. Files that were deleted before the time of the backup are not restored. In true image restore, only the correct versions of files are restored from the appropriate full or incremental backups that contain them. Previous versions are not unnecessarily restored and then overwritten.
UMI (Unique Message Identifier)	A unique code that is associated with an error reported in the job log, or on some alerts. These codes contain hyperlinks that you can click to go to the Symantec Technical Support Web site. You can access technical notes and troubleshooting tips that are related to a specific error.

vault store	A disk-based container for the archived data that the Backup Exec Archiving Option archives from one server.
vault store partition	The physical location on a disk where archived items are stored, if the Backup Exec Archiving Option is installed. Backup Exec creates one vault store partition in each vault store by default. As the data in a vault store grows, you can create more vault store partitions to provide additional capacity.
virtual disk	A logical disk that you configure on a storage array to provide storage to the Backup Exec server.

Index

A

- accelerator
 - defined 1404
- accessibility
 - dialog boxes 1425
 - keyboard navigation 1425
 - keyboard shortcuts 1404
 - Backup and Restore tab 1405
 - Backup and Restore tab for the Active Alerts view 1412
 - Backup and Restore tab for the Backup Sets view 1410
 - Backup and Restore tab for the Credentials view 1412
 - Backup and Restore tab for the Job History view 1409
 - Backup and Restore tab for the Jobs view 1408
 - Home tab 1405
 - Job monitor 1413
 - Reports tab 1424
 - Storage tab 1416
 - Storage tab for the Active Alerts view 1423
 - Storage tab for the Backup Sets view 1422
 - Storage tab for the Job History view 1421
 - Storage tab for the Jobs view 1420
 - overview 1403
 - settings 1427
- active alerts
 - responding to 293
- Active Directory
 - backing up in Exchange 946
- Active File Exclusion 950
- active jobs
 - about managing and monitoring 253
 - canceling 257
 - holding 258
 - removing a hold 259
 - statuses 277
 - viewing job activity 257
 - viewing properties 257
- Add
 - Remote Media Agent for Linux 1252
- adding tapes by importing media 459
- administration console
 - overview 133
 - role in backup process 38
- Advanced Disk-based Backup Option
 - about 1155
 - baseline
 - setting 1156
 - best practices for offhost backup 1165
 - host computer
 - defined 1163
 - offhost backup options 1159, 1161
 - offhost backup overview 1163
 - setting default options 1159
 - transportable snapshots
 - defined 1163
 - true image restore
 - overview 1161
- Advanced Open File
 - configuring options for backup jobs 540
- Agent for Hyper-V
 - adding a Hyper-V host 869
 - and GRT 881
 - backing up 873
 - delayed catalogs 883
 - highly available virtual machines 883
 - how byte count is calculated for delayed catalog 884
 - installation overview 864
 - installing Agent for Windows 871
 - overview 860
 - protecting Microsoft application data 881
 - push-installing Agent for Windows 872
 - requirements 862
 - restoring 885
 - setting default backup options 878
 - viewing details about resources 870
 - with GRT and vhd files 865
 - with Microsoft incremental backups 866
 - with Windows Server 2012 867

- Agent for Hyper-V *(continued)*
 - with Windows Server 2012 R2 868
 - with Windows Server 2012/2012 R2 865
- Agent for Linux
 - about backing up 1204
 - about establishing trust 780
 - about exclusions from backup 1196
 - backing up Novell OES components 1206
 - backup job options 1205
 - beoper group, defined 1192
 - configuration options in the ralus.cfg file 1197
 - configuring the ralus.cfg file 1195
 - creating the beoper group 1192
 - default options 1209
 - editing configuration options in the ralus.cfg file 1196
 - editing default options 1209
 - establishing trust relationship 1193
 - installing 1190
 - Novell OES, requirements for backup 1207
 - publishing to Backup Exec servers 1194
 - push-installing 1189
 - requirements 1188
 - restore options 1208
 - restoring 1207
 - runtime scripts 1213
 - saving the installatoin log 1190
 - starting the Linux agent daemon 1214
 - stopping the Linux agent daemon 1215
 - troubleshooting 1215
 - uninstalling 1211
 - uninstalling manually 1212
 - using SSH 1189
- Agent for Mac
 - about the ralus.cfg file 1226
 - default options 1235
 - editing default options 1235
 - installing 1221
 - manually starting 1232
 - manually stopping 1233
 - requirements 1219
 - restore options 1238
 - restoring 1238
 - supported backup methods 1235
 - troubleshooting 1238
 - uninstalling 1223
- Agent for Microsoft Active Directory
 - about 1086
 - about restoring individual objects 1092
- Agent for Microsoft Active Directory *(continued)*
 - Granular Recovery Technology (GRT)
 - overview 1088
 - passwords 1093
 - recreating purged objects 1092
 - requirements 1087
 - tombstones 1092
- Agent for Microsoft SharePoint
 - about 971
 - adding a farm 974
 - backing up SharePoint data 974
 - disabling or enabling communication between web servers and Backup Exec 983
 - disaster recovery for SharePoint 2010/2013 985
 - disaster recovery SharePoint 2007 989
 - installing 971
 - overview 971
 - requirements 971
 - restoring SharePoint data 981
 - system requirements 971
 - using with SharePoint Server 2007 and Windows SharePoint Services 3.0 973
 - using with SharePoint Server 2010/2013 and SharePoint Foundation 2010/2013 972
- Agent for VMware
 - adding VMware vCenter and ESX/ESXi hosts 834
 - backing up Microsoft application data 847
 - backup defaults 843
 - backup methods 838
 - dynamic inclusion 847
 - Granular Recovery Technology
 - about 847
 - GRT requirements 847
 - installing 833
 - installing Agent for Windows 836
 - overview 830
 - requirements 831
 - restoring resources 852
 - selecting transport method for VMDK file 842, 846
 - unsupported characters 831
 - unsupported GRT configurations 847
 - with Windows Server 2012 832
- Agent for Windows
 - about 777
 - about establishing trust 780
 - Backup Exec Agent Utility 780
 - hardware requirements 778

- Agent for Windows *(continued)*
 - installation methods 86
 - installing in an Active Directory network 94
 - installing on a Microsoft cluster 714
 - installing on Hyper-V virtual machines 871
 - installing on virtual machines 836
 - installing updates 93
 - installing using a command script 101
 - installing using the command prompt 98
 - licenses 778
 - publish to Backup Exec servers 784
 - push-installing to remote computers 86
 - push-installing to VMware virtual machines 837
 - push-installing to Hyper-V virtual machines 872
 - requirements 778
 - stopping and starting 779
 - uninstalling using a command script 102
 - uninstalling using command prompt 100
- agents
 - Backup Exec
 - trial version 118
 - upgrading 125
- Alert History by Backup Exec server report 653
- Alert History report 652
- alerts
 - categories 286
 - clearing informational alerts 294
 - configuring categories 302
 - configuring defaults 306
 - configuring groups for notification 299
 - copying text 292
 - defined 286
 - deleting from alert history 291
 - deleting recipients 301
 - disabling pop-ups 307
 - enabling email and text messages 297
 - enabling pop-ups 307
 - filters 292
 - notification 295
 - responding to 293
 - sending job complete notification 304
 - setting up notification 294
 - severity 286
 - showing on the Home tab 289
 - SNMP traps 308
 - stopping notification for recipient 305
 - viewing job log 293
 - where to find 287
- All Servers server group
 - about 154
- allocated media
 - overwriting 395
- Alternate location
 - setting an SDR 747
- Any virtual disk storage device pool
 - description 1278
- append period
 - changing 392
 - defined 382, 390
 - editing 388
 - setting for media set 390
- archive bit
 - using to determine backed up status 200
- Archive Job Success Rate 653
- Archive logging
 - Lotus Domino 1073
 - Lotus Domino recovery 1081
- Archive Selections by Archive Rules and Retention
 - Categories report 654
- archive settings
 - overview 1363
- archives
 - deleting 1357
 - deleting items with expired retention periods 1357
 - editing properties 1356
 - overview 1356
- Archiving Option
 - archives, overview 1356
 - arranging mailbox groups for provisioning 1360
 - assigning a vault store 1345
 - audit log entries 1358
 - backing up components 1377
 - backing up from remote Backup Exec
 - server 1377
 - best practices 1339
 - configuring archiving rules for file system
 - selections 1369
 - data not archived 1338
 - deleting a vault store 1353
 - deleting archives 1357
 - deleting data after archiving 1363
 - deleting items from archives 1375
 - deleting items with expired retention
 - periods 1346, 1357, 1377
 - Disabled job status 1336
 - disabling backup mode 1381
 - editing archive properties 1356

Archiving Option (*continued*)

- editing default retention category 1346
- editing retention categories 1367
- editing vault store partition properties 1355
- editing vault store properties 1352
- enabling single instance storage 1346
- granting permissions on Exchange Server 1320
- installing 1332
- installing Enterprise Vault 1336
- item deletion mode 1352
- items not supported 1314
- managing mailbox groups 1346, 1359
- overview 1313, 1337
- overview of archive settings 1363
- overview of components 1377
- overview of mailbox groups 1358
- overview of retention categories 1366
- redirecting restores for Directory database 1384
- redirecting restores of all components 1383
- redirecting restores with Backup Exec Utility 1383–1384
- reinstalling 1335
- reports 1388
- requirements 1314
- restoring from remote Backup Exec server 1377
- restoring items from archives 1374
- running Backup Exec Utility 1386
- running consistency checks on databases 1381
- running Enterprise Vault services 1336
- searching archives for data 1375
- selecting administrative shares 1342
- selecting file system shares and folders 1342
- setting defaults 1346
- setting rules for archiving mailbox groups 1361
- specifying archive settings 1364
- specifying retention period 1368
- synchronizing permissions and settings 1346, 1350
- troubleshooting 1388
- uninstalling 1335
- updating SQL Server name 1383
- vault store partitions, overview 1354
- vault stores, overview 1351
- viewing the Enterprise Vault event log 1388
- viewing vault store partition status 1355
- viewing vault store status 1352

ARCserve media

- about restoring data from 246

audit log

- about 624
- Archive Option entries 1358
- configuring 624
- removing entries 625
- saving to a file 626
- viewing 625

Audit Log report 655

- automatic exclusion of SQL data during volume level backups 899

B

- back up and delete the files method
 - freeing disk space 202

- backp definition
 - creating 158

- Backup
 - using the Remote Media Agent for Linux 1257

- backup
 - creating 158
 - overview 158

- Backup and Restore tab
 - list of servers 152

- backup definition
 - creating 161
 - creating from an existing backup definition 161
 - defined 158
 - editing 206
 - excluding selections 178
 - including selections 181
 - one-time 158
 - selecting data 169

- Backup Exec
 - overview
 - how it works 38

- Backup Exec Agent Utility
 - activity status
 - viewing 782
 - command line applet 789
 - switches 790
 - using 789
 - database access
 - configuring 787

- Event Viewer
 - open 780

- job template name for DBA-initiated jobs 1002

- Linux
 - configure Oracle instance on 1004

- Backup Exec Agent Utility *(continued)*
 - port
 - configure for Oracle operations 1005
 - publish to Backup Exec servers 784
 - publishing
 - adding Backup Exec servers 784
 - editing Backup Exec server information 785
 - removing Backup Exec servers 786
 - Real Application Cluster (RAC)
 - publish to Backup Exec server 1001
 - refresh interval
 - setting 783
 - Registry Editor
 - open 780
 - security
 - removing certificate 789
 - Services
 - open 780
 - starting 781
 - starting automatically 783
 - update credentials for Linux instances 1002
 - view status 782
 - Windows
 - configure Oracle instance on 999
- Backup Exec Archiving Site
 - backing up 1377
- Backup Exec diagnostic application
 - generating 702
- Backup Exec Migrator
 - about 1041
 - about retrieving Enterprise Vault data 1058
 - about staged migrations 1046
 - about the Backup Exec Backup Sets view 1057
 - Backup Exec server
 - working with 1052
 - best practices 1061
 - communicating with Enterprise Vault 1054
 - configuring 1049
 - data migration process 1046
 - Enterprise Vault retention periods 1048
 - events
 - about 1047
 - how it works 1042
 - log file location 1048
 - logs
 - about 1047
 - migrated files
 - about deleting 1048
 - Migrator for Enterprise Vault options 1053
- Backup Exec Migrator *(continued)*
 - requirements 1042
 - retrieving Enterprise Vault data 1058
 - troubleshooting 1062
- Backup Exec server 38
 - viewing properties 627
- Backup Exec services
 - changing service account credentials 621
 - changing startup options 623
 - stopping and starting 621
- Backup Exec settings
 - changing preferences 565
 - database maintenance 567
 - DBA-initiated jobs 597
 - discover data to back up 571
 - Granular Recovery Technology (GRT)
 - options 596
 - network and security 576
- Backup Exec Utility
 - redirecting restores for Archiving Option Directory
 - database 1384
 - redirecting restores of Archiving Option 1383
 - running for Archiving Option 1386
- backup job
 - creating 158, 161
 - creating from an existing backup definition 161
 - deduplication 822
 - editing 206
 - excluding selections 178
 - excluding selections globally 178
 - including selections 181
 - one-time 158
 - pre/post commands 547
 - preparing for 150
 - required user rights 151
 - running the next scheduled instance 205
 - selecting data 169
- Backup Job Success Rate report 655
- backup methods
 - about 188
 - advantages and disadvantages 191
 - configuring 196
 - delete selected files and folders after successful
 - backup 202
 - differential 190
 - duplicate 188
 - full 188
 - incremental 190
 - specific types of data 193

- backup methods *(continued)*
 - VMware data 838
- backup network
 - overview 574
- Backup Recommendations report 655
- Backup Resource Success Rate report 656
- backup selections
 - about 169
 - changing the order 175
 - critical system components 172
 - multiple servers or applications 172
 - using fully qualified domain names 174
- backup sets
 - about 342
 - about duplicating 221
 - about verifying 227
 - automatically deleting expired 335
 - cataloging 251
 - changing the expiration dates 345
 - duplicating 222
 - expiring 345
 - prevent from expiring 346
 - releasing from retention 347
 - retaining 346
 - verifying 227
 - viewing contents 348
 - viewing properties 348
- Backup Sets by Media Set report 656
- Backup Size By Resource report 657
- backup strategies
 - increase throughput with Agent for Windows 777
- backup-to-disk folder
 - as read only legacy storage 350
 - changing the location of 352
 - importing 354
 - recovering with Simplified Disaster Recovery 350
 - recreating 353
- barcode labels
 - default 400
 - overview 400
 - robotic library support 400
- barcode media, finding in drives and portals 452
- baseline
 - setting for synthetic backup 1156
- beoper group
 - Agent for Linux, about 1192
 - creating 1192
- blink feature
 - about 1310

- blink feature *(continued)*
 - how to identify the physical disks 1310
- Boot managers
 - restoring in SDR 764
- buffer count
 - setting for tape drives 364
- buffer size
 - setting for disk cartridge devices 331
 - setting for tape drives 363
- byte count
 - how it is calculated for delayed catalog operations 884
 - incorrect 693

C

- calendar
 - excluding dates 562
 - viewing scheduled backup jobs 219
- CASO
 - alerts 1132
 - Backup Exec server
 - changing to a managed Backup Exec server 1113
 - Backup Exec server pool
 - adding managed Backup Exec servers 1138
 - creating 1138
 - deleting 1139
 - overview 1136
 - removing a managed Backup Exec server 1139
 - selecting for backup 1137
 - Backup Exec Utility
 - running 1153
 - catalog locations 1118
 - central administration server
 - setting for a managed Backup Exec server 1113
 - centralized restore
 - multiple storage devices 1140
 - overview 1140
 - changing to a central administration server 1112
 - communication thresholds 1131
 - deleting managed Backup Exec server 1114
 - disabling communications 1132
 - disaster recovery 1150
 - enabling communications 1132
 - installing 1101
 - installing across a firewall 1107
 - installing managed Backup Exec server 1102

- CASO *(continued)*
 - job delegation 1134
 - managed Backup Exec server
 - configuration options 1106
 - viewing properties 1145
 - network interface cards
 - using any available 1133
 - network traffic
 - reducing 1118
 - notifications 1132
 - overview 1096
 - pausing managed Backup Exec server 1144
 - recovering failed jobs 1142
 - requirements 1098
 - restoring data from multiple devices 1141
 - resuming paused managed Backup Exec server 1144
 - settings for managed Backup Exec servers 1120
 - starting Backup Exec services 1144
 - status 1131
 - stopping Backup Exec services 1144
 - storage and media data 1099
 - troubleshooting 1151
 - uninstalling Backup Exec from central administration server 1153
 - uninstalling Backup Exec from managed Backup Exec server 1154
 - upgrading 1110
- catalog
 - configuring delayed catalog operations 536
 - defined 248
 - delayed catalog 248
 - editing default options 249
 - levels 251
 - media with encrypted backup sets 409
- catalog operation errors
 - DLT tape drive hangs 690
- cataloging
 - media 452
- catalogs
 - using to determine backed up status 201
- centralized catalogs in CASO 1118
- centralized restore
 - best practices 1141
- CHECKCATALOG utility 895
- CHECKDB utility 895
- checkpoint restart
 - about 543
 - configuring 543
- checkpoint restart *(continued)*
 - configuring default settings 544
 - considerations 544
 - supported technologies 543
- Circular logging
 - Lotus Domino 1073
 - recovery of Lotus Domino server 1084
- cleaning a drive 458
- cleaning slots
 - defining for robotic libraries 373
- client-side deduplication
 - overview 821
- cluster failover error-handling rule 276
- clusters
 - disaster recovery 728
 - entire cluster manually 731
 - nodes using SDR 729
 - using SDR to prepare 729
 - Microsoft 718
 - adding or removing a failover node 717
 - all drives pool 716
 - BEUtility 718
 - changing the order in which nodes fail over 717
 - configurations 719–722, 724
 - disaster recovery 732–733
 - disaster recovery of Backup Exec on a cluster using SDR 730
 - failover restart 712
 - installation 713, 715
 - overview 727
 - uninstalling Backup Exec 716
 - troubleshooting 734
 - using with Backup Exec 710
- command line
 - installing Backup Exec 108
 - installing Remote Administrator 103
 - switches for installation 109
- command prompt
 - uninstalling Agent for Windows 100
- common encryption keys 585
- compact view 138
- completed jobs
 - statuses 277
- configuration settings
 - copying to another server 626
- Configure Storage Wizard
 - changing or adding hot spares 1308
 - configuring a storage array 1270

- Configure Storage Wizard *(continued)*
 - overview 428
- configuring
 - holidays 562, 565
- consistency check options
 - Exchange Agent 953, 958
 - SQL Agent 894
- continuing Exchange backup if consistency check fails 953, 959
- control connection with the Remote Media Agent for Linux 1242
- conversion to virtual machines
 - adding a stage for 498
 - after a backup job 492
 - backup methods 479
 - conversion of disks larger than 2TB 479
 - from point in time 503
 - how backup selections are processed 479
 - one-time conversion 505
 - options 479
 - overview 479
 - requirements 487
 - setting default options 509
 - simultaneously with backup job 488
 - Windows Server 2012 Hyper-V host 479
- Copy text 136
- credentials
 - creating for backup sources 213
 - creating for jobs 213
 - replacing for a backup source 212
 - replacing for jobs 212
 - testing for backup sources 211
 - testing for jobs 211
- Credentials pane
 - about 211, 214
 - deleting backup sources 214
- critical system components
 - about 183
 - restore scenarios 186
 - selecting 172

D

- Daily Device Utilization report 658
- damaged tape media
 - removing 406
- DAOS
 - .nlo files 1066
 - about the Lotus Domino Agent and DAOS 1066
 - DAOS-enabled databases 1066

- data connection to remote computers 1242
- data discovery
 - used with Exchange Agent 932
- data lifecycle management, overview 335
- database maintenance
 - about 567
 - configuring 567
- database server
 - in Microsoft clusters 718
- Database snapshots
 - SQL 909, 919
- DBA-initiated jobs
 - creating a template 598
 - deleting a template 599
 - duplicate job settings 606
 - editing a template 599
 - general options 605
 - network options 605
 - storage options 600
 - templates 597
- Debug Monitor 707
- deduplication
 - Windows 168
- Deduplication Device Summary report 659
- deduplication disk storage
 - changing logon account password 816
 - overview 809
- deduplication disk storage devices
 - disaster recovery 827
 - editing properties 811
 - requirements 798
 - restoring 827
- Deduplication Option
 - about backing up 822
 - about copying deduplicated data to tapes 826
 - changing logon account password for
 - deduplication disk storage 816
 - client-side deduplication overview 821
 - copying data between OpenStorage devices or
 - deduplication disk storage devices 823
 - deduplication disk storage overview 809
 - deduplication disk storage properties 811
 - deduplication methods for agents 797
 - direct access
 - editing properties 818
 - selecting storage devices 817
 - disaster recovery of deduplication disk
 - storage 827
 - disaster recovery of OpenStorage devices 828

- Deduplication Option *(continued)*
 - installing 803
 - OpenStorage device overview 804
 - OpenStorage device properties 806
 - overview 795
 - requirements 798
 - restoring a deduplication disk storage device 827
 - sharing devices 821
 - with encryption 826
- Deduplication Summary report 660
- default backup options
 - Hyper-V 878
- Default options
 - Simplified Disaster Recovery settings 748
- default options
 - Agent for VMware 843
 - conversion to virtual machines 509
 - NDMP Option 1183
- default preferred configuration settings for tape drives 364
- delayed catalog
 - and Hyper-V virtual machine backups 883
 - how byte count is calculated for Hyper-V 884
- delayed catalog operation
 - configuring 536
 - for GRT-enabled jobs 590
- deleting
 - tape media 407
- destination Backup Exec server
 - adding 626
 - importing a list 626
- Details pane
 - Hyper-V 870
- Device Summary report 660
- devices
 - adding iSCSI-attached 358
 - OpenStorage overview 804
 - reconnecting USB tape devices 358
- diagnostic file
 - command line switches 704
- differential backups
 - about 190
 - advantages and disadvantages 192
- direct access
 - editing properties 818
 - selecting storage devices 817
- DirectCopy to tape
 - copying data 230
- DirectCopy to tape *(continued)*
 - overview 229
- directories
 - about including and excluding for NDMP servers 1176
- Directory database
 - backing up for Archiving Option 1377
- disabled job status 1336
- disabling backup mode in Archiving Option 1381
- disaster preparation
 - Disaster Preparation Plan (DPP) 144
 - Exchange Server 968
 - hardware protection 145
 - off-site storage 146
 - overview 144
- Disaster recovery
 - alternate path in SDR 749
 - Lotus Domino Agent 1079
 - setting path locations
 - disaster recovery information file 748
- disaster recovery
 - clusters
 - Backup Exec on a Microsoft cluster using SDR 730
 - entire cluster manually 731
 - nodes using SDR 729
 - overview 728
 - using SDR to prepare 729
 - deduplication disk storage 827
 - Exchange Server 968
 - Microsoft clusters
 - Backup Exec 733
 - data files 732
 - shared disks 732
 - Microsoft SharePoint 2007 989
 - Microsoft SharePoint 2010/2013 985
 - OpenStorage devices 828
- discover data to back up
 - about 571
 - adding servers 571
 - configuring 571
- disk cartridge storage
 - about 328
- disk storage
 - creating 318
 - editing properties 321
- Disk Storage Summary report 661
- disk-based storage
 - about 313

- distributed catalogs in CASO 1118
- DLT tape
 - drive hangs when cataloging 690
- domain controller
 - installing by redirected restore 242
 - restoring System State 241
- drive pools
 - creating in a Microsoft cluster 716
- duplication between OpenStorage devices or
 - deduplication disk storage devices 823
- dynamic inclusion
 - for Hyper-V 880

E

- editions of Backup Exec
 - listed and described 37
- eject media 458
 - after job completes 531
- email notification
 - configuring 295
- encrypted files
 - about cataloging media 409
- encrypted SQL database restore 921
- encryption
 - about 583
 - hardware 584
 - restoring encrypted SQL databases 921
 - software 584
 - types 583
 - with deduplication 826
- encryption keys
 - 128-bit AES 583
 - 256-bit AES 583
 - about deleting 590
 - common 585
 - creating 587
 - deleting 590
 - encryption types 583
 - managing 586
 - overview 585
 - pass phrases 585
 - replacing 588
 - restoring encrypted data 245
 - restricted 585
- Enterprise Vault
 - running services for 1336
 - viewing event log 1388
- Enterprise Vault Agent
 - about redirecting a restore job 1037
- Enterprise Vault Agent *(continued)*
 - about restoring 1031
 - About restoring individual files and folders 1036
 - automatic redirection of Enterprise Vault
 - components 1036
 - available backup methods 1025
 - Backup Exec Migrator
 - about 1041
 - about deleting migrated files 1048
 - about events 1047
 - about logs 1047
 - about retrieving Enterprise Vault data 1058
 - about staged migrations 1046
 - about the Backup Exec Backup Sets
 - view 1057
 - best practices 1061
 - communicating with Enterprise Vault 1054
 - configuring 1049
 - data migration process 1046
 - Enterprise Vault retention periods 1048
 - how it works 1042
 - log file location 1048
 - Migrator for Enterprise Vault options 1053
 - requirements 1042
 - retrieving Enterprise Vault data 1058
 - troubleshooting 1062
 - VxBSA logs 1047
 - working with a Backup Exec server 1052
 - Backup Exec server
 - log file location 1048
 - logs 1047
 - best practices 1041
 - collections
 - configuring 1050
 - vault store partition properties 1051
 - installing 1024
 - migration
 - vault store partition properties 1056
 - non-operational state 1031
 - Partition Recovery Utility
 - about 1059
 - finding an archive ID 1060
 - log file location 1048
 - logs 1047
 - requirements 1059
 - running 1060
 - troubleshooting 1062
 - ready-to-use state 1031
 - redirection options 1038

- Enterprise Vault Agent *(continued)*
 - requirements 1023
 - restore options 1033
 - selecting a backup method 1024
- Environment Check
 - results 58
 - running before installing or upgrading 58
- error codes
 - Unique Message Identifier
 - viewing 270
- error-handling rules
 - cluster failover rule 276
 - creating 273
 - custom rules
 - defined 272–273
 - custom rules for recovered jobs 275
 - default rules
 - defined 272
 - deleting a custom rule 274
 - enabling for a failed job 275
 - enabling or disabling 274
 - overview 272
 - recovered jobs custom rule 273
- Error-Handling Rules report 662
- ESX/ESXi host, adding 834
- Event Recipients report 663
- Exchange Agent
 - Active Directory
 - backing up 946
 - automatic exclusion of files during volume level backups 950
 - backing up
 - Exchange 950
 - recommended selections 946
 - backup methods 955–956, 960–961
 - best practices 943
 - data discovery feature 932
 - databases
 - configuring 962
 - disaster recovery 968
 - Exchange 2007 snapshot backup method 949
 - Exchange high availability server option 954, 959
 - Exchange Web Services
 - overview 947
 - excluding files during volume level backups 950
 - Granular Recovery Technology (GRT)
 - overview 947
 - requirements for 932
 - setting for backup 957, 962
- Exchange Agent *(continued)*
 - installation 939
 - Internet Information Service (IIS) metabase
 - backing up 946
 - mailbox access requirements 944
 - offhost backup
 - with Granular Recovery Technology (GRT) 947
 - overview 930
 - protecting Exchange using VSS 948
 - redirecting data 962
 - requirements 931
 - restore of individual items
 - requirements 932
 - restore requirements 962
 - restoring data from snapshot backups 962
 - restoring data to server 962
 - restoring Exchange 2007 with Recovery Storage Group 962
 - restoring individual public folder messages from tape 962
 - services account 931
 - snapshot backup
 - configuring 948
 - snapshot technology 948
 - storage groups
 - backing up 950
 - strategies for backing up 945
 - system state
 - backing up 946
 - troubleshooting snapshot and offhost jobs 950
 - volume level backups
 - automatic exclusion of files 950
- Exchange Mailbox Archiving Option
 - overview 1313
- Exchange Mailbox Group Archive Settings report 664
- Exchange Web Services
 - using with the Exchange Agent 947
- exclude dates
 - deleting dates 564
 - exporting dates to another server 565
 - importing a list of dates 562
 - selecting dates for all backups 562
 - selecting dates for individual backup jobs 218
- executing a command
 - after backup 550
 - before backup 550
- exporting expired media 466
- exporting media 466

F

- failback
 - defined 712
- Failed Archive Jobs report 664
- Failed Backup Jobs report 665
- failover
 - adding or removing a failover node 717
 - changing the order in which nodes fail over 717
 - defined 710
 - restart 712
- FAT
 - partition 693
- father/son media rotation strategy 412
- file history
 - enabling for NDMP Option 1183
- File System Archive Settings report 666
- File System Archiving Option
 - overview 1313
- files
 - about including and excluding for NDMP servers 1176
- files and folders options
 - configuring for backup jobs 552
- filters
 - for alerts 292
- fingerprint database
 - backing up for Archiving Option 1377
 - for vault stores 1351
- firewall
 - Backup Exec ports 580
 - browsing systems through 579
 - enabling a SQL instance behind 583
 - using Backup Exec with 578
- formatting media 457
- full backups
 - about 188
 - advantages and disadvantages 191

G

- grandfather media rotation strategy 412
- Granular Recovery Technology
 - about using with VMware 847
 - requirements for VMware 847
 - unsupported configurations for VMware 847
 - with the Agent for Hyper-V 881
- Granular Recovery Technology (GRT)
 - about restoring individual items 590
 - Exchange data 947
 - offhost backup 947

Granular Recovery Technology (GRT) *(continued)*

- recommended devices for 593
- requirements 594
- setting default options 596
- staging 593
- using Exchange Web Services 947
- Group Policy Object, configuring 97
- groups
 - configuring to receive notifications 299
- GRT
 - about using with VMware 847
 - requirements for VMware 847
 - unsupported configurations for VMware 847

H

- hardware
 - protection in case of disaster 145
 - troubleshooting 687
- high water count
 - setting for tape drives 364
- highly available virtual machines
 - about backing up and restoring 883
- Home tab
 - about 139
 - configuring 139
 - Layout items 139
 - restoring the default configuration 139
 - Support items 139
 - System Health items 139
- hot key
 - defined 1404
- hot spare
 - best practices 1307
 - changing or adding 1308
 - description 1307
 - specifying 1270
- Hyper-V host
 - adding to list of servers 869

I

- IBM computers
 - recovering with Simplified Disaster Recovery 763
- imported media
 - overwriting 395
- imported tape media
 - labeled by Backup Exec 400
- importing media 459

- include dates
 - creating 216
- incremental backups
 - about 190
 - advantages and disadvantages 192
- index locations
 - backing up for Archiving Option 1377
- initializing a robotic library 456
- installation
 - additional options 76
 - Agent for Windows 86
 - Agent for Windows from command prompt 98
 - Agent for Windows with command script 101
 - checklist 57
 - command line switches 109
 - configuring Group Policy Object 97
 - creating a transform 94
 - Environment Check
 - overview 58
 - from command line 108
 - Media Agent for Linux 1244
 - Microsoft SQL Server 2008 R2 Express 60
 - Migration Report 129
 - NDMP Option 1172
 - overview 54
 - parameter files
 - creating 117
 - post-installation tasks 129
 - pre-upgrade checklist 127
 - push-installing Agent for Windows 86
 - push-installing to remote computers 78
 - Remote Administrator 102
 - Remote Administrator from command line 103
 - standard features 61
 - system requirements 63
 - trial version 118
 - typical 66
 - uninstalling Agent for Windows using command prompt 100
 - uninstalling Agent for Windows with command script 102
 - uninstalling Backup Exec 130
 - uninstalling options from local Backup Exec server 131
 - updates to Agent for Windows 93
 - Windows Management Instrumentation
 - performance counter 311
 - Windows Management Instrumentation SNMP provider 312

- installation log 119
 - Agent for Linux 1190
- installation overview
 - Agent for Hyper-V 864
- installation parameter file
 - creating 117
- Installation Summary Report 119
 - viewing 120
- installed updates
 - viewing 123
- installing
 - additional Backup Exec options on a Microsoft cluster 715
 - Backup Exec in a Microsoft cluster 713
 - SharePoint Agent 971
 - to an existing Microsoft SQL Server 2005 or 2008 instance 61
- Internet Information Services (IIS) metabase
 - backing up 946
- inventory
 - robotic libraries when Backup Exec services start 368
- IPv4 576
- IPv6 576
- iSCSI-attached devices
 - adding 358

J

- job activity 257
- job defaults
 - backing up multiple servers or applications 566
 - backup jobs 515
 - changing 515
 - exclude dates 562
 - exclude selections 178
 - schedule for rule-based and run now jobs 559
- job history 255
 - about duplicating 221
 - about verifying 227
 - deleting jobs 265
 - duplicating 224
 - overview 264
 - running a job 266
 - verifying 228
 - viewing 265
- job log 267
 - configuring default options 271
 - finding text 267
 - linking to technical support web site 270

- job log *(continued)*
 - printing 269
 - saving 269
 - viewing from an alert 293
 - with vertical applications 270
- job monitor 253, 255
- job progress indicators
 - displaying 566
- job queue
 - holding 260
 - removing the hold 260
- job status 277
- job status and recovery 282
- job template
 - defined 158
- jobs
 - canceling 257
 - changing priority for scheduled 262
 - configuring error-handling rules 273
 - deleting from Job History 265
 - deleting scheduled 263
 - holding 258
 - holding the job queue 260
 - holiday scheduling 562, 565
 - managing and monitoring 253
 - removing a hold 259
 - removing hold on the job queue 260
 - running from Job History 266
 - running scheduled job 261
 - sending notification when complete 304
 - setting status and recovery options 282
 - viewing the job log 267
- Jobs Summary report 666

K

- keyboard navigation
 - dialog boxes 1425
 - standards 1425
- keyboard shortcuts
 - Backup and Restore tab 1405
 - Backup and Restore tab for the Active Alerts view 1412
 - Backup and Restore tab for the Backup Sets view 1410
 - Backup and Restore tab for the Credentials view 1412
 - Backup and Restore tab for the Job History view 1409
 - Backup and Restore tab for the Jobs view 1408

- keyboard shortcuts *(continued)*
 - Home tab 1405
 - Job Monitor tab 1413
 - Reports tab 1424
 - Storage tab 1416
 - Storage tab for the Active Alerts view 1423
 - Storage tab for the Backup Sets view 1422
 - Storage tab for the Job History view 1421
 - Storage tab for the Jobs view 1420

L

- labeling media 398
- labeling tape media
 - imported media label 400
 - renaming 401
 - using barcode labels 400
- Latest Backup Exec Device Drivers 360
- Least Free Space First storage device in pool 419
- Library Expansion Option
 - about 358
 - SCSI addresses for hardware 367
 - setting up hardware 367
- license information
 - viewing 123
- licenses 54
 - Agent for Windows 778
- list boxes
 - navigation 1426
- list of servers
 - about 152
 - adding discovered servers 573
 - adding servers 153
 - removing servers 154
 - server groups 154
- list view 138
- LiveUpdate
 - about 121
 - running manually 122
 - scheduling automatic updates 121
- local Backup Exec server
 - breaking connection with 106
- local server properties
 - about viewing 627
 - viewing 627
- logon accounts
 - about 610
 - changing default 618
 - changing the password 616
 - checking 570

- logon accounts *(continued)*
 - copying to another server 620
 - creating 612
 - default
 - about 610
 - deleting 617
 - editing 614
 - replacing 616
 - restricted 612
 - scheduling a test 570
 - SQL databases 892
 - system logon account 611
 - testing 620
- logon information
 - copying to another server 620
- Lotus Domino
 - transaction logs
 - about 1072
- Lotus Domino Agent
 - Active File Exclusion 1071
 - APIs 1068
 - archive logging 1073
 - automatic exclusion of files
 - volume level backups 1071
 - backup options 1070
 - circular logging 1073
 - DAOS NLO
 - restore redirection 1079
 - database backup overview 1068
 - database backup requirements 1065
 - disaster preparation 1079
 - disaster recovery
 - archive logging 1081
 - circular logging 1084
 - of server 1079
 - Microsoft Cluster Server 1065
 - restoring 1076
 - monitor change journal
 - disabling 1081
 - re-enabling 1083
 - overview 1064
 - redirecting restore 1078
 - requirements 1064
 - restore options 1076
 - selecting for restore 1073
 - supported configurations 1071
- Lotus Domino transaction logs
 - viewing 1068

- low disk space thresholds
 - editing for a virtual disk 1280

M

- Mac Agent
 - manual install and uninstall 1225
- mailbox access requirements for Exchange 944
- mailbox groups
 - arranging for provisioning 1360
 - managing 1346, 1359
 - overview 1358
 - setting rules for archiving 1361
- maintenance contract information
 - about 124
 - managing customer numbers 125
- majority node in a cluster 712
- managed Backup Exec server
 - changing settings 1120
 - copying jobs to 1134
 - installing 1102
 - network interface card
 - using any available 1133
 - pools 1136
 - upgrading 1110
- Managed Backup Exec Servers report 667
- master database (SQL)
 - backup 894
- media
 - erasing 407
 - overwrite options 395
 - overwriting allocated or imported 395
 - retired
 - defined 380
 - scratch
 - defined 381
- Media Agent for Linux
 - beoper group 1244
 - installing 1244
- Media Audit report 669
- Media Errors report 670
- media ID
 - defined 398
- media overwrite protection level
 - defined 394
- Media Required for Recovery report 670
- media rotation
 - strategies
 - father/son 412
 - grandfather 412

- media rotation *(continued)*
 - strategies *(continued)*
 - son 412
- media set
 - creating 386
 - default 379
 - deleting 394
 - media vault rules 389
 - overwrite and append properties 389
 - renaming 394
 - vault rule properties 404
- Media Summary report 671
- Media Vault Contents report 671
- media vaults, about 402
- messages
 - error 691
- Microsoft Cluster Server
 - using with Backup Exec 711
- Microsoft incremental backups 866
- Microsoft SharePoint data
 - backing up 974
 - restoring 981
- Microsoft SQL Server 2008 R2 Express
 - installing 60
- Microsoft Virtual Hard Disk files
 - about managing 169
- Migration Report 129
- mnemonic
 - defined 1404
- modified time
 - using to determine backed up status 200
- Most Free Space First storage device in pool 419
- Move Media to Vaultreport 672
- MSCS
 - using with Backup Exec 711
- MSDE
 - 2008 components
 - installed with Backup Exec 60

N

- named transaction
 - restore up to 921
- navigation
 - list boxes 1426
 - Tabbed pages 1426
- NDMP Option
 - adding NDMP servers 1172
 - backing up NDMP servers 1174
 - duplicate backed up data 1179

- NDMP Option *(continued)*
 - how to use patterns 1177
 - installing 1172
 - overview 1170
 - redirecting restored data 1179
 - requirements 1171
 - restoring data 1179
 - setting default options 1183
 - sharing storage on NDMP servers 1174
 - viewing server properties 1184
 - viewing storage device properties 1185
- network
 - configuring options for backup jobs 204
 - overview of backup networks 574
- network and security
 - configuring 576
- network traffic
 - reducing in CASO 1118
- network-attached storage NDMP server
 - backing up 1170
- nodes
 - configurations in a Microsoft cluster 719
 - defined 710
 - disaster recovery using SDR 729
 - Microsoft
 - adding or removing a failover node 717
 - changing the order in which nodes fail over 717
- notification
 - configuring 295
 - configuring group recipient 299
 - editing recipient properties 301
 - enabling 297
 - removing recipient from group 300
 - sending for completed jobs 304
 - setting up 294
 - stopping 305
- notifications
 - overview 286
- Novell OES
 - about restoring 1208
 - requirements for back up 1207
 - supported components 1206
- NTFS
 - partition 693

O

- off-site storage of backups 146

- offhost backup
 - best practices 1165
 - configuring 1157
 - host computer
 - defined 1163
 - issues with hardware providers 1168
 - overview 1163
 - setting default options 1159
 - single volume snap 1159, 1161
 - transportable snapshots
 - defined 1163
 - troubleshooting 1166
- one-time conversion to virtual machine 505
- open files
 - unable to back up 692
- OpenStorage devices
 - disaster recovery of 828
 - editing properties 806
 - overview 804
 - requirements 798
- Operations Overview report 673
- optimized duplication 823
- Oracle Agent
 - authentication credentials 1005
 - deleting 1007
 - setting 1006
 - authentication credentials options 1007
 - authentication for Oracle operations 1005
 - back up with 1008
 - Backup Exec Agent Utility options 998
 - backup options 1011
 - configuring 994
 - DBA-initiated backup 1010
 - DBA-initiated job settings
 - create template for 597
 - DBA-initiated jobs
 - job template name for 1002
 - DBA-initiated restore 1014
 - default options 996
 - features 992
 - install 993
 - Linux servers
 - configuring an Oracle instance 1001
 - deleting an Oracle instance 1003
 - editing an Oracle instance 1003
 - enabling database access 1004
 - viewing an Oracle instance 1002
 - multiple data streams
 - specify 1012

- Oracle Agent *(continued)*
 - Oracle Net Service name 997
 - port
 - configure for Oracle operations 1005
 - publish Oracle databases on Linux 1002
 - Real Application Cluster (RAC) 1001, 1009
 - recovery catalog 997, 1002
 - redirected restore 1014
 - restore 1012
 - update credentials for instances 996, 1002, 1008
 - Windows computers
 - configuring an Oracle instance 995
 - deleting an Oracle instance 999
 - editing an Oracle instance 998
 - enabling database access 999
 - viewing an Oracle instance 997
- Overnight Archive Summary report 675
- Overnight Summary report 676
- overwrite protection period
 - changing 392
 - defined 383, 390
 - editing 388
 - setting for media set 390

P

- parameter files
 - creating 117
- partition
 - creating for robotic library 373
 - FAT 693
 - NTFS 693
 - removing or reconfiguring 377
- Partition Recovery Utility
 - about 1059
 - finding an archive ID 1060
 - log file location 1048
 - logs
 - about 1047
 - requirements 1059
 - running 1060
 - troubleshooting 1062
- pass phrases 585
- password
 - changing for logon account 616
- patterns in NDMP Option excludes 1177
- performance
 - increase during backups of remote Windows computers 778

- physical disk
 - creating a physical disk group 1270
 - viewing properties 1272
- PHYSICAL_ONLY utility 895
- point in time
 - conversion to virtual machine 503
- point in time log restore option
 - SQL Agent 921
- pop-up alerts
 - enabling or disabling 307
- Port number
 - changing for Remote Media Agent for Linux 1253
- ports used by Backup Exec
 - default 580
 - listening 582
- post-job command
 - configuring 548
 - for backup jobs 547
 - setting for backup job 550
- Post-Migration Report 129
- pre-installation checklist 57
- pre-job command
 - configuring 548
 - for backup jobs 547
 - setting for backup job 550
- pre-upgrade checklist 127
- preferred server configurations for Exchange DAGs
 - about 940
 - creating 941
 - deleting 942
 - designating a default 942
 - removing as default 942
- priority
 - about 262
 - changing for scheduled job 262
- Problem Files report 677
- properties
 - active job 257
- Publish
 - Linux computers to Backup Exec servers 1194
- publish
 - to Backup Exec servers
 - using Agent for Windows 784

R

- ralus.cfg
 - about, for the Agent for Linux 1195
 - about, for the Remote Media Agent for Linux computer 1248

- ralus.cfg (*continued*)
 - configuration options 1197
 - editing configuration options in 1196
 - for the Agent for Mac 1226
- Recently Written Media report 678
- recipients
 - configuring groups 299
 - deleting 301
 - editing 301
 - enabling email and text messages 297
 - removing from a group 300
 - stopping notification 305
- reclaiming disk space automatically 335
- Recover This Computer Wizard
 - requirements 764
 - running 764
- recovered jobs
 - setting thresholds 282
- recovered jobs custom error-handling rule 273
- Recovery Storage Group 962
- redirected restore
 - Exchange data 962
- Remote Administrator
 - installing 102
 - installing using the command line 103
 - running 105
- remote computers
 - push-installing 78
- Remote Media Agent for Linux
 - adding to Backup Exec database 1252
 - backing up data 1257
 - changing port number 1253
 - creating a simulated tape library 1259
 - deleting simulated tape library 1261
 - determining server status 1255
 - how it works 1242
 - ICMP ping 1253
 - managing simulated tape libraries from the command line 1262
 - requirements 1242
 - restoring data 1257
 - Simulated Tape Library options 1259
 - simulated tape library properties 1261
 - Tape Library Simulator Utility 1258
 - Command line switches 1263
 - troubleshooting 1264
 - uninstalling 1246
 - viewing properties 1255

Remote Media Agent for Linux *(continued)*

- viewing properties of simulated tape libraries 1260

renaming

- tape media labels 401

repair option 120

replicated catalogs in CASO 1118

reports

- Alert History 652

- Alert History by Backup Exec server 653

- Archive Job Success Rate 653

- Archive Selections by Archive Rules and Retention Categories 654

- Audit Log 655

- Backup Job Success Rate 655

- Backup Recommendations 655

- Backup Resource Success Rate 656

- Backup Sets by Media Set 656

- Backup Size By Resource 657

- copying custom reports 642

custom

- changing graph options 641

- changing grouping or sorting 640

- previewing 641

- custom, changing filters 638

- custom, creating 633

- Daily Device Utilization 658

- Deduplication Device Summary 659

- Deduplication Summary 660

- deleting 644

- Device Summary 660

- Disk Storage Summary report 661

- editing 644

- Error-Handling Rules 662

- Event Recipients 663

- Exchange Mailbox Group Archive Settings 664

- Failed Archive Jobs 664

- Failed Backup Jobs 665

- File System Archive Settings 666

- Jobs Summary 666

- list of standard reports 647

- Managed Backup Exec Servers 667

- Media Audit 669

- Media Errors 670

- Media Required for Recovery 670

- Media Summary 671

- Media Vault Contents 671

- Move Media to Vault 672

- Operations Overview 673

reports *(continued)*

- Overnight Archive Summary 675

- Overnight Summary 676

- overview 629

- printing 643

- Problem Files 677

- re-running completed report 644

- Recently Written Media 678

- Resource Protected Recently 678

- Resource Risk Assessment 679

- Restore Set Details by Resource 680

- Retrieve Media from Vault 680

- Robotic Library Inventory 681

- running 631

- saving 642

- Scheduled Server Workload 682

- scheduling 631

- Scratch Media Availability 682

- setting defaults 645

- Test Run Results 683

- Vault Store Usage Details 684

- Vault Store Usage Summary 685

- viewing completed 643

- viewing properties 646

Requirements

- Lotus Domino Agent 1064

requirements

- Agent for Hyper-V 862

- Backup Exec 63

- Central Admin Server Option 1098

- conversion to virtual machines 487

- Exchange Agent 931

- NDMP Option 1171

- off-host backup 1164

- Remote Media Agent for Linux 1242

- synthetic backup 1157

- user rights for backup jobs 151

- Resource Protected Recently report 678

- Resource Risk Assessment report 679

Restore

- creating restore jobs for the Remote Media Agent for Linux 1257

- Lotus Domino Agent

- redirecting Lotus Domino 1078

- Lotus Domino databases 1073

- Restore Set Details by Resource report 680

- Restore Wizard 232

restoring

- about restoring data 232

restoring *(continued)*

- ARCserve tapes 246
- canceling a restore job 247
- domain controller into an existing domain 242
- encrypted data 245
- Exchange data 962
- file system data 238
- from a backup set 236
- from a completed backup job 235
- from a server 235
- from storage device media 236
- media created with other backup software 246
- online restore of a Windows computer 239
- searching for data to restore 234
- Shadow Copy Components 244
- SQL master database 923
- starting the Restore Wizard 232
- System State data 240
- System State to a domain controller 241
- UEFI system partitions 245
- utility partitions 245
- restricted encryption keys
 - defined 585
- restricted logon accounts
 - about 612
- retensioning a tape 456
- retention categories
 - default retention category 1346
 - editing 1367
 - overview 1366
 - specifying properties 1368
- retired media
 - defined 380
- Retired Servers server group
 - about 154
 - adding servers 157
- retired tape media
 - moving damaged media 406
- Retrieve Media from Vault report 680
- robotic library
 - cleaning slot 373
 - creating partitions 373
 - example configuration 368
 - initializing when Backup Exec services start 372
 - inventory when Backup Exec services start 368
 - reconfiguring partitions 377
 - removing partitions 377
 - setting up hardware 367
 - using with Backup Exec 366

- Robotic Library Inventory report 681
- runtime scripts, for Agent for Linux 1213

S

SAN

- hardware errors 696
- resetting the SAN 696
- troubleshooting 693
- troubleshooting offline storage devices 694

scan

- detecting storage arrays 1309

schedule

- about 215
- availability window 215
- configuring for backup jobs 519
- conflicts 216
- deleting exclude dates 564
- exclude dates for all backups 562
- exclude dates for individual backup jobs 218
- exporting exclude dates 565
- importing a list of dates to exclude 562
- including and excluding dates 216
- including dates 216

scheduled jobs

- about managing and monitoring 253
- changing priority 262
- deleting 263
- holding 258
- removing a hold 259
- running immediately 261
- statuses 277

Scheduled Server Workload report 682

scratch media

- creating 395
- defined 381

Scratch Media Availability report 682

SCSI

- pass-through mode for tape drives 365
- setting address for robotic library drives 367

SCSI bus

- configuring for tape devices in a Microsoft cluster 722

Search 234

Search Wizard 234

Section 508 of the Rehabilitation Act

- compliance 1403

seeding a deduplication disk storage device 823

server groups

- about 154

- server groups *(continued)*
 - backing up 158
 - creating 155
 - editing 156
 - hiding 155
 - removing 158
 - retiring servers 157
 - viewing 155
- Server properties
 - Remote Media Agent for Linux 1255
- server properties
 - about viewing 627
 - viewing 628
- service account
 - changing credentials 621
- services
 - changing service account credentials 621
 - changing startup options 623
 - starting and stopping 621
- SGMon 707
- Shadow Copy Components
 - file system 187
 - restoring 244
- SharePoint Agent
 - about 971
 - adding a farm 974
 - backing up SharePoint data 974
 - disabling or enabling communication between
 - web servers and Backup Exec 983
 - disaster recovery for SharePoint 2007 989
 - disaster recovery for SharePoint 2010/2013 985
 - installing 971
 - overview 971
 - requirements 971
 - restoring SharePoint data 981
 - system requirements 971
 - using with SharePoint Server 2007 and Windows
 - SharePoint Services 3.0 973
 - using with SharePoint Server 2010/2013 and
 - SharePoint Foundation 2010/2013 972
- SharePoint farms
 - adding 974
 - viewing properties 984
- silent mode installation 108
- Simplified Disaster Recovery
 - Advanced Disk Configuration
 - about 773
 - boot managers 764
- Simplified Disaster Recovery *(continued)*
 - clusters
 - recovering Backup Exec on a Microsoft
 - Cluster 730
 - recovering nodes 729
 - Contents of the Simplified Disaster Recovery disk
 - image 760
 - editing the default path 748
 - enabling backups for 743
 - installing 736
 - Microsoft Exchange Server
 - recovering 771
 - Microsoft Hyper-V hosts 771
 - Microsoft SQL Server
 - recovering 771
 - OS/2 boot manager
 - restoring 764
 - overview 736
 - Recover This Computer Wizard
 - encrypted backup sets 764
 - restoring from a remote Backup Exec
 - server 764
 - recovering IBM computers 763
 - recovery requirements in SDR 764
 - requirements 737
 - setting an alternate location 747
 - SharePoint Portal Server
 - recovering 771
 - storage pools and storage spaces
 - recovering 769
- Simplified Disaster Recovery disk image
 - contents 760
- Simulated tape library
 - creating 1259
 - delete 1261
 - viewing properties 1260
- single block mode
 - setting for tape drives 364
- snapshot technology
 - using with Exchange Agent 948
- SNMP
 - configuring system service for Windows 311
 - installing WMI provider 312
 - object identifier prefix 308
 - traps
 - defined 308
 - traps for alerts 308
- son media rotation strategy 412
- Sort, filter, and copy 136

- splash screen
 - show at startup 566
- SQL
 - restore to named transaction 921
- SQL Agent
 - backing up
 - backup methods 909, 919
 - consistency check after backup 907, 917
 - consistency check recommendations 894
 - databases 899
 - strategies for 893
 - Windows registry 894
 - consistency check 895
 - recommendations 894
 - Database Consistency Check (DBCC)
 - recommendations 894
 - database snapshots
 - overview 897
 - disaster recovery 926
 - manual 928
 - features 891
 - installation 893
 - logon accounts 892
 - overview 891
 - requirements 892
 - restoring
 - master database 923
 - point in time log restore option 921
 - redirecting restores 921
 - TDE-encrypted database backups 921
 - very large databases 921
 - snapshot technology
 - using 896
 - strategy recommendations 893
- stages
 - about 219
 - adding 219
 - defined 158
 - editing 221
 - types 220
- stalled jobs
 - setting thresholds 282
- standard view 138
- storage
 - about sharing 448
 - about tape drives and robotic libraries 357
 - backup set retention by storage type 344
 - change to online 450
 - configuring for backup jobs 527
 - storage *(continued)*
 - disabling 455
 - editing global settings 439
 - enabling 455
 - Hot-swappable Device Wizard 358
 - Library Expansion Option 358
 - pausing 455
 - renaming 451
 - seleting 450
 - sharing deduplication devices 821
 - Symantec Device Driver Installation Wizard 360
 - unpausing 455
 - Virtual Tape Library Unlimited Drive Option 358
 - storage and media data
 - location of in CASO 1099
 - storage array
 - about identifying physical disks 1310
 - blink 1310
 - capacity 1272
 - configuring 1270
 - configuring virtual disks 1285
 - detecting 1309
 - hardware health 1272
 - hardware status 1272
 - identifying the physical disks 1310
 - renaming 1309
 - status 1272
 - viewing components 1270
 - viewing properties 1272
 - Storage device pools
 - any virtual disk storage device pool 1278
 - storage device pools
 - about 416
 - adding or removing devices 420
 - changing defaults 419
 - creating 416
 - Least Free Space First device 419
 - Most Free Space Firsst device 419
 - properties 421
 - system-defined 416
 - storage devices
 - installing 58
 - storage operations
 - cataloging 452
 - cleaning a drive 458
 - ejecting media 458
 - exporting expired media 466
 - exporting media 466
 - format WORM 456

- storage operations *(continued)*
 - importing media 459
 - initializing a robotic library 456
 - inventorying 453
 - inventorying and cataloging 454
 - locking the front portal 469
 - overview 424
 - retensioning 456
 - scanning 452
 - unlock front portal 469
- storage pools and storage spaces, recovering with SDR 769
- Storage Provisioning Option
 - configuring in CASO 1269
 - description 1268
 - detecting storage arrays 1309
 - installing 1269
 - requirements 1269
 - upgrading 1269
- storage trending 316
- Symantec Device Driver Installation Wizard 360
- Symantec Endpoint Protection
 - using with Backup Exec 545
- Symantec Help Tool 702
- Symantec Knowledge Base
 - searching 699
- Symantec RSS Reader
 - customizing 143
 - options 143
 - overview 143
 - removing default RSS feed 143
 - viewing articles 143
- synchronizing archiving permissions and settings 1350
- synthetic backup
 - baseline 1156
 - encryption
 - requirements for 1157
 - requirements 1157
- system logon account
 - about 611
 - creating 619
- system requirements
 - Backup Exec 63
- System State
 - restoring 240
 - restoring to a domain controller 241

T

- Tabbed dialog boxes
 - navigation 1426
- Tabs 135
- tape drives
 - buffer count 364
 - buffer size 363
 - default settings 364
 - high water count 364
 - statistics 366
- Tape Library Simulator Utility
 - creating a simulated tape library 1259
 - deleting library 1261
 - overview 1258
 - running from command line 1262
 - viewing properties 1260
- tape media
 - associating with a media set or vault 409
 - damaged 406
 - deleting 407
 - scanning barcode labels 406
 - with excessive errors 406
- tape media label
 - barcodes 400
 - imported 400
 - overview 398
 - renaming 401
- tape media operations
 - associating media with media sets 409
- tapeinst.exe 360
- tapes
 - DLT tape drive 690
- TCP/IP
 - required for Mac agent 1219
- technical support
 - contacting 700
- test run job
 - about 226
 - configuring 533
 - running 226
- Test Run Results report 683
- text message notification
 - configuring 295
- ThreatCon levels 545
- transaction logs
 - Lotus Domino DBIID 1073
 - reuse
 - Lotus Domino 1071
- transform, creating 94

- Transparent Database Encryption
 - SQL Agent 921
- tree view 138
- trial version
 - agents and options 118
- Troubleshooting
 - Remote Media Agent for Linux 1264
- troubleshooting
 - Backup Exec performance
 - improving 698
 - backup issues 691
 - clusters 734
 - error messages 691
 - hardware-related issues 687
- true image restore
 - overview 1161
- trust
 - establishing 780
 - establishing for a remote Linux computer 1193
 - establishing for a Remote Media Agent for Linux
 - computer 1248
- typical installation 66

U

- unconfigured virtual disk
 - configuring 1285
 - hardware health 1285
 - hardware status 1285
 - viewing properties 1285
- uninstallation
 - Backup Exec 130
 - Backup Exec options from local Backup Exec
 - server 131
 - using command line 131
- uninstalling
 - Backup Exec from a Microsoft cluster 716
- Unique Message Identifier (UMI) error code
 - viewing 270
- unrecognized media 379
- updates
 - installing to Agent for Windows 93
 - viewing what is installed 123
- upgrades
 - checklist 127
 - overview 125
- USB tape devices
 - reconnecting 358
- utility partitions
 - restoring 245

V

- vault
 - scan barcode labels to move tape media 406
- vault rules for media sets 404
- vault store
 - assigning 1345
 - changing item deletion mode 1352
 - deleting 1353
 - editing properties 1352
 - viewing status 1352
- vault store group
 - backing up for Archiving Option 1377
- vault store partitions
 - backing up for Archiving Option 1377
 - editing properties 1355
 - open and closed 1354
 - overview 1354
 - viewing open and closed states 1355
- Vault Store Usage Details report 684
- Vault Store Usage Summary report 685
- vault stores
 - backing up for Archiving Option 1377
 - fingerprint database 1351
 - overview 1351
- verify operation
 - configuring for backup jobs 534
- version of Backup Exec, displaying 136
- vhd files
 - about managing 169
- vhd files
 - with GRT 865
- Viewing information on the administration console 138
- virtual disk
 - blink 1310
 - capacity 1297
 - concurrent jobs 1297
 - configuring 1285
 - creating 1270
 - description 1278
 - editing default options 1280
 - editing general properties 1292
 - editing low disk space thresholds 1280
 - hardware health 1297
 - hardware status 1297
 - identifying the physical disks 1310
 - number of files 1297
 - renaming 1309
 - status 1297
 - unconfigured virtual disk 1285

- Virtual Disk Service
 - installing for the Storage Provisioning Option 1269
- virtual machine conversion
 - adding a stage for 498
 - after a backup job 492
 - from point in time 503
 - overview 479
 - requirements 487
 - setting default options 509
 - simultaneous with backup job 488
- virtual machines
 - automatic protection for Hyper-V 880
- virtual tape library
 - DirectCopy to physical devices 229–230
- Virtual Tape Library Unlimited Drive Option, about 358
- VMware vCenter Server, adding 834
- VMware virtual machines
 - push-installing Agent for Windows 837
- volume level backups
 - automatic exclusion of SQL data 899
- VSS
 - perform consistency check before Exchange backup 953, 958
 - using to protect Exchange data 948

W

- Windows change journal
 - defaults set by Backup Exec 202
 - resolving errors 202
 - using to determine backed up status 201
- Windows deduplication volume
 - backing up 168
- Windows Management Instrumentation (WMI)
 - adding WMI capability 311
- Windows registry
 - backing up with SQL Agent 894
- Windows Server 2003
 - backing up 187
- Windows Server 2008
 - backing up 187
 - Read Only Domain Controller 63
 - Server Core 63
- Windows Server 2012
 - with the Agent for Hyper-V 867
- Windows Server 2012 R2
 - with the Agent for Hyper-V 868
- Windows user rights 151

- WMI
 - installing performance counter provider 311
 - installing SNMP provider 312
 - uninstalling performance counter provider 312
 - uninstalling SNMP provider 312
- WORM media, about 401