Veritas Data Insight Administrator's Guide

6.0

June 2017

VERITAS™

Documentation version: 6.0.0

Legal Notice

Copyright © 2017 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third party software for which Veritas 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. Refer to the third party legal notices document accompanying this Veritas product or available at:

https://www.veritas.com/about/legal/license-agreements

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 Veritas Technologies LLC 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. VERITAS TECHNOLOGIES LLC 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, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. 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.

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

http://www.veritas.com

.

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

https://www.veritas.com/support

You can manage your Veritas account information at the following URL:

https://my.veritas.com

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)	CustomerCare@veritas.com
Japan	CustomerCare_Japan@veritas.com

Contents

Section 1 Getting started 13

Chapter 1 Introduction to Veritas Data Insight administration 14

About Veritas Data Insight administration 14
Operation icons on the Management Console 14
Data Insight administration tasks 15
Supported file servers and platforms 17

Chapter 2 Configuring Data Insight global settings 20

Overview of Data Insight licensing 21
Managing Data Insight licenses 23
Configuring SMTP server settings 23
About scanning and event monitoring 24
Configuring scanning and event monitoring 28
About filtering certain accounts, IP addresses, and paths 33
About exclude rules for access events 33
About Exclude rules for Scanner 34
Adding exclude rules to Data Insight 35
About archiving data 37
About purging data 38
Configuring data retention settings 38
About Data Insight integration with Symantec Data Loss Prevention (DLP) 41
About configuring Data Insight to integrate with Data Loss Prevention (DLP) 42
Configuring Symantec Data Loss Prevention settings 44
Importing SSL certificate from the DLP Enforce Server to Data Insight Management Server 45
About Symantec Data Loss Prevention (DLP) integration with Data Insight 46
Importing sensitive files information through CSV 47
Configuring advanced analytics 48
Choosing custom attributes for advanced analytics 49

About open shares 50
Configuring an open share policy 51
Configuring file groups 52
Configuring Workspace data owner policy 53
Configuring Management Console settings 54
About bulk assignment of custodians 55
Assigning custodians in bulk using a CSV file 56
Assigning custodians based on data ownership 56
Configuring Watchlist settings 57

Section 2 Configuring Data Insight 59

Chapter 3 Configuring Data Insight product users 60

About Data Insight users and roles 60
Reviewing current users and privileges 61
Adding a user 62
Add or edit Data Insight user options 62
Editing users 63
Deleting users 64
Configuring authorization for Symantec Data Loss Prevention users 64

Chapter 4 Configuring Data Insight product servers 65

About Data Insight product servers 66
Adding a new Data Insight server 66
Managing Data Insight product servers 67
Viewing Data Insight server details 70
About node templates 71
Managing node templates 71
Adding or editing node templates 72
Adding Portal role to a Data Insight server 73
Adding Classification Server role to a Data Insight server 74
Assigning Classification Server to a Collector 74
Viewing in-progress scans 75
Configuring Data Insight services 76
Configuring advanced settings 78
Monitoring Data Insight jobs 96
Rotating the encryption keys 97
Viewing Data Insight server statistics 98
About automated alerts for patches and upgrades 99
Viewing and installing recommended upgrades and patches 100

Deploying upgrades and patches remotely 100
Using the Upload Manager utility 102
About migrating storage devices across Indexers 103
Viewing the status of a remote installation 104

Chapter 5 Configuring saved credentials 105

About saved credentials 105
Managing saved credentials 105
Handling changes in account password 107

Chapter 6 Configuring directory service domains 108

About directory domain scans 108
Adding a directory service domain to Data Insight 109
Add/Edit Active Directory options 109
Add/Edit LDAP domain options 110
Add/Edit NIS domain options 113
Add/Edit NIS+ domain options 113
Managing directory service domains 114
Fetching users and groups data from NIS+ scanner 115
Configuring attributes for advanced analytics 116
Deleting directory service domains 118
Scheduling scans 119
Configuring business unit mappings 119
Importing additional attributes for users and user groups 120

Chapter 7 Configuring containers 122

About containers 122
Adding containers 122
Add new container/Edit container options 123
Managing containers 123

Section 3 Configuring storage devices in Data Insight 125

Chapter 8 Configuring NetApp file server monitoring 126

About configuring NetApp file server monitoring 127
Prerequisites for configuring NetApp file servers 127
Credentials required for configuring NetApp filers 128
Credentials required for configuring NetApp NFS filers 132
Configuring SMB signing 133

About FPolicy 134
Preparing Data Insight for FPolicy 135
Preparing the NetApp filer for Fpolicy 136
Preparing the NetApp vfiler for Fpolicy 138
Preparing a non-administrator domain user on the NetApp filer for Data Insight 141
Enabling export of NFS shares on a NetApp file server 144
Excluding volumes on a NetApp file server 145
Handling NetApp home directories in Data Insight 145

Chapter 9 Configuring clustered NetApp file server monitoring 147

About configuring a clustered NetApp file server 147
About configuring FPolicy in Cluster-Mode 149
Pre-requisites for configuring clustered NetApp file servers 149
Credentials required for configuring a clustered NetApp file server 151
Preparing a non-administrator local user on the clustered NetApp filer 153
Preparing a non-administrator domain user on a NetApp cluster for Data Insight 155
Preparing Data Insight for FPolicy in NetApp Cluster-Mode 157
Preparing the ONTAP cluster for FPolicy 158
About configuring secure communication between Data Insight and cluster-mode NetApp devices 159
Generating SSL certificates for NetApp cluster-mode authentication 160
Preparing the NetApp cluster for SSL authentication 162
Enabling export of NFS shares on a NetApp Cluster-Mode file server 165

Chapter 10 Configuring EMC Celerra or VNX monitoring 167

About configuring EMC Celerra or VNX filers 167
About EMC Common Event Enabler (CEE) 168
Preparing the EMC filer for CEPA 168
Preparing Data Insight to receive event notification 171
Credentials required for configuring EMC Celerra filers 173

Chapter 11 Configuring EMC Isilon monitoring 175

About configuring EMC Isilon filers 175
Prerequisites for configuration of Isilon file server monitoring 176

Credentials required for configuring an EMC Isilon cluster 177
Configuring audit settings on EMC Isilon cluster using OneFS GUI console 178
Configuring audit settings on EMC Isilon cluster using the OneFS CLI 180
Configuring Isilon audit settings for performance improvement 182
Preparing Veritas Data Insight to receive event notifications from an EMC Isilon cluster 183
Creating a non-administrator user for an EMC Isilon cluster 184
Using SmartConnect mapping for access zones 185
Purging the audit logs in an Isilon filer 186

Chapter 12 Configuring Hitachi NAS file server monitoring 188

About configuring Hitachi NAS 188
Credentials required for configuring a Hitachi NAS EVS 189
Creating a domain user on a Hitachi NAS file server for Data Insight 189
Preparing a Hitachi NAS file server for file system auditing 190
Advanced configuration parameters for Hitachi NAS 191

Chapter 13 Configuring Windows File Server monitoring 193

About configuring Windows file server monitoring 193
Credentials required for configuring Windows File Servers 194
Using the installcli.exe utility to configure multiple Windows file servers 196
Upgrading the Windows File Server agent 198

Chapter 14 Configuring Veritas File System (VxFS) file server monitoring 199

About configuring Veritas File System (VxFS) file servers 199
Credentials required for configuring Veritas File System (VxFS) servers 200
Enabling export of UNIX/Linux NFS shares on VxFS filers 202

Chapter 15 Configuring monitoring of a generic device 204

About configuring a generic device 204
Credentials required for scanning a generic device 205

Chapter 16 Managing file servers 207

About configuring filers 208
Viewing configured filers 208
Adding filers 210
Add/Edit NetApp filer options 211
Add/Edit NetApp cluster file server options 215
Add/Edit EMC Celerra filer options 220
Add/Edit EMC Isilon file server options 223
Add/Edit Windows File Server options 225
Add/Edit Veritas File System server options 229
Add/Edit a generic storage device options 232
Add/Edit Hitachi NAS file server options 233
Custom schedule options 235
Editing filer configuration 236
Deleting filers 238
Viewing performance statistics for file servers 238
About disabled shares 239
Adding shares 239
Add New Share/Edit Share options 240
Managing shares 241
Editing share configuration 245
Deleting shares 245
About configuring a DFS target 245
Adding a configuration attribute for devices 246
Configuring a DFS target 246
About the DFS utility 247
Running the DFS utility 248
Importing DFS mapping 249

Chapter 17 Configuring SharePoint monitoring 250

About SharePoint server monitoring 250
Credentials required for configuring SharePoint servers 251
Configuring a web application policy 252
About the Data Insight web service for SharePoint 253
Installing the Data Insight web service for SharePoint 254
Viewing configured SharePoint web applications 255
Adding web applications 256
Add/Edit web application options 257
Editing web applications 260
Deleting web applications 261
Adding site collections 261
Add/Edit site collection options 262

Managing site collections 263
Removing a configured web application 267

Chapter 18 Configuring cloud content sources 268

About configuring Box monitoring 268
Using a co-admin account to monitor Box resources 269
Viewing configured cloud sources 270
Configuring Box monitoring in Data Insight 272
Add/Edit Box account 272
Managing cloud storage accounts 274

Chapter 19 Renaming storage devices 276

About renaming a storage device 276
Viewing the device rename status 277
Considerations for renaming a storage device 277

Section 4 Health and monitoring 279

Chapter 20 Using Veritas Data Insight dashboards 280

Viewing the system health overview 280
Viewing the scanning overview 283
Viewing the scan status of storage devices 285
Viewing the scan history of storage devices 286

Chapter 21 Monitoring Data Insight 288

Viewing events 288
About high availability notifications 289
Monitoring the performance of Data Insight servers 289
Configuring email notifications 291
Enabling Windows event logging 292
Viewing scan errors 293

Section 5 Alerts and policies 295

Chapter 22 Configuring policies 296

About Data Insight policies 296
Managing policies 298
Create Data Activity Trigger policy options 299
Create User Activity Deviation policy options 301

Create Data Activity User Whitelist-based policy options 303
Create Data Activity User Blacklist-based policy options 305
Create Real-time Sensitive Data Activity policy options 308
Managing alerts 311

Section 6 Remediation 313

Chapter 23 Configuring remediation settings 314

About configuring permission remediation 314
Managing and configuring permission remediation 315
Configuring exclusions for permission recommendation 317
About managing data 318
About configuring archive options for Enterprise Vault 318
Adding new Enterprise Vault servers 320
Managing Enterprise Vault servers 320
Mapping file server host names 322
Viewing and managing the status of an operation 323

Section 7 Reference 326

Appendix A Backing up and restoring data 327

Selecting the backup and restore order 327
Backing up and restoring the Data Insight Management Server 327
Backing up and restoring the Indexer node 329

Appendix B Data Insight health checks 332

About Data Insight health checks 332
Services checks 333
Deployment details checks 333
Generic checks 334
Data Insight Management Server checks 338
Data Insight Indexer checks 339
Data Insight Collector checks 342
Data Insight Windows File Server checks 345
Data Insight SharePoint checks 347
Understanding Data Insight best practices 347

Appendix C Command File Reference 350

fg.exe 351
indexcli.exe 353

reportcli.exe 359
scancli.exe 363
installcli.exe 368

Appendix D Data Insight jobs 370

Scheduled Data Insight jobs 370

Appendix E Troubleshooting 380

About general troubleshooting procedures 380
About the Health Audit report 381
Location of Data Insight logs 381
Downloading Data Insight logs 383
Migrating the data directory to a new location 384
Enterprise Vault exceptions and their meanings 385
Troubleshooting FPolicy issues on NetApp devices 389
    Viewing FPolicy-related errors and warnings 390
    Resolving FPolicy connection issues 390
Troubleshooting EMC Celera or VNX configuration issues 391
Troubleshooting EMC Isilon configuration issues 394
Troubleshooting SharePoint configuration issues 395
Troubleshooting Hitachi NAS configuration issues 399

Index 400

Section 1

Getting started

- Chapter 1. Introduction to Veritas Data Insight administration
- Chapter 2. Configuring Data Insight global settings

Chapter 1

Introduction to Veritas Data Insight administration

This chapter includes the following topics:

- About Veritas Data Insight administration

About Veritas Data Insight administration

You administer the Veritas Data Insight system through the Management Console. The console has components for system administration, viewing data access information, configuring policies and alerts, and generating reports, which are accessible from the tabs located on the header panel. Navigate to the **Settings** tab on the console to carry out the various Data Insight administration tasks.

The Console is automatically installed with the Management Server. You access the Console through a Web browser that has a network connection to the Management Server. By default, the Management Server runs on HTTPS port 443. To access it, in the Web browser's address field, type `https://ms-host/`.

Provided a valid license is installed, the Server Administrator user can see and access all parts of the administration console. Other users can see only the parts to which their roles grant them access. The user account under which you are currently logged on appears at the footer of the Management Console screen.

Operation icons on the Management Console

Table 1-1 shows the operation icons that are located on the console screen:

Table 1-1 Operation icons on the Management Console

Icon	Description
	The settings icon is used in assigning custodians.
	Screen refresh. Veritas recommends using this refresh button instead of your browser's **Refresh** or **Reload** button.
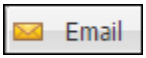	Email the data on the current screen to one or more recipients. If the current screen's data cannot be sent as an email, the icon is unavailable.
	Exports all data on a panel on the current screen to a `.csv` file.
	Exports all data on the current screen to a `.csv` file.
	Submits request to the Enterprise Vault server to archive the selected folders.
	The action selector icon displays a menu with the following two options: ■ Archive files using Enterprise Vault. ■ Submit request to invoke a custom action on selected paths.
	Submit request to invoke a custom action on selected paths.

Data Insight administration tasks

Table 1-2 summarizes the tasks to be performed to set up a new Data Insight installation:

Table 1-2 Data Insight administration tasks

Action	Description
Configure SMTP server settings.	See "Configuring SMTP server settings " on page 23.
Setup notification policies.	See "Configuring email notifications " on page 291.

Table 1-2 Data Insight administration tasks *(continued)*

Action	Description
Configure directory service domain.	See "Adding a directory service domain to Data Insight" on page 109.
Configure data retention settings.	See "Configuring data retention settings" on page 38.
Configure Exclude Rules.	See "Adding exclude rules to Data Insight " on page 35.
Install license.	See "Managing Data Insight licenses " on page 23.
Configure Data Insight nodes either individually or configure multiple nodes by applying node templates.	See "About node templates" on page 71. See "Managing Data Insight product servers" on page 67.
If monitoring events for NetApp file servers, configure Fpolicy service on collectors.	See "Preparing Data Insight for FPolicy " on page 135.
If monitoring events for a clustered NetApp file server, install the DataInsightFPolicyCmod service on the collectors.	See "Preparing Data Insight for FPolicy in NetApp Cluster-Mode" on page 157.
If monitoring events for EMC Celerra file servers, configure Celerra service on collectors.	See "About EMC Common Event Enabler (CEE)" on page 168.
If monitoring events for EMC Isilon file servers, configure EMC Common Event Enabler (CEE) on the collectors.	See "Preparing Veritas Data Insight to receive event notifications from an EMC Isilon cluster" on page 183.
If monitoring events for Windows file servers, upload agent packages to collectors.	See "About configuring Windows file server monitoring" on page 193.
If monitoring events for a generic device, use web APIs to collect access event information.	See "About configuring a generic device" on page 204.
If monitoring events for SharePoint servers, install the Data Insight Web service on the SharePoint server.	See "Installing the Data Insight web service for SharePoint" on page 254.
If monitoring events from cloud storage sources, configure authorization to the cloud storage account	See "About configuring Box monitoring" on page 268.

Table 1-2 Data Insight administration tasks *(continued)*

Action	Description
Configure file servers. Configure the SharePoint web applications.	See "Adding filers" on page 210. See "Adding web applications" on page 256.
Configure advanced analytics	See "Configuring advanced analytics " on page 48.
Configure archiving settings	See "About configuring archive options for Enterprise Vault" on page 318.
Set up classification	Do the following: ■ Install the Classification Server or assign classification role to a Data Insight server. See "Adding Classification Server role to a Data Insight server" on page 74. ■ Enable classification and configure safeguard settings for the Classification Server. For details, see the *Veritas Data Insight Classification Guide* ■ Configure settings for content scanning. See "Configuring advanced analytics " on page 48.
Configure permission remediation settings	See "Managing and configuring permission remediation" on page 315.
Configure workflows to remediate DLP incidents, review permissions on data resources, confirm ownership of files and folders, and to classify sensitive files for retention.	Do the following: **1** Install the Self-Service Portal. For details, see the *Veritas Data Insight Installation Guide*. **2** Configure workflow templates. **3** Create remediation workflows and submit them for further action on the Self-Service Portal.

Supported file servers and platforms

Table 1-3 lists the Network Attached Storage (NAS) devices and SharePoint platforms that Data Insight supports.

Table 1-3 Supported file servers and platforms

Device	Version
Hitachi NAS	Hitachi NAS 12.x
NetApp ONTAP 7-Mode	Version 7.3.5 or higher
NetApp ONTAP Cluster-Mode	CIFS - ONTAP 8.2.x or higher NFS - ONTAP 8.2.3 or higher and ONTAP 8.3.1 or higher
EMC	EMC Celerra version 5.6.45 or higher
	EMC Isilon OneFS version 7.1.0.6 or higher
	VNX version 7.1.71.1 or higher
Windows File Server	Windows Server 2008, or 2008 R2, 32 bit and 64 bit Windows Server 2012, or 2012 R2 64 bit **Note:** For Windows 2008 32-bit, use Data Insight Windows Filer Server Agent version 4.5 in backward compatibility mode.
Veritas File System (VxFS) server	6.0.1 or higher, configured in standalone or clustered mode using Cluster Server (VCS) **Note:** For VCS support, Clustered File System (CFS) is not supported.
Microsoft SharePoint	Microsoft SharePoint Server 2007 Microsoft SharePoint Server 2010 Microsoft SharePoint Server 2013 Microsoft SharePoint Server 2016
Box (Cloud-based content management platform)	-
Symantec Data Loss Prevention (DLP)	Versions 12.0.1, 12.5, 14.0, 14.5, 14.6, and 14.6 MP1 **Note:** Data Insight 6.0 does not support Box integration with DLP 14.6 and 14.6 MP1.
Enterprise Vault	Versions 11.0, 11.0.1, 12.0, and 12.1

Note the following:

- Veritas strongly recommends that you upgrade your NetApp filer to the latest available firmware. Veritas recommends ONTAP 7.3.5 or higher.

- For all supported versions of 7-mode NetApp filers, Data Insight supports CIFS protocol over NTFS and NFS protocol v3. NFS v4 is not supported.

  For supported versions of Cluster-Mode NetApp filers, Data Insight supports the following volume/qtree styles:

  - NTFS and Mixed for CIFS protocol.
  - UNIX and Mixed for NFS protocol on 7-mode NetApp filers only.
  - NFS exports on the NetApp cluster.
- For all supported versions of EMC Celerra/VNX and EMC Isilon, Data Insight supports only CIFS protocol over NTFS. Data Insight supports Common Event Enabler (CEE), version 8.2 or higher. Data Insight still supports the older version of CEE and VEE, but Veritas recommends that you move to the latest EMC Common Event Enabler, which you can download from the EMC website.
- To use the Self-Service Portal to remediate DLP incidents, ensure that Symantec Data Loss Prevention (DLP) version 12.5 or higher is installed. Data Insight uses the DLP Smart Response Rules to remediate incidents, which are introduced in DLP version 12.5.

Chapter 2

Configuring Data Insight global settings

This chapter includes the following topics:

- Overview of Data Insight licensing
- Configuring SMTP server settings
- About scanning and event monitoring
- About filtering certain accounts, IP addresses, and paths
- About archiving data
- About Data Insight integration with Symantec Data Loss Prevention (DLP)
- Importing sensitive files information through CSV
- Configuring advanced analytics
- About open shares
- Configuring file groups
- Configuring Workspace data owner policy
- Configuring Management Console settings
- About bulk assignment of custodians
- Configuring Watchlist settings

Overview of Data Insight licensing

You must install a license key file that covers the Data Insight features that you want to implement. To use all Data Insight functionality, you must have the following licenses:

- Data Insight base license – The base product license enables you to access basic product features including viewing audit and scan data, reports, classifying content, and adding resources such as filers and SharePoint web applications.

  The Data Insight base license is available in two types - based on the number of users and based on the amount of storage that is monitored by Data Insight. When the storage limit-based license is applied, Data Insight computes the total storage that it monitors on filers, generic devices, and SharePoint servers. It also includes the storage space that is occupied by disabled shares. However, it does not include the storage monitored on DFS paths and cloud storage resources such as Box.

  For paths on generic devices, Data Insight considers the storage occupied by parent and child shares independently. This behavior is observed because for generic devices, Data Insight is not able to distinguish between the physical and the logical paths. As a result, during the computation of the total storage capacity, Data Insight considers the nested storage as independent storage.

- Self-Service Portal license – A separate add-on license is needed if you want to create workflow templates and workflows. The Self-Service Portal license enables you to access the portal to take action on remediation workflows. You will also not be able to access the Self-Service Portal until valid base and portal licenses are applied.

Until Data Insight 5.0.x, the product functionality operates on Trust and Verify model. A message in the footer indicates whether the base as well as Self-Service portal license has expired, but product features are fully functional if the license has expired or in the absence of a valid license.

From Data Insight 5.1 onwards, if you do not have a valid base license or when the license expires, the software continues to run in a restricted mode. The following functionality is disabled on the Management Console:

- The **Workspace**, **Policies**, and **Reports** tab are hidden on the Management Console. Only the **Workflows** and **Settings** tabs are visible on the UI. However, you will not be able to take any remediation actions from the **Workflows** > **Audit** page.

- The options on the console to add new filers, shares, SharePoint web applications, site collections, cloud storage resources, and directory services are disabled. However, Data Insight lets you edit properties of the resources that are already configured in Data Insight.

- Auto-discovery does not result in discovering new shares or site collections that are created on filer or web application if the license on the Data Insight system is not valid.
- Creation of workflow templates and workflows is disabled, even if portal license is valid but the base license has expired.

Note that if the portal license is not valid:

- When you click **Workflows** > **Self-Service Portal Settings**, **Workflows** > **Templates**, or **Workflows** for the first time, a pop-up is displayed in the bottom right corner indicating that a valid Self-Service Portal license must be applied to use the feature.
- The options to add workflow templates and workflows are disabled.
- Custodians for a workflow continue to receive reminder emails for existing workflows. However, when you click the portal link in the email, they see a message indicating that the license has expired and must be renewed.
- The footer of the reminder email contains a message indicating that the license has expired.

Depending on the state of the license, the following warnings and errors are displayed in the footer of the Management Console:

- In case of a license in grace period, the footer shows an icon and message indicating that the license has expired and is in grace period. On clicking the footer, a detailed message indicating the number of grace days and the fact that features will be disabled after expiry is displayed. Additionally, a Data Insight event is raised daily at the WARNING level indicating that the license is in grace period.

  After the grace period ends, a warning period will then be in effect until the license completely expires.
- In case of a license in warning period, Data Insight displays an icon and a message indicating that the license is in the warning period and is due to expire in a certain number of days. The Licensing team configures the warning period. Click the footer icon or message to view the detailed message. Additionally, a Data Insight event is raised daily at the INFO level indicating that a warn policy is in effect for the license.
- On license expiry, the footer displays the message indicating that the license has expired. On clicking the message, a detailed message is displayed indicating that some of the features are disabled. Additionally, a Data Insight event is raised daily at the CRITICAL level stating that the license has expired.

  After the license expires, certain functionality is disabled and the application runs in a restricted mode.

- In case of a storage limit-based license, if the storage monitored by Data Insight exceeds the licensed limit then the footer displays a warning message indicating that the storage monitored by Data Insight has exceeded. However, Data Insight continues to function with uninterrupted access to all functionality. Click the footer message to view the detailed message.

In case of multiple error conditions, the error message pop-up populates the messages based on their severity status.

For more information about licenses, contact your Licensing team.

Managing Data Insight licenses

When you purchase Veritas Data Insight, you must install the Data Insight license file. License files have names in the format `name.slf`.

To install a license

1. Obtain the new license file.
2. In the Management Console, click **Settings** > **Licensing**.
3. On the **Licensing** page, click **Add/Update License**.
4. On the **Add new license** pop-up, browse to the new Data Insight license file that you have downloaded, and click **Upload**.

Configuring SMTP server settings

Before Data Insight can send email notifications for events, reports, workflows, and alerts you must configure SMTP details for the Management Server.

To edit the SMTP settings

1. In the Management Console, click **Settings** > **SMTP Settings**.
2. On the SMTP settings page, click **Edit**.
3. Enter the following details:
   - A valid SMTP server hostname or IP address.
   - The port number for the SMTP mail server used to send email notifications. The default is 25.
   - The username for the email server (optional).
   - The password for the email server (optional).
   - The address from which emails are sent (optional).

- Maximum attachment size. This information is used when Data Insight sends report notifications. Data Insight will not send reports as attachments, if the size of the report is over the specified limit.

4 Click **Save**.

About scanning and event monitoring

Data Insight scans the file system hierarchy to collect information related to permissions and file system metadata from the monitored storage devices.

Event monitoring is an operation that keeps track of the access events happening on a file system. During event monitoring if Data Insight detects an event such as create, write or file system ACL level permission changes, it uses this information to perform incremental scans for the paths on which events are reported.

Data Insight uses asynchronous APIs, such as FPolicy for NetApp filers, the CEE framework for EMC filers, and filter driver for Windows File Servers to collect access events.

By default, Data Insight initiates event monitoring every 2 hours. You can disable event monitoring for the individual storage devices. To turn off event monitoring, navigate to **Settings** > **Filers**. In the edit page for filer, uncheck the option **Enable file system event monitoring**.

Note: Data Insight scans only share-level permission changes when event monitoring is turned off.

To fetch file system metadata, Data Insight performs the following types of scans:

Full scan

During a full scan Data Insight scans the entire file system hierarchy. A full scan is typically run after a storage device is first added to the Data Insight configuration. Full scans can run for several hours, depending on the size of the shares. After the first full scan, you can perform full scans less frequently based on your preference. Ordinarily, you need to run a full scan only to scan those paths which might have been modified while file system auditing was not running for any reason.

In case of large shares, a full scan can take long time to complete. To reduce the time it takes to scan certain large shares, you can configure parallel scanning which uses multiple threads to scan the share. By default, the single thread scanner runs on the shares regardless of their size. To use the parallel scanner feature, you must configure it for each share.

Note: Parallel scanner supports only full scans on CIFS.

To configure parallel scanning for a Collector or a filer from the **Data Insight Servers** > **Advanced Settings** tab.

See "Configuring advanced settings" on page 78.

To configure multiple threads to scan a share:

See "Add New Share/Edit Share options " on page 240.

By default, each Collector node initiates a full scan at 7:00 P.M. on the last Friday of each month.

For SharePoint, the default scan schedule is 11:00 P.M. each night.

Figure 2-1 Scanner - Single thread and parallel threads

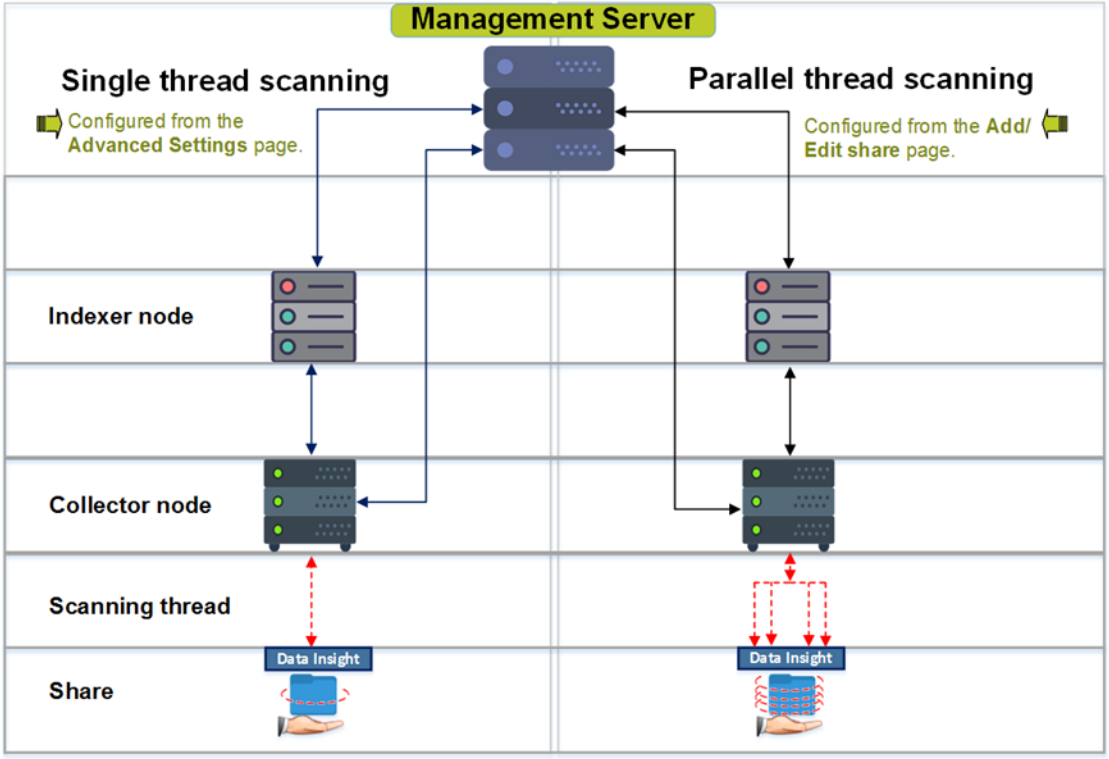

Incremental scan

During an incremental scan, Data Insight re-scans only those paths of a share that have been modified since the last full scan. It does so by monitoring incoming

access events to see which paths had a create, write, or a security event on it since the last scan. Incremental scans are much faster than full scans.

Note: For Data Insight versions before version 5.0, incremental scans were triggered only when Data Insight detected any events during event monitoring.

Incremental scans are not available for SharePoint web applications and for the cloud-based storage from Box.

By default, an incremental scan is scheduled once every night at 7:00 P.M. You can initiate an on-demand incremental scan manually by using the command line utility `scancli.exe`. It is recommended to run the IScannerJob before you execute the utility.

See "Scheduled Data Insight jobs" on page 370.

Path re-confirmation scan

After Data Insight completes indexing the full scan data, it computes the paths that no longer seem to be present on the file system. A re-confirmation scan confirms if a path which is present in the indexes, but appears to be no longer present on the file system, is indeed deleted. A re-confirmation scan is automatically triggered, when Data Insight detects potentially missing paths on the file system during a full scan.

You can turn off re-confirmation scan for any Indexer, using the **Advanced Setting** for that Indexer. When the re-confirmation scan is turned off, Data Insight readily removes the missing paths from the indexes without carrying out a re-confirmation.

See "Configuring advanced settings" on page 78.

At a global level, full scans are scheduled for individual Collectors or Windows File Server agents. The Table 2-1 gives you the details of all the entities for which you can schedule a full scan.

Table 2-1 Entities having configurable scan schedules

Entity	Scan schedule settings location	Scope	Details
Collector or Windows File Server agents	**Settings** > **Data Insight Servers** > **Advanced Setting** > **File System Scanner settings**.	Applies to all the storage devices associated with the Collector, for which a schedule is defined.	See "Configuring advanced settings" on page 78.

Table 2-1 Entities having configurable scan schedules *(continued)*

Entity	Scan schedule settings location	Scope	Details
Filers, web applications, and cloud sources	In case of a filer, **Settings** > **Filers** > **Add New Filer**. In case of a SharePoint web application, **Settings** > **SharePoint Web Applications** > **Add SharePoint Web Application**. In case of a cloud storage account, **Settings** > **Cloud Sources** > **Add New Cloud Source**. **Note:** You can also configure scanning at the time of editing filers, web applications, and cloud sources.	Applies to filers, SharePoint web applications, or cloud sources for which schedule is defined. This setting overrides the scan schedule defined for the Collector associated with the filer, web applications, and cloud sources.	See "Adding filers" on page 210. See "Adding web applications" on page 256. See "Configuring Box monitoring in Data Insight" on page 272.
Shares and site collections	**Settings** > **Filers** > **Monitored Shares** > **Add New Share**. **Settings** > **SharePoint Web Applications** > **Monitored Site Collections** > **Add Site Collection**. **Note:** You can also configure scanning at the time of editing shares and site collections.	Applies to the entire share or site collection for which schedule is defined. Overrides the scan schedules defined for the filer or the web application associated with the share or the site collection.	See "Adding shares" on page 239. See "Adding site collections" on page 261.

You can override all the full scan schedules and initiate an on-demand full scan for configured shares or site collections. See "Managing shares" on page 241.

Sometimes for maintenance and diagnostic purposes, you may need to disable all the scans. You can disable all scans:

- At the time of adding or editing a storage device.
  See "Adding filers" on page 210.
  See "Adding web applications" on page 256.
- Or from the **Settings** > **Scanning and Event Monitoring** page of the Management Console.
  See "Configuring scanning and event monitoring " on page 28.

If you disable scanning for any device, you will not be able to view any permissions data for that device. However, you may still see some stale metadata like size,

permissions etc., which was collected before the scanning was disabled. If you run a report on the paths for which scanning is disabled, you may get a report with stale data.

You can specify pause schedules for both full and incremental scans to indicate when scanning should not be allowed to run. You can configure a pause schedule from the **Settings** > **Data Insight Servers** > **Advanced Settings** page. See "Configuring advanced settings" on page 78. to know more about configuring a pause schedule.

You can view the details of the current and historical scan status for your entire environment from the scanning dashboard. To access the scanning dashboard, from the Data Insight Management Console, navigate to **Settings** > **Scanning** > **Overview**. See "Viewing the scanning overview" on page 283. to know more about the scanning dashboard.

Configuring scanning and event monitoring

You can configure Data Insight to globally turn on or off receipt of event notifications and safeguards related to FPolicy communication.

Data Insight allows you to disable scanning of all file systems. The scans can be enabled at any convenient time.

You can also configure whether you want the Scanner to fetch the Access Control Lists (ACLs) defined on folders and ownership information for files and folders. For Windows File Servers, you can configure safeguards that prevent an agent from taking up too much disk space.

To configure scanning and event monitoring

1. In the Management Console, click **Settings** > **Scanning and Event Monitoring**.

   You can view the state of scanning and event monitoring.

2. To change the state of a process to Enabled or Disabled, click **Edit**.
3. Do one of the following:
   - Select the check box for the process that you want to enable.
   - Clear the check box for the process that you want to disable.
4. Click **Save** to save the changes.

Table 2-2 Scanning and Event Monitoring options

Option	Description
Scan File System meta-data	Clear the check box to turn off all future file system scanning on all filers. Once you save the setting, it will also stop all currently running scans.
Get Folder ACLs	Clear the check box if you do not want Scanner to fetch Access Control List (ACLs) for folders during scanning. If you disable this option, the **Workspace** > **Permissions** tab in the Console is disabled and permission related reports will not produce any data. If you do not need permissions data, you can disable this option to make the scans run faster.
Get Ownership information for files and folders	Clear the check box if you do not want Scanner to fetch the Owner attribute for files and folders. Ownership information is used to determine ownership for data when access events are not available. If you do not need this information, you can disable this option to make scans run faster.
Throttling for NetApp filers	Select **Throttle scanning based on latency of the filer** to enable throttling of Data Insight scans for NetApp 7-mode and Cluster-Mode file servers. This option is not selected by default. Data Insight collects latency information from NetApp file servers. It can use this information to throttle scanning, if latency of the file server increases above a certain level. This ensures scanner does not put additional load on the file server during peak load conditions. You can configure the following parameters to enable throttling for NetApp file servers: ■ Latency threshold - Specify latency in milliseconds, which when crossed, should throttle scanning for the file server. ■ Minimum pause - Specify the minimum duration (in milliseconds) for which the scanner should pause between paths when in throttling mode. ■ Back off value - If increased latency is sustained, pause interval will be increased by the Back off value specified (in milliseconds). ■ Maximum pause - Specify the maximum pause interval for the scanner (in milliseconds). If exceeded, pause interval is no longer incremented by Back off value.
Monitor file system access events	Clear the check box to stop Data Insight from collecting access events from all file servers. In case of NetApp, it means all collector nodes will disconnect their Fpolicy connections to file servers.

Table 2-2 Scanning and Event Monitoring options *(continued)*

Option	Description
Disk safeguard settings for Windows File Server agents	Select **Enable node safeguard** check box to monitor the disk usage on the Windows File Server node, and to prevent it from running out of disk space by implementing safeguards. You can specify the threshold for disk utilization in terms of percentage and size. The DataInsightWatchdog service initiates the safeguard mode for the Windows File Server node if the free disk space falls under the configured thresholds. The DataInsightWatchdog service automatically resets the safeguard mode when the free disk space is more than the configured thresholds. You can edit the threshold limits as required. If you specify values in terms of both percentage and size, then the condition that is fulfilled first is applied to initiate the safeguard mode. **Note:** It is recommended that you configure the same safeguard thresholds for the Windows File Server and the Classification Server (if present in your environment). For more information on the DataInsightWatchdog service, see the *Veritas Data Insight Installation Guide*.

Table 2-2 Scanning and Event Monitoring options *(continued)*

Option	Description
Fpolicy Safeguard settings	Select the following: ■ **Enable FPolicy 7-Mode Safeguard** to initiate safeguard mode for the NetApp file servers. ■ Select **Enable FPolicy 7-Mode safeguard for VFilers** to initiate safeguard mode for virtual file servers by monitoring latency statistics for the associated physical NetApp filers. The safeguard mode for the virtual filers is initiated only if the details of the physical filer corresponding to the virtual filer are provided while adding the virtual filer. ■ Select **Enable FPolicy Cluster Mode Safeguard** to initiate safeguard mode for the NetApp cluster mode file servers. This Fpolicy Safeguard settings are not selected by default. Data Insight collects latency information from NetApp file servers. It can use this information to initiate safeguard mode, if latency of the file server increases above or falls below a certain level. When the safeguard is in effect, Data Insight drops its Fpolicy connection to the filer. This ensures event collection does not put additional load on the file server in peak load conditions. If the latency (CIFS or NFS) on the physical file server increases above the configured threshold, Data Insight disconnects from the associated virtual file server even if monitoring for that type (CIFS/NFS) was not enabled on the virtual filer. This information is also displayed on the Data Insight **System Overview** dashboard. Configure the following values: ■ The high and low thresholds for CIFS and NFS. ■ The number of samples to be considered to calculate the average latency. ■ The minimum time to wait before Data Insight reconnects to the filer after a disconnection. See “Add/Edit NetApp filer options” on page 211.

You can choose to override the global scan throttling and FPolicy safeguard settings for specific Collector nodes.

See “Configuring advanced settings” on page 78.

Considerations for running a parallel scan

You can configure a collector to scan multiple shares of filers in parallel. By default, each scan process uses a single thread to perform the scanning. For certain large

shares, scanning with a single thread can take a long time to complete. For such shares, Data Insight lets you use multiple parallel scanning threads within a scan process to enable faster scans.

To configure parallel scanning for a collector or a filer from the **Advanced Settings** page:

See "Configuring advanced settings" on page 78.

To configure multiple threads to scan a single share:

See "Add New Share/Edit Share options " on page 240.

Enabling multiple threads for a scan process has certain performance implications for the collector and the filer.

The following considerations apply when you enable parallel scanning for one or more shares on a collector:

- It is recommended that the collector meets the following system requirements for running multi-threaded scans:
  - 32 cores
  - 64GB RAM
  - Windows Server 2012. The operating system must be 64-bit.
- As a precautionary measure to avoid excessive load on the filers, Data Insight runs a maximum of four multi-threaded scans in parallel on a collector regardless of the value configured on the **Advanced Settings** page.
- Based on your scanning requirements and available system resources, you can set the number of scan threads that you want Data Insight to use in parallel within a share from the **Add New Share/Edit Share** page. Data Insight interprets the scan thread count as follows:
  - Automatic: Data Insight computes the number of parallel scan threads based on the number of processors and the amount of workload. The maximum number of threads that Data Insight runs is between 1-16.
  - 1-128: Data Insight runs the specified number of scan threads.

  Note: For optimal parallel scan performance, it is recommended to select the scan option as **Automatic**.

- When a parallel scan is configured for certain shares, those shares are scanned faster. However, higher number of parallel scan threads may introduce a resource overhead on the filers. Thus, ensure that you do not set an excessively high value.

- It is recommended that the `DataInsightComm` service should be stopped only after the parallel scans are completed. Or, on the next restart of the service, the parallel scan automatically starts a fresh scan.
- If you edit a share to use parallel scanner instead of the single-threaded scanner while scanning is in the running or paused state, the single-threaded scan continues to run till it completes. For the subsequent full scans, the parallel scanner is executed. If you need to start using the parallel scanner immediately, then you must cancel the current scan from the **Settings** > **Scan Status**> **In-Progress Scan** page and restart the scan.

Limitations

The following limitations apply for the parallel scanner.

- Support for incremental scan is not available. Only full scans are supported.
- Support for NFS shares is not available.
- Support for filtering out shares based on the **Exclude Rules** configuration is not available.
- Support for throttling of parallel scans for NetApp 7-mode and Cluster-Mode file servers is not available.
- The **Scan History** sub-tab on the **Scanning** dashboard does not display the historical details of a parallel scan.
- For Windows File Server agents version older than 5.2, the parallel scanner cannot be executed. Even if it is configured, the single thread scan runs.

About filtering certain accounts, IP addresses, and paths

You can configure Veritas Data Insight to filter data accesses by specific users, IP addresses, file system paths, and URLs. You can combine these criteria together or use them individually to create a filter.

You can create separate exclude rules for file servers and SharePoint servers. For each of these, Data Insight supports two types of filters:

- Exclude rules for access events
- Exclude rules for Scanner

About exclude rules for access events

You can configure the following types of exclude rules for access events:

Filters for account names or SIDs	Typically, used to mask service accounts from registering data accesses into Veritas Data Insight. For example, if an antivirus software performs scans on a mounted share using a specific user account, you can add that user account to a filter. Data Insight omits all accesses made by that service user account.
Filters for IP addresses	Used to filter data accesses from specific IP addresses. Such filters are useful if you have file system scanners configured on certain machines in your environment, whose accesses you want to ignore. **Note:** IP addresses for events are not available for Windows File Servers and SharePoint.
Filters for path names	Filters for path names are of two types, file extension based and path based. The file extension based filter specifies the file extensions to be filtered. The path based filter specifies the path of a folder and filters out all events which have that path prefix. For path-based filtering, you must specify a fully qualified path prefix or a path relative to the root of each share. You can also use the wildcard (*) with a prefix string to exclude paths that have the string in the path name.
Filter for URLs	Used to filter data accesses from specified web applications or from SharePoint sites. Use fo wildcard (*) in exclude rules for access events is not supported.

About Exclude rules for Scanner

Scanner supports filtering out a top-level folder for all shares. You can define rules to exclude the scanning of the specified share, or SharePoint URL by the Scanner process.

Scanner does not support excluding folders under a top-level folder. Also, use of wildcard (*) in exclude rules for access events is not supported.

Adding exclude rules to Data Insight

You must create a rule for every filter you want to add to Data Insight. The rule must contain a value for at least one criterion that you want to exclude.

To add an exclude rule

1. In the Console, click **Settings** > **Exclude Rules**.
2. Click **Add Exclude Rule for File System** or **Add Exclude Rule for SharePoint**, as the case may be.

   From the drop-down, select **Exclude access** or **Exclude scanning**.
3. On the Add Exclude Rule screen, enter the Exclude rule properties.
4. Click **Save**.

Add/Edit Exclude rule for access events options

Use this dialog box to add a new exclude rule for access events to Data Insight or to edit the an existing exclude rule.

Table 2-3 Add/Edit file system Exclude rule for access events options

Field	Description
Rule name	Enter a logical name for the Exclude rule.
Usernames/SIDs	Enter the username or SIDs that you want to exclude. **Note:** The usernames must be present in the Data Insight users database, before they can be added to a exclude rule.
IP Addresses	Enter the IP addresses that you want to exclude. This filter only applies to NetApp and EMC Celerra file servers.

Table 2-3 Add/Edit file system Exclude rule for access events options *(continued)*

Field	Description
Exclude patterns	When defining a file system rule, enter the file extensions or paths that you want to exclude. A CIFS file system path must be fully qualified path in the format, `\\filer\share\folder` or relative to each share, for example, `<name of folder>`. A NFS path must be a fully qualified physical path on the actual file system in the format, `/path/in/the/physical/filesystem`. The logical operator OR is used create a rule with multiple values of the same dimension and the logical operator AND is used to combine values across dimensions in a rule. For example, if you create a rule to ignore user_foo1, user_foo2, and IP_10.209.10.20, it means that all accesses from IP_10.209.10.20 AND (user_foo1 OR user_foo2) will be ignored. You can also specify the wildcard (*) in an exlude rule for paths. Data Insight allows the use of wildcard (*) in the following formats in an exclude rule: ■ <prefix string> - Events for paths that start with the specified <prefix string> are excluded . ■ <prefix string>* - Events on paths that start with the specified <prefix string> are excluded. ■ *<string> - Events on paths which start with anything followed by the specified string are excluded. ■ *<string>* Events on paths which have the specified string somewhere in the path name are excluded. For example, if you specify *`<abc>`*, events on all paths that have the string `abc` anywhere in the path name will be excluded. When defining a SharePoint rule, enter the URL of the SharePoint web application or the site. You can use the wildcard (*) to exclude events for URLs that contain a specified sting in its name. For example, if you specify <abc>*, events on all URLs that have the string `abc` anywhere in the path name are excluded.
Pattern Type	Select PREFIX or EXTENSION from the **Pattern Type** drop-down. This field is only available for a file system rule.
Rule is enabled	Select the **Yes** radio button to enable the rule and the **No** radio button to disable it.

Add/Edit Exclude rule for Scanner options

Use this dialog box to add a new exclude rule for access events to Data Insight or to edit the an existing exclude rule.

Table 2-4 Add/Edit file system Exclude rule for Scanner options

Field	Description
Rule name	Enter a logical name for the Exclude rule.
Exclude Patterns	When defining a CIFS file system rule, specify the name of the folder to exclude as `/<name of first level folder>`. For NFS file system rule, specify the name of the folder to exclude as `/<name of first level folder>` When defining a SharePoint rule, enter the URL of the SharePoint web application or the site.
Rule is enabled	Select the **Yes** radio button to enable the rule and the **No** radio button to disable it.

About archiving data

Data Insight stores system events, alerts, and access events on the Indexer worker node in a pre-determined folder. You can configure Data Insight to automatically archive access events older that the specified interval to another folder to save space. Once the data is archived, it is no longer available for querying. You can, however, restore the data back to the original location on the Indexer node, if needed.

By default, archived data is automatically moved to `$data/indexer/archive` folder on each Indexer worker node. You can also configure a different archive folder on the Indexer nodes. The archive folder is organized by YEAR/MONTH to make restoring easy. Once data is moved to this folder based on the configured archive policy, you can do one of the following:

- Backup the archive folder and delete the archived files from the Indexer node.
- Or, configure a file system archiving solution like Enterprise Vault File System Archiving to archive all files in the archive folder.

If you want to restore archived data at a later time, you must bring back the appropriate segments from the backup folder to their original location in the archive folder and use the indexcli utility to restore these segments. Once restored, segments are not archived or purged by the data retention policy till they are marked for re-archiving.

Note: You can disable archiving of access events information by enforcing a legal hold on shares and site collections.

See "About purging data" on page 38.

See "Configuring data retention settings" on page 38.

About purging data

If you want Data Insight to automatically delete data, such as access events, system events, report outputs, and alerts older than specified interval, you can configure a purging policy. The purging policy must take into account your organization's requirements to retain data for legal compliance reasons.

Use the `indexcli` utility if you want to purge data at a more granular level than what you can configure on the **Data Retention** page on the Management Console. Purged data cannot be restored back at a later time.

Note: You can disable purging of access events information by enforcing a legal hold on shares and site collections.

See "About archiving data" on page 37.

See "Configuring data retention settings" on page 38.

Configuring data retention settings

Data Insight enforces the data retention policy twice a month. Archived index segments can be restored using a command line `indexcli.exe` utility. The utility is also useful to enforce a more granular archiving or purging policy, if the global option is not sufficient for your needs.

See indexcli.exe on page 353.

You can configure the duration for which you want Data Insight to retain various types of data and the duration after which you want to purge data. Automatic archiving and purging of data is not enabled by default.

To configure the data retention period

1. Click **Settings** > **Data Retention**.
2. On the **Data Retention** page, click **Edit**.
3. Enter the following information:

Purge workflows data automatically

Use the **Purge workflows data automatically** option to delete workflow records that are older than the configured age for purging to comply with your organization's data retention policy.

Do the following:

1. Select the check box to enable automatic purging of historical workflow records from Data Insight.
2. Enter the age of the workflow data that you want to retain so that the workflow data older than that will be purged (deleted).

Note: This option is disabled by default to ensure that workflow data is not deleted. When you enable it, the data for workflows which are in Completed, Canceled, Error, Failure, or Failed state is deleted.

Purge report outputs automatically

Use the setting to delete old copies of reports to free up storage space on the Management Server. Unless you configure this global setting or the per report setting to preserve only a certain number of reports, Data Insight keeps unlimited copies of report outputs in its database.

Select one or both options:

- **Purge older report outputs if count exceeds** - Enter a value that denotes the number of report outputs that you want to retain.

  Note: The number of reports to be preserved is also specified when configuring reports. The local setting for a report specified in the **Maximum reports to preserve** field takes precedence over the global count-based setting configured here. If you have not configured any local setting then Data Insight uses the global setting.

- **Purge report data older than <*number of months*>** - Enter the duration (in months). You can choose to purge reports that are older than the configured time period.

Note: If both, the local and the global settings are configured, then the global time-based option takes precedence for purging, and then the local count-based setting is considered.

Archive access data automatically	Do the following: **1** Select the check box to enable archiving of file system or SharePoint events. **2** Enter the age of the data (in months) after which the data must be archived. **3** Enter the path of the archive folder. In case, you want to archive access data immediately, execute the following command on the Indexer node: `INSTALL_DIR/bin/configcli execute_job DataRetentionJob`
Purge access data automatically	Select the check box to enable purging of file system or SharePoint access events, and enter the age of the data (in months) after which the data must be deleted.
Purge Data Insight system events automatically	Select the check box to enable purging of Data Insight system events, and enter the age of the data (in months) after which the data must be deleted. Data Insight system events are displayed on the **Settings** > **Events** page.
Purge alerts automatically	Select the check box to enable purging of alerts, and enter the age of the alerts (in months) after which they must be deleted.
Purge classification requests	Select the check box to delete all classification requests completed before the specified number of months.
Automatically purge data for deleted shares or site collections	Select the check box to enable purging of data pertaining to deleted shares. This option is enabled by default. If you clear the check box, the index data for the share or site collection remains on the disk until the global archive or purge takes place. Thus, if you decide to decommission a filer, or site collection, the data continues to be available for reporting.

4 Click **Save**.

5 In case, you want to purge workflows or audit data immediately, execute the following command on Management Server:

```
INSTALL_DIR/bin/configcli execute_job DataRetentionJob
```

About Data Insight integration with Symantec Data Loss Prevention (DLP)

Symantec Data Loss Prevention (DLP) helps you detect potential breach of sensitive data from your environment. The DLP Network Discover Server scans files on your NAS devices and generates incidents with details of the files that violate DLP policies.

Figure 2-2 Data Insight- DLP integration

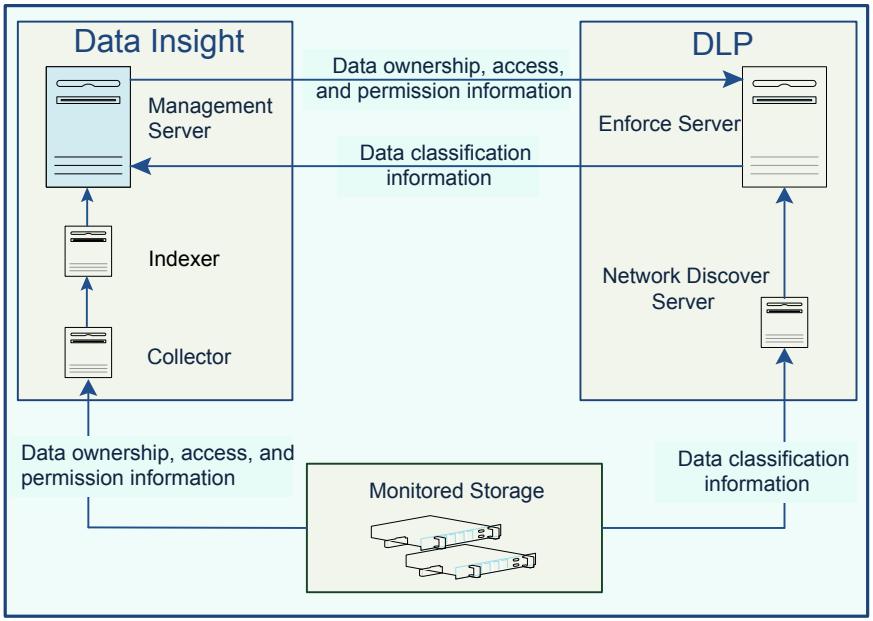

Data Insight has a bi-directional integration with DLP. Based on your requirement you can integrate the two products in any or both of the following ways:

- Configure DLP in Data Insight:
  DLP provides Data Insight the information about sensitive files in a storage environment monitored by DLP. Data Insight uses this information to raise alerts

in response to configured DLP policies. Data Insight runs the DLPSensitiveFilesJob at 12:00 A.M. every night to retrieve a list of sensitive files from DLP.

The information about sensitive files and DLP policies is used to calculate the risk-scores for storage resources and their users. The risk-scores and related information are displayed on the dashboard view of the Data Insight Management Console. You can use this information to find the high-risk shares and the folders that violate important DLP policies. Additionally, you can use the information from DLP to define DLP Incident Remediation workflow to take action on the files that violate certain DLP policies.

See "About configuring Data Insight to integrate with Data Loss Prevention (DLP)" on page 42.

- Configure Data Insight in DLP:

  Data Insight provides DLP ownership, access and permission-related information. DLP uses this information to generate folder risk reports and to display ownership information against incidents.

  See "About Symantec Data Loss Prevention (DLP) integration with Data Insight" on page 46.

For a visual representation of the integration refer Figure 2-2

About configuring Data Insight to integrate with Data Loss Prevention (DLP)

To configure Data Insight to fetch sensitive files information from DLP, complete the following tasks:

- Configure DLP settings from the Data Insight Management Console.

  See "Configuring Symantec Data Loss Prevention settings" on page 44.

- Import the SSL certificate from the DLP Enforce Server to Data Insight.

  See "Importing SSL certificate from the DLP Enforce Server to Data Insight Management Server" on page 45.

To view the complete workflow for setting up Data Insight to integrate with DLP, refer to Figure 2-3

Figure 2-3 Data Insight - DLP integration workflow

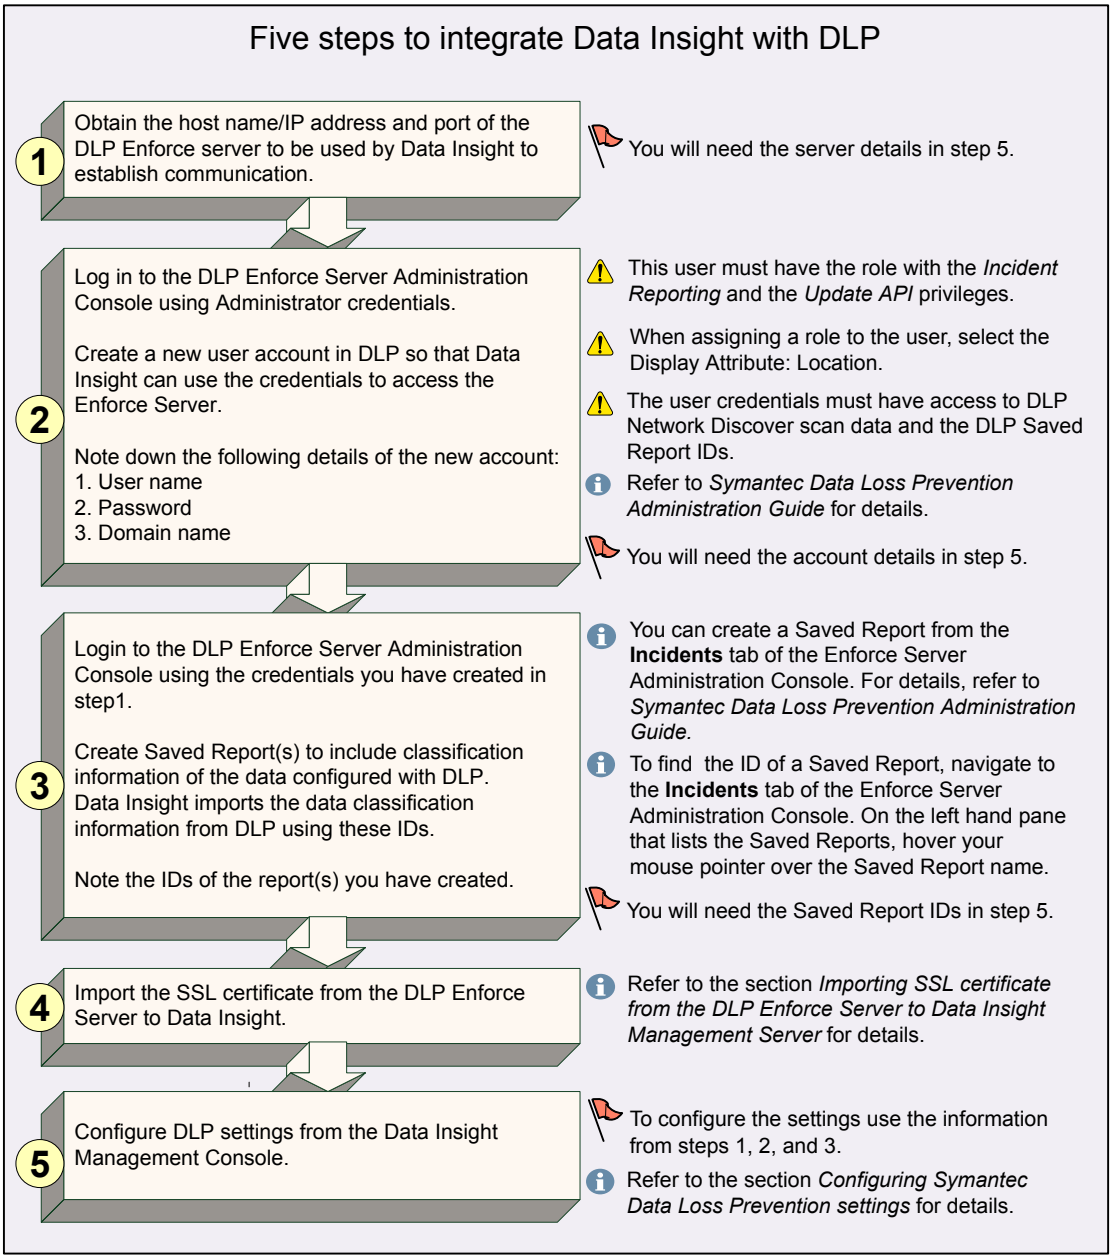

Configuring Symantec Data Loss Prevention settings

You must configure the settings that allow Data Insight to communicate with Symantec Data Loss Prevention.

To configure Data Loss Prevention settings

1. In the Management Console, click **Settings** > **Data Loss Prevention**.
2. Click **Edit**, and enter the following details:
   - **Hostname /IP address of DLP Server**
     The host name or IP address of the DLP Enforce Server. DLP host names are case-sensitive.
   - **Port**
     The port through which Data Insight connects to the DLP Enforce Server. The default port number is 443.
   - **Username**
     The user name of the account that is used to access the DLP Enforce Server. DLP user names are case-sensitive.

     Note: Ensure that the credentials belong to an existing DLP user assigned the Incident Reporting and Update API role. Also ensure that when assigning a role to the user, the Display Attribute **Location** is selected. This attribute allows Data Insight to view the complete path of a file.

     The user credential being used must have access to DLP Network Discover scan data and DLP Saved Report IDs.
   - **Password**
     The password of the account that is used to access the DLP Enforce Server.
   - **Domain**
     The name of the domain to which the user belongs. DLP domains are case-sensitive. Specifying the domain is optional for a user who is a DLP administrator.
   - **DLP Role**
     Specify the role you want to use to log on to DLP. DLP roles are case-sensitive.
     Users who are assigned more than one role can only log on under one role at a time.
   - **Configure storage resources automatically**
     By default, Data Insight discards classification information for paths on storage devices that it does not monitor. Select this option to add the

unmonitored devices as generic devices in Data Insight and include classification information for all paths on those devices. Note that only data classification information is available for the paths. To enable Data Insight to get audit and metadata information for the paths, you must manually configure the devices in Data Insight.

- **Saved Report IDs**
  Enter a comma-separated list of Saved Report IDs in DLP.
  Data Insight fetches classification information for all the paths present in the DLP Saved Report. Data Insight uses the DLP Reporting API Web service to request a list of incident IDs by specifying a saved report ID. A Data Insight process then fetches the sensitive files corresponding to the incident IDs.

3. Click **Test Connection** to verify the connection to the DLP Enforce Server.
4. Click **Save** to save the settings.

See "Importing SSL certificate from the DLP Enforce Server to Data Insight Management Server" on page 45.

Importing SSL certificate from the DLP Enforce Server to Data Insight Management Server

The DLP Enforce Server administration console requires SSL transport for all communication. Data Insight must be able to negotiate the SSL connection with the Enforce Server. For this purpose, you must import the certificate to the keystore used by Data Insight.

To import the SSL certificate from the DLP Enforce Server to Data Insight using Firefox

1. Type the URL to connect to a DLP Enforce Server Administration console.
2. On the security certificate warning page, click **I understand the risks**.
3. Click **Add Exception**.
4. On the Add Security Exception page, click **View** to view the certificate details.
5. Click the **Details** tab and click **Export**.
6. From the Save as type drop-down, select X.509 Certificate (DER).
7. Click **Save**.

To import the SSL certificate from the DLP Enforce Server to Data Insight using Internet Explorer

1. Type the URL to connect to a DLP Enforce Server Administration console.
2. On the security certificate warning page, click **Certificate Error** next to address bar.
3. Select **View certificates**.
4. Click the **Details** tab, and select the appropriate certificate.
5. Click **Copy to File**
6. In the Certificate Export Wizard, select **DER encoded binary**.
7. Click **Next**.
8. Enter the name of the file and browse to the location where you want to save the file.
9. Click **Next**
10. Click **Finish** to save the file.

After the SSL certificate is imported, complete the following steps to import the SSL certificate on the Data Insight server.

To import the SSL certificate on the Data Insight server

1. From the Windows Start menu, select **Run** and type `cmd` in the dialog box to open a command prompt window.
2. Run the following command:

   ```
   cd C:\Program Files\DataInsight\jre\bin
   ```

   ```
   .\keytool -importcert -alias dlp -keystore c:\
   DataInsight\data\keys\commd.keystore -trustcacerts -file <file
   path of SSL certificate>
   ```

   Specify `changeit` as the password for the keystore.

   You can now pull a list of sensitive files from Symantec Data Loss Prevention (DLP).

See "Configuring Symantec Data Loss Prevention settings" on page 44.

About Symantec Data Loss Prevention (DLP) integration with Data Insight

Complete the following tasks in DLP to configure DLP to pull access data from Data Insight:

- Configure a connection between the DLP Enforce Server and Data Insight.
- Configure the Data Insight Lookup Plug-in to retrieve data owner information.
- Configure other lookup plug-ins to populate the Data Owner email field in Data Insight.
  Refer to the *Symantec™ Data Loss Prevention Administration Guide* for details on configuring these plug-ins.
- On the Enforce Server, create custom attributes for each file detail that you want retrieved from the Data Insight Management Server.
- Map the custom attributes that you have created to the details from the Data Insight Management Server.
- Restart the DLP Enforce services.

The steps mentioned in this section are applicable for DLP users who want to pull data ownership, permissions and access information from Data Insight. For the detailed steps, see the *Data Loss Prevention Data Insight Implementation Guide*.

See "About Data Insight integration with Symantec Data Loss Prevention (DLP)" on page 41.

Importing sensitive files information through CSV

Data Insight pulls information about sensitive files in your storage environment from Symantec™ Data Loss Prevention (DLP).

However, if you use a third-party application other than DLP, Data Insight provides you the ability to identify sensitive files in your storage environment using a .csv file. A scheduled Data Insight process reads the .csv file to retrieve the list of sensitive files from Data Insight.

To use a .csv file to classify sensitive files in Data Insight

1. Log in to the Data Insight Management Server.
2. Create a .csv file, in which each line indicates the path of the sensitive file, and policy names which that particular file violates.

   For example, you have a `file /foo/bar/info.txt` which violates the policies *Personal Information* and *Hipaa*. And another file `/foo/ssn.pdf` which violates the policy *Personal Information* In this case, create a .csv file as follows:

   ```
   data /foo/bar/info.txt,Personal Information, Hipaa /foo/ssn.pdf,
   Personal Information
   ```

3. Edit the dlp_db.conf file to add the following lines:

   ```
   dlp.csv.enabled=true
   ```

   ```
   external.file.path=<full path to the location of
   .csv file on Management Server>
   ```

   The dlp_db.conf file is located in the <datadir>\conf directory.

Note: If the `dlp.csv.enabled` property is set to true in the `dlp_db.conf` file, the Data Insight process uses the .csv file to identify sensitive files, even if DLP is configured in Data Insight.

Configuring advanced analytics

Data Insight provides several advanced data analytics features in the form of Dashboard reports and Social Network Map. These features let you make intelligent decisions about your data.

Use the **Advanced Analytics Configuration** page to configure the following criteria:

- The period for which the activity information should be considered to calculate the device and share statistics. This period is also considered for analyzing collaborative activity on a share when displaying Social Network Map, and for generating data displayed on the **Workspace** views.
- The depth of the folder hierarchy with respect to the root of the share that must be evaluated to compute the control points within a share. The default folder depth for computing control points within a share is 5. This means that by default Data Insight evaluates the folder hierarchy 5 levels deep to calculate the control points within a share.

Data Insight displays the information about control points on the **Workspace** > **Data Sources** view of the Management Console.

- The interval for refreshing the data that is displayed **Workspace** > **Dashboard** view of the Management Console.
  You must ensure that you decide the frequency of refreshing the data judiciously, because the statistics are calculated for all the configured devices, shares, and site collections.

Note: The Data Insight dashboard does not display any data, if a summary report has not run even once.

The page also displays information about refresh cycles that have failed. Click on the **Failed** link to download the logs that you can use to troubleshoot the cause for failure.

To configure advanced analytics

1. In the Management Console, click **Settings** > **Advanced Analytics**. The existing analytics settings display by default.
2. Click **Edit** to change the appropriate settings.
3. Click **Save** to save the changes.
4. Click **Compute Now** to refresh the data on the Data Insight dashboard.

See "About open shares" on page 50.

Choosing custom attributes for advanced analytics

Data Insight lets you configure the custom attributes that you can use for the purpose of advanced analytics, reporting, and remediation.

See "Configuring attributes for advanced analytics" on page 116.

The configured attributes are available for selection on the Advanced Analytics Manager. Use this page to mark the attributes that you want to consider for analytics and for grouping of users based on the primary attribute.

The primary attribute is used to best identify users in the Social Network Map cluster groups. It is also used to distinctly categorize the users in the cluster groups. Ensure that the attribute that you choose as the primary attribute is the one that is populated for the majority of the users in your directory service. For example, good choices for primary attribute would be department or cost code. Attributes such as, employee ID or title are different for each user, and should ideally not be used as the primary attribute.

To choose attributes for advanced analytics

1. On the **Settings** tab, click **Advanced Analytics**.
2. Click **Attributes** to display the **Advanced Analytics Manager**.

   The **Available Attributes** panel displays all the configured custom attributes.
3. Select an attribute, and click the right arrow to select an attribute. Similarly, you can click the left arrow to remove an attribute from the list of selected attributes.
4. Use the up and down arrows to set the priority of the attributes for computing the analytics data.
5. From the **Primary grouping attribute** drop-down, select the attribute that you want to use as the primary attribute for identifying users and for creating attribute-based filters.

About open shares

A share is considered to be open due to the permissions assigned on it. You can define the parameters that determine whether a share should be considered as open. However, the open share policy does not allow you to define specific permissions that render a share open.

One of the following parameters determine whether a share is open:

- If certain group has access to a share, either directly or as a nested group.
- If more than a certain number of users have access to the share.

Additionally, you can also specify granular criteria for examination, such as:

- The level at which Data Insight should start examining the paths.
- The number of levels deep to examine the path permissions.

For example, you can define the following criteria to consider a share to be open:

- List of groups: Domain Users or Everyone.
- No of users who have access to the share: 500.
- Level to start examining permissions: 1.
- Depth to examine: 3.

According to the above specification, any share that is accessed by the Domain Users/Everyone group or by more than 500 users is considered to be open. For this purpose, the ACLs are examined from level 1 (root being level 0), and all folders three levels down are examined.

Defining an open share policy helps you to review the permissions assigned at share level and revisit entitlement decisions. You can view the details of the open shares, the size of the data in open shares, and the number of files in open shares on the **Dashboard** tab on the Management Console.

See "Configuring an open share policy" on page 51.

Configuring an open share policy

You can define the parameters for an open share on the **Open Share Policy** page of the Management Console. The paths on a storage device that meet the criteria specified in the policy are considered as open.

See "About open shares" on page 50.

To configure an open share policy

1. Click **Settings** > **Advanced Analytics**. The **Configuration** page opens by default.
2. Click the **Open Share Policy** sub-tab.

   The page displays the existing policy that defines the parameters that determine whether a share is open.
3. To change the existing policy, click **Edit**.
4. Enter the criteria for considering a share to be open:
   - Enter the threshold for users that have access to the share. Or select the user(s) or group(s) that have access to the share. You can choose to select both conditions simultaneously.
   - You can choose to specify the users and groups that have permissions on the share. Use the users and groups selection widget or upload a .csv file with the users and groups information. The selected entities are listed in the Selected Users/Groups pane.

     You can type a name in the search bar to search for a user or group. You can also type a domain name in the **Domain Filter** field to narrow your search to users in a specific domain. You can also filter a user or group from the **Select Filter** drop-down. Select the **All Filtered Users** check box in the Selected Users/Group pane to include all filtered users in the policy definition.

     All the configured users and groups are displayed on the open share policy definition page.
5. Use the Up and Down arrows to define the level in the share hierarchy that the policy should be applied. You can also examine the depth starting from level 0, that is the root of the share.

6. The depth in terms of number of levels of the folder heirarchy for which the permissions should be examined.
7. Click **Save** to save the policy and close the window.

You can use the report.exe utility to exclude certain paths from the open share computation for the dashboard.

See reportcli.exe on page 359.

Configuring file groups

By default, Data Insight sorts files into 18 file groups based on the extension of the files. The file group information is used for reporting on ownership, access pattern, and space consumption on storage devices. Data Insight uses file group information when generating the following reports:

- Inactive Data by File Group
- Consumption by File Group
- Consumption by File Group and Owner

You can modify the default file groups by adding custom extensions or by deleting extensions from the file groups. You can also create new file groups in **File Groups** view from the Management Console or use the `fg.exe` command from the command line interface. When a new extension is added, all indexes update their extension information. This process can take several hours depending on size of your indexes. Thus, new extensions will not show in your report output till all indexes have been updated.

You can issue fthe ollowing command on your Indexer nodes to monitor the status of updating the indexes:

```
indexcli -j
```

You can search for extensions and file groups by using the filter at the top right corner of the screen.

For detailed information on `fg.exe`, see the *Command File Reference*.

To configure file groups

1. In the Management Console, click **Settings** > **File Groups**.
2. To add a new file group, click **Add new file group**.
3. Enter a logical name for the file group.
4. In the Extension field, enter the extension that you want to include in the file group.

5. Click **Add Extension**.

   You can add multiple file extensions to a single file group.

6. To modify an existing file group, click the file group.
7. Click the delete icon corresponding to the extension you want to remove, or enter the extension that you want to include in the file group, and click **Add Extension**.
8. To delete an entire file group, click the corresponding delete icon.

Configuring Workspace data owner policy

By default, Data Insight infers owners of files or folders based on the access history. The most active user of a file is considered to be the data owner for the purpose of efficient remediation and data management. The Workspace data owner policy is used to infer the probable owner when configuring workflows and reports.

However, you can define a global policy to infer the owners of files or folders based on one of the following criteria:

- The number of read events on the file or folder.
- The number of write events on the file or folder.
- The cumulative count of all activities except for the permission change events on the file or folder.
- The creator of the file or folder.
- The user account which last accessed the file or folder.
- The user account which last modified the file or folder.
- Owner of the parent folder.
  If Data Insight is not able compute the owner on the basis of read, write, last accessed, last modified, or create events, the owner of the immediate parent folder is displayed as the owner of the file or folder.

For example, you can define a policy to consider the count of read and write events on a file to determine the data owner. In this case, the user with the most read and write accesses is considered to be the data owner.

The criteria that is defined in the global Workspace data owner policy are also considered for determining the data owner when you configure reports and workflows. However, you can choose to override the policy when you create an **Inferred Data Owner** report.

For information about configuring workflows and reports, see the *Veritas Data Insight User's Guide*.

To configure a Workspace data owner policy

1. In the Management Console, click **Settings** > **Workspace Data Owner Policy**.
2. The Data Owner Policy details page displays the current policy.
3. Use the up and down arrow keys to arrange the criteria in the preferred order.

   Data Insight serially evaluates the criteria in the list to compute the data owner.
4. Select the check boxes to exclude the following types of users when Data Insight calculates the data owner:

   - Disabled users - These are the users whose accounts are disabled in the directory service, but are still present in the Data Insight configuration.
   - Deleted users - These are the users whose accounts are removed from the directory service.
     Accounts of users are typically disabled or deleted in the directory service when a user leaves an organization.
   - Unresolved SIDs - These are accounts for which the corresponding user information is missing in the directory service.

   For example, you define a policy that considers the number of write events on a file to compute the data owner. If the user account with the most write events on a particular file is either deleted or disabled, Data Insight excludes such user from the data owner computation. It then considers the next criteria in the list to determine the owner of the file.
5. Select the groups or users that you want to exclude from the scope of the policy. Double-click the group or user to select it. The selected data set is listed in the **Exclusion List** pane.

   Use the Domain Filter drop-down to search for specific users, users in a particular domain, built-in user accounts, and SIDs that have been migrated to another domain.
6. Click **Save**.

Configuring Management Console settings

In the Console Settings view, you can configure global settings that apply to various tasks that you carry out on the Management Console.

To configure the Console settings

1 Click **Settings** > **Console Settings**.

2 Click **Edit**.

You can edit any of the following settings:

Session Timeout	Your login session on the Management Console times out after certain period of inactivity. The default timeout period is one hour. To configure the session timeout period, enter the time in minutes.
Report Footer Text	You can choose to add a footer to all the reports that you run in the Console. Enter the sentence string that you want to appear in the footer of the report. For example, Proprietary and Confidential.

3 Click **Save**.

For more information about creating reports, see the *Veritas Data Insight User's Guide*.

About bulk assignment of custodians

Data Insight uses the information about data custodians to distribute remediation requests and reports. The custodians are responsible for taking action on the data assigned to them. To know more about the Data Insight custodians, you can refer to the *Veritas Data Insight User's Guide*.

You can assign custodians to a data resource from the **Workspace** tab. However, assigning custodians on a path at a time can be tedious and time-consuming. You can use the **Custodian Manager** to easily assign multiple custodians using only a few steps.

You can bulk-assign custodians by using any of the two options:

- **Assign by CSV** – You can use a CSV file that contains the information about the paths and their respective custodians for to assign custodians on paths. See "Assigning custodians in bulk using a CSV file" on page 56.
- **Assign by owner method** – You can specify the criteria for computing the possible owner of the selected paths, and assign the computed owners as the custodians.

See "Assigning custodians based on data ownership" on page 56.

Assigning custodians in bulk using a CSV file

To assign custodians by using a CSV file

1. Prepare a `.csv` file containing the following details separated by a comma:
   - Paths for which a custodian is to be assigned.
     For example: `\\30.219.81.23\Managers\Reportees\`
   - Custodian to be assigned to that path.
     To specify a custodian, you can use the user name along with the domain name. For example: *Custodian_name@domain_name*. Alternatively, you can specify the SID of the user.
2. On the Data Insight Management Console, navigate to **Settings** > **Global Settings** > **Custodian Manager**.
3. From the drop-down menu, select **Assign by CSV**.
4. Click **Choose File**.
5. Browse to the location where you have saved the CSV file. Select the CSV file and click **Open**.
6. Click **Assign**.

Note: The custodian assignment in Data Insight can take some time depending on the number of paths. You can view the status of the operation on the **Settings** > **Events** page of the Management Console.

See "About bulk assignment of custodians" on page 55.

Assigning custodians based on data ownership

Data Insight computes the possible owners for any path based on parameters such as read counts, write counts, last modifiers, last accessors, creators etc. Based on your requirements, you can select any of these parameters and assign them priorities for ownership computation. For computing the owner, Data Insight by default considers the activity period configured on the **Advanced Analytics** page.

To assign custodians in bulk based on data ownership

1. From the Data Insight Management Console, navigate to **Settings** > **Custodian Manager**.
2. From the drop-down menu, select **Assign by owner method**. The pane expands to display the **Custodian Selection Criteria** panel.
3. Note that the **Use default data owner policy** check box is selected by default. Do any of the following:
   - Click the **Use default data owner policy** check box to let Data Insight use the data owner policy and exclusion rules as defined under the **Workspace Data Owner Policy** setting.
   - Clear the **Use default data owner policy** check box to define your own set of criteria to calculate the data owner. Use the right arrows and the left arrows to select or deselect a criterion. Use the up arrow and the down arrow to change the priority of the criteria.

     Note: When you clear the **Use default data owner policy** check box, Data Insight still enforces the exclusion rules for deleted, disabled, and unresolved users as defined under the **Workspace Data Owner Policy** setting.

   See "Configuring Workspace data owner policy" on page 53.
4. From the **Paths** panel, select the paths for which you want to assign custodians. The selected paths are displayed under the **Selected resources** tab.
5. Click **Assign** to complete the request.

Note: The custodian assignment in Data Insight can take some time depending on the number of paths. You can view the status of the operation from the **Settings**> **Events** page of the Management Console.

See "About bulk assignment of custodians" on page 55.

Configuring Watchlist settings

Data Insight lets you add users and groups to a watchlist. You can add such users or groups to the watchlist that you want to monitor closely for one of the following reasons:

- A long-serving employee who is leaving the organization.

- A user or group who has access to a number of content repositories that store sensitive data.
- An outlier that is suggested by the Social Network Map.

You can review the risk score assigned to a user and track the activities of the user or group that is added to a watchlist from the Data Insight dashboard.

For information about the Data Insight dashboard, see the *Veritas Data Insight User's Guide*.

To add a user or group to a watchlist

1. From the Management Console, navigate to **Settings** > **Watchlist Settings**. The **Configure Watchlist** page opens.
2. You can add the users or groups to a watchlist in one of the following ways:

   You can use the domain filter search bar to filter users or groups according to domains.

   - Under the **User Selection** tab, click either the **Users** or the **Groups** radio buttons to view the configured users or groups respectively. Click a user or a group to select it.

     Alternatively, you can upload the users and groups information through a CSV file. Click **Upload CSV**.
   - Under the **User Inclusion Using Attributes** tab, click **Add Condition**. From each of the drop-down menu, select the criteria to build the query. The query is used to select users based on their Active Directory attributes.

     You can add multiple conditions to select users. Data Insight evaluates a query as follows:

     - Data Insight uses the logical OR operation for same attribute condition with different values.
     - It uses the logical AND operation between multiple attribute conditions.
3. Similarly, you can exclude specific users from a watchlist in the following ways:
   - Under **Exclusion list** tab, select the users or groups that you want to exclude from the watchlist.

     Or click **Upload CSV** to upload a list of users or groups in bulk.
   - Under the **User Exclusion Using Attributes** tab, click **Add Condition**. From each of the drop-down menu, select the criteria to build the query. The query is used to exclude users based on their Active Directory custom attributes.
4. Click **Save**.

Section 2

Configuring Data Insight

- Chapter 3. Configuring Data Insight product users
- Chapter 4. Configuring Data Insight product servers
- Chapter 5. Configuring saved credentials
- Chapter 6. Configuring directory service domains
- Chapter 7. Configuring containers

Chapter 3

Configuring Data Insight product users

This chapter includes the following topics:

- About Data Insight users and roles
- Reviewing current users and privileges
- Adding a user
- Editing users
- Deleting users
- Configuring authorization for Symantec Data Loss Prevention users

About Data Insight users and roles

Before a user can log in to Veritas Data Insight, you must add an account for that user. The user can then use that account to log in to the Console. The user account can be any account that is valid on the Management Server system. This includes local system accounts as well as users belonging to the domain which the Management Server is a part of.

When you create an user account, a role (set of access privileges) is associated with the account. Roles specify access privileges to the Veritas Data Insight system. For example, a role might let users view access and permissions data, but prevent them from adding or deleting filers. Data Insight role-based access control governs access to product features and functionality. Roles consist of the user privileges that determine what a user can see and do in the Management Console.

The Data Insight administrator (a user mapped to the predefined Server Administrator role) assigns roles to users. Users can be mapped to one role only. Data Insight ships with predefined roles that you can assign to user accounts.

Table 3-1 summarizes the various Data Insight roles.

Table 3-1 Veritas Data Insight roles

Role name	Description
Server Administrator	Allows the user to perform all actions in the product GUI that includes setting up all infrastructure (including filers, users, and others) and view all the access and permissions data.
Product Administrator	Allows the users to manage filer settings and optionally to view all the access and permissions data for the given filers. Product administrator role, configured for a select set of filers/web applications, is not allowed to add new filers or delete configured filers.
Report Administrator	Allows the user to view and edit all reports that are configured in the system and all data on the **Workspace** tab. Additionally, the Report Administrator may or may not be configured to take post processing actions on a report output. A Report Administrator does not have access to the **Settings** tab. Only a user with the Server Administrator role can add a user with the Report Administrator role.
User	Allows the users to view all permissions data. However, user in this role may or may not be allowed to view activity data of other users and certain reports. See "Add or edit Data Insight user options " on page 62. Users assigned this role also do not have access to any tasks on the **Settings** tab.
Storage User	Allows the users to view storage-related data in the **Workspace** tab, but does not allow them to view permissions data or audit logs. Users in this role do not have access to the **Settings** tab.

Reviewing current users and privileges

You can review the current Data Insight users and the roles assigned to them on the Product Users list page. On this page you can also review the filers and web applications that these users are allowed to monitor.

To review current users and privileges

1 In the Console, double-click **Settings** > **Product Users**.

The following information is displayed:

- The name of the user.
- The role assigned to the user.
- The scope of the access available to the user.
- The time when the user last logged into Data Insight.

2 Click the Export icon at the bottom of the page to save the data to a `.csv` file.

Adding a user

This section describes how to add users to Veritas Data Insight.

To add new a Data Insight user

1 In the Console, click **Settings** > **Data Insight Users** to display the Product Users listing page.

2 Click **Add New Data Insight User**.

3 On the Configure new product user page, enter the user properties, click **Add New Data Insight User**.

Add or edit Data Insight user options

Use this dialog box to add a new user to Data Insight, or edit the properties of an existing user.

Table 3-2 Add/Edit Data Insight user options

Field	Description
User name	Enter the user name for the user.
Domain name	Enter the name of the domain to which the user belongs.
Role	From the drop-down, select the role you want to assign the user. See "About Data Insight users and roles" on page 60.

Table 3-2 Add/Edit Data Insight user options *(continued)*

Field	Description
Allow access to Workspace data	Select the check box to enable the user to view the screens on the **Workspace** and the **Reports** tabs This option is only available if the user is assigned the Product Administrator role.
Resources/Containers to grant access to	Select one of the following: ■ **All filers/Web applications (Includes the ones added in the future)** ■ **Selected Filers/Web applications** ■ **Selected Shares/Site Collections** ■ **Selected DFS paths** ■ **Containers** If you select **Selected filers/Web Applications** , **Selected Shares/Site Collections**, **Selected DFS paths**, or **Containers** , the system displays a list of the appropriate configured entity. Use the arrows to select the entities you want the user to monitor. **Note:** A user, assigned the Server Administrator role, has the scope set to **All Filers/Web Applications**, by default. The scope by DFS paths is applicable only for User and Storage User roles.

Editing users

After you add a user to Data Insight, you can edit the user properties. For example, you might need to edit any of the following:

- The role assigned to the user
- The view option for the user
- The filers and/or web applications that the user is allowed to monitor

To edit the properties of a user

1. In the Console, double-click **Settings** > **Data Insight Users** to display the Product Users listing page.
2. Click the **Edit** button for the corresponding user.
3. On the Edit Data Insight user page, make changes. as necessary, and click **Save**.

   See "About Data Insight users and roles" on page 60.

Deleting users

You can delete Data Insight users.

To delete an user

1. In the Console, double-click **Settings** > **Data Insight Users** to display the Product Users listing page.
2. Select the user, and click **Delete**.
3. Click **OK** on the confirmation message.

Configuring authorization for Symantec Data Loss Prevention users

Symantec Data Loss Prevention makes web service calls into Data Insight to obtain ownership information for sensitive files and folders. However, you must first provision a Data Insight account for Symantec Data Loss Prevention in Data Insight.

You can provision a Active Directory service account OR a local system account and assign it the Server Administrator privilege. Symantec Data Loss Prevention can use this account to access Data Insight data.

Chapter 4

Configuring Data Insight product servers

This chapter includes the following topics:

- About Data Insight product servers
- Adding a new Data Insight server
- Managing Data Insight product servers
- Viewing Data Insight server details
- About node templates
- Adding Portal role to a Data Insight server
- Adding Classification Server role to a Data Insight server
- Assigning Classification Server to a Collector
- Viewing in-progress scans
- Configuring Data Insight services
- Configuring advanced settings
- Monitoring Data Insight jobs
- Rotating the encryption keys
- Viewing Data Insight server statistics
- About automated alerts for patches and upgrades
- Deploying upgrades and patches remotely

- Using the Upload Manager utility
- About migrating storage devices across Indexers
- Viewing the status of a remote installation

About Data Insight product servers

A Data Insight product server is any server which has Data Insight software installed. This includes the Management Server, zero or more Collectors, zero or more Indexers, and zero or more Windows File Server agents.

For more information about the Data Insight servers, see the *Veritas Data Insight Installation Guide*.

You can view information about configured product servers, check the status of running scans, and change advanced settings from the **Settings** tab of the Management Console.

Adding a new Data Insight server

You can add a new Data Insight from the Management Console.

To add a Data Insight server

1. In the Console, click **Settings** > **Data Insight Servers** to display list of configured product servers.
2. Click **Add new server**.
3. On the **Add New Server** page, enter the following details:
   - The host name or IP address of the server.
   - Select a credential from the **Select Saved Credential** drop-down. The credential must belong to a user with Administrator privileges on the node where the software needs to be installed.
     See "Managing saved credentials " on page 105.
   - The location where you want to store the product data. Select a location with enough free space and high-performance disks. By default the data directory is located at **C:\DataInsight\data**.
   - The directory in which you want Data Insight to be installed. By default, the destination directory is `C:\Program Files\DataInsight`.
   - Depending on your deployment scenario, select **Collector** only or **Indexer and Collector**, **Self-Service Portal**, or **Classification Server** as the installation option.

- The **Communication Service Port**. By default, the Communication Service connects through service port 8383.
- The **Configuration Service Port**. By default, the Configuration Service connects through service port 8282.
- If you to install the Self-Service Portal, enter the port numbers for the Self-Service Portal service and the Workflow service. By default the port numbers are 443 for the **Portal Webserver Port** and 8686 for the **Workflow Service Port**.

4 Click **Install**.

You can view the progress and status of the installation on the Installation Status page.

See "Viewing the status of a remote installation" on page 104.

See "Configuring Data Insight services" on page 76.

Managing Data Insight product servers

On the Data Insight servers listing page you can do following tasks:

- View detailed information about all configured servers.
- Add a new Data Insight server.
  See "Adding a new Data Insight server" on page 66.
- Add the Portal role to the existing Management Server or an existing Collector node.
  See "Adding Portal role to a Data Insight server" on page 73.
- Adding the Classification role to the existing Management Server or Collector node.
  See "Adding Classification Server role to a Data Insight server" on page 74.
- Get a list of currently running scans.
  See "Viewing in-progress scans" on page 75.
- Edit the server's configuration.
- Apply a node template.
- View and install recommended patches for a server node.
  See "Viewing and installing recommended upgrades and patches" on page 100.
- Remotely upgrade and push-install rolling patches on Collector and Windows file server agent nodes.
  See "Deploying upgrades and patches remotely" on page 100.

- Run a specific Data Insight process.
  See "Monitoring Data Insight jobs" on page 96.
- Delete the server.

Use the filter on the left to filter the list of servers based on the health of the server or role of the server. You can also search for a server by entering the name or IP address of the server in the **Filter** text box. The displayed list of servers changes automatically as you enter the term in the **Filter** text box or when you select the check boxes. For example, you can search for all Healthy servers that are assigned the Collector role. Data Insight displays all the Collectors in your environment that are in the Healthy state.

The following parameters determine the health of a Data Insight server:

- Whether a node is online.
- Whether all the required services are running on a node.
- The disk space on the server.
- Whether the node has error files in the `err` folder.
- CPU and Memory usage on the server

The health of a server is not known for five minutes after starting the DataInsightWatchdog service, or if the DataInsightWatchdog service is stopped.

The Servers pie chart on the **System Overview** dashboard gives a high-level overview of the number of Data Insight severs in Faulted, at Risk, and Healthy state.

See "Viewing the system health overview" on page 280.

To view configured product servers

1. In the Console, click **Settings** > **Data Insight Servers** to display list of configured product servers.
2. Use the provided dynamic search filter to search for configured servers based on the name of the server.
3. Review the following information about the servers:
   - The ID of the server.
   - The name of the server.
   - The role of the server.
   - In case of a node with the role Management Server, Collector, or Collector and Indexer, the Classification Server assigned to the node.
   - The status of the server — whether the server is online or offline.
   - A list of matching templates for this server.

- The version of the Data Insight software that is installed on the server.
- The credentials used to install the server.
- The links to recommended patches and hot fixes.

Data Insight also displays the recommendation to upgrade to the latest version, if available.

To view server events

1. In the Console, click **Settings** > **Data Insight Servers**.
2. Click the **Select Action** drop-down for the corresponding server in the servers listing table, and select **Event Log**.

   Or, on the details page for a server, click **Event Log**.

   The event log for that server appears.

You can create a node template to change one or few settings on multiple nodes. Using node templates is useful when multiple nodes need to inherit the same settings, for example, more number of indexer threads for all indexers in your environment.

See "Managing node templates" on page 71.

To apply a template to a Data Insight node

1. In the Console, click **Settings** > **Data Insight Servers**.
2. Select the server to which you want to apply a template, and from the **Apply a node template** drop-down, select a configured template.
3. Click **Yes** to confirm.

To delete a server

1. In the Console, click **Settings** > **Data Insight Servers**.
2. Click the **Select Action** drop-down for the corresponding server in the servers listing table, and select **Delete**.

Note: Data Insight does not allow you to delete a server if it is associated with a storage device.

See "About automated alerts for patches and upgrades" on page 99.

See "Configuring advanced settings" on page 78.

Viewing Data Insight server details

You can view server details on the **Overview** tab and perform various tasks, such as, changing th eadvanced settings, monitor in-progress scans, view server statistics, and check event logs for the server.

To review server details

1. In the Console, click **Settings** > **Data Insight Servers**
2. Click the server that you want to review, or click **Select Action** and select **View**.

   The **Overview** tab of the server details screen appears. It displays the following information:

   - **Name**
     This is the address of the server configured when the server was added. Remote product servers use this address when communicating with this server. At this time, Data Insight does not support changing the address of the server.
   - **Roles**
     This indicates the roles that the server plays. Possible server role values are Management Server, Indexer, Collector, Classification Server, and Windows File Server Agent.
   - **Filer Name**
     If the server is a Windows File Server Agent, the name of the associated file server is displayed here.
   - **Data Insight Version**
     Indicates the version of Data Insight installed on this server.
   - **Operating System**
     Indicates the operating system installed on this server.
   - **CPUs**
     Indicates the number of CPUs available on this server.
   - **Memory**
     Indicates the RAM installed on this server in MBs.
   - **Associated Windows File Server**
     This detail is available only if you select the server with the Windows Filer Server agent role. Indicates the host name or IP address of the Windows File Server that the agent is monitoring.
   - **Server Health**

Indicates the current state of the server - whether a server is in a healthy state, is faulted, or at risk. You can also view the reasons for the current health state of the server.

- **Composition of Data Directory**
  The pie chart shows the disk space utilized by various folders under the `data` directory. These folders include `Collector`, `Indexer`, `Workflow`, `Attic`, `Inbox`, `Outbox`, `Reports`, `Classification`, and `Other`.
  You can configure purge settings to delete old data to free disk space on the server.
  See "Configuring data retention settings" on page 38.
- **Product Updates**
  The suggestion for upgrades if a newer version of Data Insight is available.

See "Configuring advanced settings" on page 78.

About node templates

A node template consists of a set of pre-defined node settings. To ensure consistency of configuration across all Data Insight server nodes, you can create templates with the required settings and apply them when configuring the servers.

Data Insight comes bundled with three standard templates which include frequently used settings to help you set up nodes for well known configurations, for example, templates for POC configuration, large Indexer, and large Collector. The standard node templates are stored in the `INSTALL_DIR\conf\node_template`folder. You cannot edit or delete the standard node templates from the console.

You can apply multiple templates when configuring a server. However, when you apply multiple templates to a node, Data Insight applies each template serially when evaluating the configuration for that node. You can also use a node template to change a single setting on all nodes. For example, if you want to change the **Ping timeout** setting for all nodes, you can create a template with the required timeout setting and apply to all Data Insight nodes in your environment.

Managing node templates

You can view configured templates, create new node templates, edit existing templates, and delete a template from the **Node Templates** list page.

To view configured node templates

1. In the Console, click **Settings** > **Data Insight Servers**.
2. From the **Node Templates** drop-down, select **Manage Node Templates**.
3. The list page displays all standard and configured node templates.

You can choose to delete configured node templates.

To delete a node template

1. In the Console, click **Settings** > **Data Insight Servers**.
2. From the **Node Templates** drop-down, select **Manage Node Templates**.
3. On the node templates list page, select the template that you want to edit, and from the **Select Action** drop-down, select **Delete**.
4. Click **OK** on the confirmation dialog to delete the node template.

See "Adding or editing node templates" on page 72.

Adding or editing node templates

You can create new templates that you can apply to multiple nodes simultaneously or edit existing templates to change the configuration settings.

To create a node template

1. In the Console, click **Settings** > **Data Insight Servers** .

   The Data Insight Servers list page displays.
2. From the **Node Templates** drop-down, select **Manage Node Templates**.
3. Click **Add New Node Template**.
4. In the node template name field , enter a unique name for the template.
5. For each configuration category, click **Edit**,.

   Select the check box for the settings that you want in the template, and configure appropriate values for the settings, as required.

   See "Configuring advanced settings" on page 78.
6. Click **Close Configuration** to save the settings as a node template.

To edit a node template

1. In the Console, click **Settings** > **Data Insight Servers**.
2. From the **Node Templates** drop-down, select **Manage Node Templates**.

3 On the node templates list page, select the template that you want to edit, and from the **Select Action** drop-down, select **Edit**.

4 On the **Edit Node Templates** page, change the required settings, and click **Close Configuration** to save the changes.

When you edit a node template, the changes in configuration do not automatically reflect on the Data Insight servers on which the template is applied. You must apply the modified template to the Data Insight server again for the configuration changes to take effect on the server.

Adding Portal role to a Data Insight server

In case of small deployments or POC environments, you can use an existing Data Insight Management Server or a Collector node as the Self-Service Portal node. You must add the Portal role to the Data Insight server you want to use as the Portal node.

You must ensure that you upgrade the node to which you want to add the Portal node to Release 4.5.1.

For more information about the Self-Service Portal node, see the *Veritas Data Insight Installation Guide*.

To add the Portal role to a Data Insight server

1 In the Console, click **Settings** > **Data Insight Servers**.

2 Click the **Select Action** drop-down for the corresponding server that you want to use as the Self-Service Portal, and select **Add Portal role**.

The option to add the Portal role to the Management Server, Collector node or Collector, and Indexer node is not available if the Portal role is already added to the server.

3 On the **Add Portal role** pop-up, enter the default port number. The DataInsightPortal service runs on port 443 by default.

However, if you want to designate the Management Server as the Portal node, you must enter a port number other than the default port number because the DataInsightWeb service also runs on port 443 on the Management Server.

4 Click **Configure** to designate the server as the Portal node and to install the DataInsightWorkflow service and the DataInsightPortal service on the server.

See "Configuring Data Insight services" on page 76.

Adding Classification Server role to a Data Insight server

Veritas recommends that you deploy a standalone Classification Server because assigning a Classification Server role to an existing Data Insight node will affect its performance.

You can assign a Management Server, Collector, or Collector and Indexer node to serve as a Classification Server instead of a designated standalone Classification Server. This indicates that the Data Insight node behaves as a Classification Server. However, you cannot add a Classification Server role to a pure Indexer, Self Service Portal node, and Windows File Server Agent node.

For more information about Classification Server, see the *Veritas Data Insight Installation Guide*.

To add the Classification Server role

1. In the Console, click **Settings** > **Data Insight Servers**.
2. Click the **Select Action** drop-down on the row corresponding to the server that you want to use as the Classification Server, and select **Add Classification role**.

   The option to add classification role is available only for the Management Server, Collector, or Collector and Indexer node. This option is not available when the role is already set to Classification Server.
3. Click **Yes** on the confirmation message to designate the server as the Classification Server. As part of this configuration, the `DataInsightVICClient` and the `DataInsightVICServer` services are installed on the designated server.

   Note that when you add a classification role to a server, it can have a impact on the performance of the server.

See "Assigning Classification Server to a Collector" on page 74.

For information about how Data Insight classifies files, see the *Veritas Data Insight Classification Guide*.

Assigning Classification Server to a Collector

Every data source is associated with a Collector node. The mapping of the Collector to the Classification Server decides to which Classification Server the request should be fanned out. When you submit paths for classification, a Data Insight server with a Collector role fans out the file paths to the assigned Classification Server which classifies those files.

Before assigning a Classification Server to a Collector, note the following:

- Make sure that a pure Classification Server is installed.
- Assign a Classification Server only to a Data Insight server with the Collector role.
- If a Collector is assigned a classification role, then it cannot be assigned as a Classification Server to another Collector.
- If there is any previously configured mapping, then it will be overwritten when you assign a Classification Server to a Collector.

See "Adding Classification Server role to a Data Insight server" on page 74.

To assign a Classification Server

1. From the Data Insight Management Console, click **Settings** > **Data Insight Servers** to display list of configured servers.
2. Select the server with the role as Management Server, Collector, or Collector and Indexer.
3. Click **Assign Classification Server**.
4. In the **Select Classification Server** window, review the information about the configured Classification Servers. Click **Select** on the row corresponding to the Classification Server with which you intend the Data Insight server to communicate. Only pure Classification Servers are listed on this pop-up.
5. Click **OK** on the confirmation message.

Alternatively, you can also assign a Classification Server from the command line interface of the Management Server by issuing the following command:

```
configdb -o -T node -k <node id of the Collector node> -J
vic.server.id -j <node id of Classification Server>
```

Deleting the assignment of a Classification Server to a Data Insight server

You can delete the mapping with the Classification Server by issuing the following command on the Management Server:

```
configdb --delete-obj-attr -J vic.server.id -T node -k <node id of
the Collector node>
```

Viewing in-progress scans

You can view a list of currently running scans on the **In-progress scans** tab.

To review the in-progress scans

1. On the product server details page, click **In progress** scans.
2. Review the following information from the in-progress scans table:
   - **Object Name** - Name of the object being scanned.
   - **Object Type** - The type of object being scanned. This can be a share, site collection or Active Directory.
   - **Task Name** - Indicates the type of the scan.
   - **Task State** -Whether the task is RUNNING or IN QUEUE.
     If none of the tasks are in RUNNING state, it usually means a scan pause window is in effect. You can configure the pause interval for the server from **Advanced Settings** page for the server. To override the pause schedule for a share or site collection and start the scan immediately, from the Action drop-down, select **Override pause schedule**.
     See "Configuring advanced settings" on page 78.
   - **Start Time** - Time the scan started if it's in RUNNING state.
   - **Time elapsed** - Indicates how long the scan has been running.
   - **Task Statistics**- Indicates the statistics of the in-progress scans. It shows the number of folders scanned and the files or folders scanned per minute.
3. Click **Cancel** to cancel a particular scan or click **Cancel All** to cancel all scans.

Note: To view the in-progress scans across all nodes in your environment, navigate to the **Settings** > **Scanning** > **In-progress Scans** tab.

Configuring Data Insight services

You can view the status of all Data Insight processes from the **Services** tab on the Management Console. You can manage the following services from this page:

- Data Insight Web server service - The process runs on the Management Server. It starts by default when you start your Management Server.
- Data Insight Communication service - The process runs on each node in a Data Insight deployment. This service is responsible for all inter-node communication.
- Data Insight Configuration service - The process that provides interface to configuration and other product data that is stored on the local system.
- DataInsightFpolicy service - The process runs on the Collector node or the Management Server. This service is responsible for registering the Data Insight

server with the NetApp filer and enables Data Insight to receive access events from the filer.

- DataInsightFpolicyCmod service - The Windows service runs on the Collector node. It is responsible for interacting with the NetApp Cluster Management host to receive access events from the nodes in the cluster.
- DataInsightCelerra service - The process runs on the Collector Worker node or the Management Server. This service is responsible for registering the Data Insight server with the EMC Celerra filer and enables Data Insight to receive access events from the filer.
- DataInsightWatchdog service - The process runs on all nodes in a Data Insight deployment and monitors the health of the node. The service also monitors the disk usage on the Windows File Server node and prevents it from running out of disk space by implementing safeguards.
- DataInsightWinNAS service - The process runs on the Windows File Server. The service receives event information from the Windows File Server filter driver and transfers it to the Collector node that is configured for the filer.
- DataInsightGenericCollector service - The service runs on the Collector associated with a generic file server. The service collects all incoming events from generic file servers and web API clients, and copies them to a specific folder on the Collector.
- DataInsightWorkflow service - The service runs only on the Management Server. This service is responsible for managing the lifecycle of various actions initiated from the Management Server.
- DataInsightPortal service - The service runs on any server that is designated as the Portal node. It provides an interface to the portal where the custodians can log on to take remediation action. The service runs on the Management Server and the Portal node.
- DataInsightVICClient service- This service runs on the Management Server and any Data Insight server that is designated as the Classification Server. It is responsible for interactions that happen between the Data Insight servers and the Classification Server such as when a file is submitted for classification, this service routes the request to the Classification Server.
- DataInsightVICServer service - This service runs on the Management Server and anyData Insight server that is designated as the Classification Server. This service hosts the classification engine which is responsible for evaluating files against the configured policies and detecting the patterns.

For detailed information about the Data Insight services, see the *Veritas Data Insight Installation Guide*.

Depending on the type of the filers managed by the Collector, you can enable the FPolicy, EMC Celerra, or Genericcollector service on the server from this page.

To enable or reconfigure the FPolicy or EMC Celerra service

1. On the **Services** tab, click the service that you want to enable on the server.
2. From the **Select saved credential** drop-down, select the credential that the service uses to run. Ensure that the user used to configure the FPolicy service is added to the Group Policy object with the **Log on as a service** privilege in your Active Directory domain controller.
3. If configuring the DataInsightFpolicy service, enter the name of the policy.
4. If configuring the DataInsightCelerra service, select one of the following to specify the location of the server on which the EMC CAVA service is installed:
   - EMC CAVA Service is installed locally on this server
   - Remote EMC CAVA Server Pool publishes events to this server
5. Click **Configure**.

See "Configuring advanced settings" on page 78.

See "Credentials required for configuring NetApp filers" on page 128.

See "Credentials required for configuring EMC Celerra filers" on page 173.

Configuring advanced settings

You can edit various settings of the Data Insight servers from the **Settings** > **Data Insight Servers** > **Advanced Settings** page.

The advanced settings are divided into the following categories:

- **Filesystem Scanner settings** - Configures how the server scans file systems. Data Insight performs two types of scans on the configured shares:
  - Full scans
    During a full scan, Data Insight scans the complete share. These scans can run for several hours, if the share is very big. Typically, a full scan should be run once for a newly added share. After the first full scan, you can perform full scans less frequently based on your preference. Ordinarily, you need to run a full scan only to scan those paths which might have been modified while event monitoring was not running for any reason. In all other cases, the incremental scan is sufficient to keep information about the file system metadata up-to-date.
    See Table 4-1 on page 80.
  - Incremental scans

During an incremental scan, Data Insight re-scans only those paths that have been modified since the last full scan. It does so by monitoring incoming access events to see which paths had a create event or write event on it since the last scan.

See Table 4-2 on page 82.

- **Indexer settings** - Configures how the indexes are updated with new information. This setting is applicable only for Indexers.

  See Table 4-5 on page 85.

- **Audit events preprocessor settings** - Configures how often raw access events coming from file servers must be processed before they are sent to the Indexer.

  See Table 4-6 on page 87.

- **High availability settings** - Configures how this server is monitored.

  Each server periodically monitors its CPU, memory, state of essential services, number of files in its inbox, outbox, and err folders. Events are published if these numbers cross the configured thresholds. Also, each worker node periodically heartbeats with the Management Server. The Management Server publishes events if it does not receive a heartbeat from a node in the configured interval.

  See Table 4-8 on page 89.

- **FPolicy safeguard settings** - Configures the safeguards related to FPolicy communication. You can either choose to use the global settings or customize the settings for a specific Collector node. You can configure settings for FPolicy Cluster-Mode and 7-mode

  See Table 4-7 on page 88.

- **Report settings** - Configures settings for reports.

  See Table 4-11 on page 92.

- **Classification settings** - Configures how the Classification Server scans file system contents.

  Table 4-9
  Table 4-10

- **Windows File Server Agent settings** - Configures the behavior of the Windows File Server filter driver. This setting is applicable only for the Windows File Server Agent server.

  See Table 4-12 on page 93.

- **Veritas File System server (VxFS) settings** - Configures how Data Insight scans the VxFS filer.

  See Table 4-13 on page 94.

- **NFS settings** - Configures how Data Insight scans NFS shares.

  See Table 4-14 on page 94.

- **SharePoint settings** - Configures the duration for which old audit logs are kept on the SharePoint server. Audit logs that are fetched from the SharePoint server are automatically deleted from the Data Insight database. You can disable this feature at the web application level.
  See Table 4-15 on page 95.
- **Troubleshooting settings** - Configures settings that aid troubleshooting.
  See Table 4-16 on page 96.
- **Set custom properties** - Configures certain advanced properties of a Data Insight worker node. Using this facility, you can customize certain properties that are not accessible by the normal settings.

Note: Veritas recommends using the custom properties settings under the guidance of Veritas Support.

You can configure the advanced settings per node or save commonly used settings as a templates. See "About node templates" on page 71.

To configure advanced settings

1. In the Console, click **Settings** > **Data Insight Servers**.
2. Click the server, for which you want to configure the advanced settings.
3. Click **Advanced settings**.
4. Click **Edit**.
5. Make necessary configuration changes, and click **Save**.

See "Managing node templates" on page 71.

Each of the categories for the advanced settings are described in detail below.

Table 4-1 File system scanner settings - Full scan settings

Setting	Description
Maximum scans to run in parallel on this server	The Collector can perform multiple full scans in parallel. This setting puts a limit on the total number of full scans that can run in parallel on a Collector. The default value is two threads. Configure more threads, if you want scans to finish faster. The setting is disabled by default.

Table 4-1 File system scanner settings - Full scan settings *(continued)*

Setting	Description
Maximum shares per filer to scan in parallel	If multiple shares of a filer can be scanned in parallel, this setting puts a limit on the total number of shares of a filer that you can scan in parallel.
Default scan schedule	Specifies how often full scans must be performed. By default, full scans are scheduled to repeat at 19:00 P.M. on the last Friday of each month. Select the check box to override this setting at a filer or at a share level.
Pause scanner for specific times	You can configure the hours of the day when scanning should not be allowed. This setting ensures that Data Insight does not scan during peak loads on the filer. The setting is enabled by default. Scans resume from the point they were at before they were paused.
Pause scanner schedule	Configures when scanning should not be allowed to run. By default, scanning is paused from 7:00 A.M. to 7.00 P.M., Monday to Friday. You can specify multiple scanner pause schedules for different days of the week. For example, you can choose to pause scanning from 7:00 A.M. to 7:00 P.M. on weekdays and from 7:00 A.M. to 9:00 P.M. on Saturdays and Sundays. To add a scanning schedule: **1** Click **Add**. **2** On the Pausing schedule pop-up, select the time period and the days on which you want to pause scanning. **3** Click **Save**. You can also edit or delete existing scanning schedules.

Table 4-2 File system scanner settings - Incremental scan settings

Setting	Description
Maximum scans to run in parallel on this server	The Collector can perform multiple full scans in parallel. This setting puts a limit on the total number of full scans that can run in parallel on a Collector. The default value is two threads. Configure more threads, if you want scans to finish faster. The setting is disabled by default.
Maximum shares per filer to scan in parallel	If multiple shares of a filer can be scanned in parallel, this setting puts a limit on total number of shares of a filer that can be scanned in parallel. The default value is 2.
Default scan schedule	Specifies how often incremental scans must be performed. By default, incremental scans are scheduled to run at 7:00 P.M. each night. Schedule incremental scans more or less frequently based on how up-to-date you need information in Data Insight to be.
Pause scanner for specific times	You can configure hours of the day when scanning should not be allowed. This setting ensures that Data Insight does not scan during peak loads on the filer. This setting is enabled by default. Scans resume from the point they were at before they were paused.
Pause scanner schedule	Configures when scanning should not be allowed to run. By default, scanning is paused from 7:00 A.M. to 7:00 P.M. Monday to Friday.

Table 4-3 File system scanner settings - Throttling for NetApp filer

Setting	Description
Use global settings	This option is selected by default. When this option is selected, Data Insight uses the throttling thresholds defined on the **Settings** > **Global Settings** > **Scanning and Event Monitoring** page. See "Configuring scanning and event monitoring " on page 28.
Use custom settings	Select **Use custom settings** to disable scanning or override the global throttling thresholds and define custom values.

Table 4-3 File system scanner settings - Throttling for NetApp filer *(continued)*

Setting	Description
Throttle scanning based on latency of the filer	Clear the check box to disable scanning for the filers that the collector is monitoring. Select to enable throttling of Data Insight scans for NetApp 7-mode and Cluster-Mode file servers. This option is not selected by default. Data Insight collects latency information from NetApp file servers. It can use this information to throttle scanning, if latency of the file server increases above a certain level. This ensures scanner does not put additional load on the file server during peak load conditions. You can configure the following parameters to enable throttling for NetApp file servers: ■ Latency threshold - Specify latency in milliseconds, which when crossed, should throttle scanning for the file server ■ Minimum pause - Specify the minimum duration (in milliseconds) for which the scanner should pause between paths when in throttling mode. ■ Back off value - If increased latency is sustained, pause interval will be increased by the Back off value specified (in milliseconds). ■ Maximum pause - Specify the maximum pause interval for the scanner (in seconds). If exceeded, pause interval is no longer incremented by Back off value.

Table 4-4 File system scanner settings - Common settings

Setting	Description
Scanner snapshot interval	Scanning a big share can take several hours. The scanner periodically saves information to a disk so that information is visible sooner without waiting for the entire scan to finish. You can configure how often information is saved to the disk by the scanner. By default, the scanner creates a snapshot of new information every 300 seconds (5 minutes). The minimum value you can set for this parameter is 300.

Table 4-5 Indexer settings

Setting	Description
Total indexer threads	The indexer processes incoming scan and access event information for various shares and updates the per-share database. This setting configures how many databases can be updated in parallel. By default 2 threads are configured. Specify a larger value for bigger setups where indexer is not able to keep up with incoming rate of information. This is indicated when you observe too many files in the inbox of the Indexer worker node. However, you must ensure that the Indexer has adequate CPU and memory when configuring a higher number of indexer threads. You need approximately 1 GB of RAM per indexer thread.
Limit maximum events processed in memory	By default, the indexer processes all new incoming events in memory before saving the events to the disk. If your are falling short of RAM on your Indexer, you can limit the maximum number of events that the indexer processes in memory before it saves them to the disk. **Note:** Specifying a small number makes the indexing very slow.

Table 4-5 Indexer settings *(continued)*

Setting	Description
Reconfirm deleted paths when reconciling full scan information	After Data Insight indexes full scan data, it computes the paths that no longer seem to be present on the file system. Similarly, if Data Insight discovers folders or subsites missing in a SharePoint site, then it computes those subsites or folders as deleted without performing a reconfirmation scan. Select this check box to have Data Insight re-confirm if those paths are indeed deleted using an incremental scan before removing them from the index. When the check box is clear, Data Insight readily removes the missing paths from the indexes without carrying out a re-confirmation. **Note:** Re-confirm scan is not supported for site collections.
Indexer schedule	Specify how often an index should be updated with new information. By default, all new data is consumed once every four hours. Indexer gets better throughput if more information is given to it when indexing. However, if you configure a very high value, new information will not be visible in the Console for a much longer period.
Indexer integrity checking schedule	Data Insight checks the integrity of its databases once a week. If any errors are found in the database, an event is published. You can configure a different schedule if required.

Table 4-6 Audit events preprocessor settings

Setting	Description
Audit events preprocessor schedule	Incoming raw audit events from file servers must be pre-processed before sending them to the Indexer. At this stage, `collector.exe` applies various heuristics to the raw events and also removes transient events. By default, raw events are processed every 2 hours.
Batch size (MB)	The maximum size of the raw audit event files that a single Collector thread can process. The default batch size is 2 GB.
Total Collector threads	The Collector can run multiple pre-processors in parallel. This setting configures how many instances can run in parallel.

Table 4-7 FPolicy safeguard settings

Setting	Description
Use global settings	This option is selected by default. Use the FPolicy safeguard settings for NetApp 7-mode and Cluster-mode filers as defined on the **Settings** > **Global Settings** > **Scanning and Event Monitoring** page. See "Configuring scanning and event monitoring " on page 28. Data Insight collects latency information from NetApp file servers. It can use this information to initiate safeguard mode, if latency of the file server increases above or falls below a certain level. When the safeguard is in effect, Data Insight drops its FPolicy connection to the filer. This ensures event collection does not put additional load on the file server in peak load conditions. If the latency on the physical file server increases above the configured threshold, Data Insight disconnects from the associated virtual file server. This information is also displayed on the Data Insight System Overview dashboard.
Use custom settings	Select **Use custom settings** to disable the safeguard settings or to override the global safeguard thresholds and define custom values.

Table 4-7 FPolicy safeguard settings *(continued)*

Setting	Description
Enable FPolicy safeguard settings	Select one of the following, as appropriate: ■ **Select Enable FPolicy 7-Mode Safeguard** to initiate safeguard mode for the NetApp file servers. ■ **Select Enable FPolicy 7-Mode safeguard for VFilers** to initiate safeguard mode for virtual file servers by monitoring latency statistics for the associated physical NetApp filers. The safeguard mode for the virtual filers is initiated only if the details of the physical filer corresponding to the virtual filer are provided while adding the virtual filer. ■ **Select Enable FPolicy Cluster Mode Safeguard** to initiate safeguard mode for the NetApp cluster mode file servers. These FPolicy safeguard settings are not selected by default. When these check boxes are cleared, the safeguard settings are not in effect. Configure the following values: ■ The high and low thresholds for CIFS and NFS. ■ The number of samples to be considered to calculate the average latency. ■ The minimum time to wait before Data Insight reconnects to the filer after a disconnection.

Table 4-8 High availability settings

Setting	Description
Ping timeout (in minutes)	If a worker node does not heartbeat in the specified interval, Management server will publish an event to that effect. This setting is only applicable for the Management Server.
Notify when CPU continuously over (percentage)	If CPU used on this server is consistently over the specified percentage, an event is published. (Default value: 90%)

Table 4-8 High availability settings *(continued)*

Setting	Description
Notify when memory continuously over (percentage)	If Memory used on this server is consistently over the specified percentage, an event is published. (Default value: 80%)
Notify when disk usage over (percentage)	If disk usage, either for the system drive or data drive, is over the specified threshold, an event is published. (Default value: 80%)
Notify when disk free size under (MB)	If the free disk space for the system drive or data drive is over the specified threshold in megabytes, an event is published. (Default value: 500 MB)
Notify when number of files in `err` folders over	If Data Insight is not able to process an incoming file for some reason, that file is moved to an `err` folder. Data Insight publishes an event if number of files in the `err` folder crosses the specified threshold. (Default value: 50)
Notify when number of files in `inbox` and `outbox` folder over	If Data Insight is not able to process incoming data fast enough, the number of files in the transient folders, `inbox` and `outbox`, goes on building up. Data Insight publishes an event if number of files crosses the configured threshold. (Default value: 5000)

Table 4-9 Classification settings - Content scans for file system

Setting	Description
Maximum content scans to run in parallel on this server	The Classification Server can perform multiple content scans in parallel. This setting puts a limit on the total number of scans that can run in parallel on a Classification Server. The default value is two threads. Configure more threads, if you want scans to finish faster. Ensure that the server has enough resources to run the configured parallel threads.
Maximum classification threads to run in parallel on this server	When content is being classified in parallel, this setting puts a limit on the total number of classification threads that you can run in parallel.

Table 4-9 Classification settings - Content scans for file system *(continued)*

Setting	Description
Maximum shares per filer to content scan in parallel	When multiple shares of a filer are scanned in parallel, this setting puts a limit on the total number of shares of a filer that you can scan in parallel.
Pause content scanner for specific times	You can configure the hours of the day when content scanning should not be allowed. This setting ensures that Data Insight does not scan during peak loads on the filer. The setting is enabled by default. Scans resume from the point they were at before they were paused. Content scanning is paused by default from Monday to Friday from 7:00 A.M. to 7:00 P.M. To configure the pause schedule, click **Edit** on the scheduler to configure the days and time on which content scanning should be paused. You can configure more than one pause schedules. Click **Add** to add new pause schedule.

Table 4-10 Classification settings - Content scans for SharePoint

Setting	Description
Maximum content scans to run in parallel on this server	The Classification Server can perform multiple content scans in parallel. This setting puts a limit on the total number of scans that can run in parallel on a Classification Server. The default value is two threads. Configure more threads, if you want scans to finish faster. Ensure that the server has enough resources to run the configured parallel threads.
Maximum classification threads to run in parallel on this server	When content is being classified in parallel, this setting puts a limit on the total number of classification threads that you can run in parallel.

Table 4-10 Classification settings - Content scans for SharePoint *(continued)*

Setting	Description
Maximum site collections per web application to content scan in parallel	If multiple site collections can be scanned in parallel, this setting puts a limit on the total number of site collections per web application that you can scan in parallel.
Pause content scanner for specific times	You can configure the hours of the day when content scanning should not be allowed. This setting ensures that Data Insight does not scan during peak loads on the SharePoint server. The setting is enabled by default. Scans resume from the point they were at before they were paused. Content scanning is paused by default from Monday to Friday from 7:00 A.M. to 7:00 P.M. To configure a different pause schedule, click **Edit** on the scheduler to configure the days and time on which content scanning should be paused. You can configure more than one pause schedules. Click **Add** to add new pause schedule.

Table 4-11 Reports settings

Setting	Description
Maximum memory when generating report output	Specifies the maximum memory that can be used for generating a report output. By default, it is 1024 MB on a 32 bit machine and 2048 MB on a 64 bit machine
Total threads for generating report output	Configure the number of threads for generating report output (PDF/HTML/CSV) in parallel. Default value is 2. This setting applies to the Management Server.

Table 4-11 Reports settings *(continued)*

Setting	Description
Number of threads for a single report run	Configure the number of threads responsible for generating the report output database for a single report. This configuration applies to the Indexer node. This setting helps you speed up the process of report generation. For a particular Data Insight server, the thread count applies to all types of reports.
Maximum reports that can run simultaneously	By default, Veritas Data Insight Administrator's Guide executes two reports in parallel. However, you can configure a higher value to run multiple reports in parallel.

Table 4-12 Windows File Server agent settings

Setting	Description
Maximum kernel ring buffer size	The Windows File Server filter driver puts events in an in-memory buffer before the DataInsightWinnas service, consumes them. By default, it uses a 10MB buffer. You can use a bigger buffer. Data Insight publishes an event that indicates events are being dropped due to a high incoming rate. **Note:** Buffer is in kernel and is limited on a 32 bit operating system.
Ignore accesses made by Local System account	The Windows File Server filter driver ignores accesses made by processes running with Local System account. This setting ensures that Data Insight can ignore most events originating from the operating system processes or other services like antivirus and backup. Clear this check box to enable monitoring accesses made by LOCAL SYSTEM account. This is not recommended on a production file server.

Table 4-13 Veritas File System server settings

Setting	Description
Flush events on VxFS filer before audit	Set this option to true, if you want to force VxFS to flush its events to disk each time requests for information. This option is useful in Proof-of-Concept (POC) setups and enables you to see events faData Insightster.
Maximum number of audit threads	This option determines how many filers to fetch audit information from in parallel.
Maximum kernel ring buffer size (Number of records)	The access event records are saved in a log file on the VxFS filer before Data Insight consumes them. By default, 50,000 records can be saved in the log file. You can also specify a larger number. Data Insight publishes an event that indicates that events are being dropped due to a high incoming rate.

Table 4-14 NFS settings

Setting	Description
Set default credentials for NFS scanner	Set this option to true if you want to allow Data Insight to use the specified User and Group ID to log in to scan NFS shares.
User ID	The ID of the NFS user that the Data Insight uses to scan the filer. You can set the value to 0 to allow root access from the Data Insight scan hosts.
Group ID	The ID of the group that the Data Insight uses to scan the filer. You can set the value to 0 to allow root access from the Data Insight scan hosts.

Table 4-15 SharePoint settings

Setting	Description
Automatically delete audit events from SharePoint server that are older than (days)	When configuring a SharePoint web application, you can choose to let Data Insight delete audit logs that have already been fetched from SharePoint. By default, Data Insight deletes audit logs older than two days. Deletion of audit logs takes place every 12 hours.
Schedule to fetch audit events from SharePoint server	Data Insight fetches new audit events from SharePoint periodically. By default, it does so every 2 hours. You can configure a different schedule.
Scan multiple site collections of a web application in parallel	The Collector can perform multiple full scans in parallel. This setting configures how many full scans can run in parallel. The default value is 2 parallel threads. Configure more threads if you want scans to finish faster. The setting disabled by default.
Maximum site collections per web application to scan in parallel	If multiple site collections of a web application can be scanned in parallel, this setting puts a limit on the total number of site collections of a web application that you can scan in parallel.
Default scan schedule	Specifies how often scans need to be performed. You can override this setting at a web application or site collection level. By default, scans are scheduled to repeat 11:00 P.M. each night.
Pause scanner for specific times	You can configure the hours of the day when scanning should not be allowed. This ensures that Data Insight does not scan during peak loads on the SharePoint servers. The setting is enabled by default. Scans resume from the point they were at before they were paused.
Pause scanner schedule	Specify when scanning should not be allowed to run. By default, scanning is paused from 7:00 A.M. to 7:00 P.M., Monday to Friday.

Table 4-15 SharePoint settings *(continued)*

Setting	Description
Pause auto-delete for specific times	You can configure the hours of the day when auto-delete of audit events from SharePoint server should not be allowed. This feature can help you to avoid overloading the SharePoint servers during the peak hours. By default, deletion of audit logs takes place every 12 hours.
Pause schedule for auto-delete	Specify when auto-delete of the audit logs should not be allowed to run.

Table 4-16 Troubleshooting settings

Setting	Setting
Preserve intermediate files	As new data comes into a Data Insight system, it moves between various modules. In this process the original files are deleted and a new processed file is generated for the next stage of processing. To aid troubleshooting, select this check box to retain the intermediate data files. These files get stored in `attic` folder in the data directory.
Preserve raw audit event files	Events processed by the Audit Pre-processor stage are deleted once consumed. If this setting is enabled, raw audit event files will be preserved in the `attic` folder in the data directory.

See "Managing Data Insight product servers" on page 67.

Monitoring Data Insight jobs

You can monitor and execute the jobs that run on remote Data Insight servers from the Management Console. The view enables you to search for the jobs are part of the communication Web service, the DataInsightWatchdog service, and the DataInsightWorkflow service. These jobs either run continuously or are scheduled.

For more information about the different Data Insight services, see the *Veritas Data Insight Installation Guide*.

Note: You can view the SPEnableAuditJob and the SPAuditJob only if the server is configured to be the Collector for a SharePoint site collection.

To view and execute jobs from the console

1. On the Management Console, click **Settings** > **Data Insight Servers**.
2. On the list page, click **Select Action** > **View** for the server on which you want to view or execute a process or a job.

   The overview page for the server opens by default.
3. Click the **Jobs** sub-tab.

   The page lists all jobs that are scheduled on the server. Review the following details about a job:

   - The status of the job, whether stopped or running.
   - The date and time when the job was previously run.
   - The date and time of the next scheduled run.
   - The schedule of the job. For example, every day, every hour, every one minute.
   - The service the job runs on, such as the Communication service or the DataInsightWatchdog service.
4. To execute a job, do one of the following :

   - On the Jobs sub-tab, click **Select Action** > **Run** to run the job without waiting for the next scheduled run.
   - On the **Data Insight Server** list page, click the **Run Job** drop-down, and select the job you want to execute.

See "Scheduled Data Insight jobs" on page 370.

Rotating the encryption keys

Encryption keys are used to secure communication between the Data Insight servers and to decrypt encrypted passwords for the stored credentials. You can periodically change the encryption key to ensure security and to protect data.

To rotate the encryption key

1. On the Management Console, click **Settings** > **Data Insight Servers**.
2. On the list page, click the Management Server.
3. On the Management Server details page, click **Jobs**.

4 On the **Jobs** list page, search for *KeyRotationJob*.

5 On the **Select Action** drop down, click **Run** to change the encryption key.

When the *KeyRotationJob* runs, the key is regenerated and all passwords are re-encrypted with the new key.

For more information about the Data Insight security configuration, see the *Veritas Data Insight Installation Guide*.

Note: The *KeyRotationJob* is not automatically scheduled. You must manually run the job when you want to rotate the encryption keys. New keys are replicated to the other worker nodes job by the *updateconfig* job. Ensure that the network connection exists between the Data Insight nodes for receiving the new keys.

Viewing Data Insight server statistics

You can view the line graphs that indicate the health of the Data Insight servers. The DataInsightWatchdog service collects the server statistics. You can use the information to know the performance trends for each Data Insight server. For more information about the DataInsightWatchdog service, see the *Veritas Data Insight Installation Guide*.

The line graphs display hourly, weekly, monthly, and yearly data.

To view server statistics

1 Click **Settings** > **Data Insight Servers** > ***Server Name*** > **Statistics**.

2 Select the following to filter the data that is displayed on the page:

- The charts that you want to view.
- The duration for which you want to view the data.
- The type of statistics that you want to view - Average, Minimum, and Maximum.

3 The **Statistics** sub-tab displays the following high-level data:

- The number of files in the `inbox` and the `outbox` folders.
- The number of files in the `classification inbox` and `classification outbox` folders.
- The number of error files in the `scanner/err`, `collector/err`, `indexer/err`, and `classification/err` folders.
- The CPU and memory usage on the Data Insight servers.

- The disk space utilization on the Data Insight servers on the system disk, Data Insight installation disk, and the different folders within the Data Insight `data` directory.
- Incoming event rate for this server if it is a collector for one or more filers.
- Throughput for the event pre-processor.
- The raw event files being processed by the Collector, that is, the number of events that are processed by `collector.exe` per second.
- The scan and audit files being processed by the Indexer
- Internal event files being processed by the Management Server.

4 You can view each processing backlog chart from three perspectives. Click one of the following aspects of the backlog:

- The size of the backlog in MB or GB; the total size of all files that are yet to be processed.
- The count of files that are yet to be processed.
- The time lag in terms of hours or days; the time lag is the difference between the current time and the file with the oldest timestamp that is yet to be processed.

5 For the Collector and Indexer backlogs, click the drop-down to the right of the chart to view the Top 10 objects that are contributing to the backlog. In case of Collector backlog, objects are the filers or web application and case of Indexer backlog, objects mean the shares or site collections

The Top ten chart is a bar chart.

See "Configuring advanced settings" on page 78.

About automated alerts for patches and upgrades

Data Insight simplifies the task of installing upgrades and patches by providing you with automated alerts and suggestions. Data Insight fetches this information from Veritas Operations Readiness Tool (SORT) to help you keep track of the product updates applicable to your installed version.

For Data Insight to communicate with Veritas SORT, the Data Insight Management Server must have an active Internet connection to the web site,*https://sort.veritas.com/*.

Data Insight displays the following upgrade recommendations for each of its product servers:

- Rolling Patches (RPs) that are available for the installed version. Only the latest rolling patch is displayed.
- Recommendation for a new product version appears at the footer of the Data Insight servers page.

See "Viewing and installing recommended upgrades and patches" on page 100.

Viewing and installing recommended upgrades and patches

You can view the recommendations regarding patches and the product version upgrades from the **Data Insight Servers** page. Using the download links you can download the recommended patched and install them on your **Data Insight Servers**.

To view and install a patch for your product server:

1. In the Management Console, navigate to **Settings** > **Data Insight Servers** to display a list of configured product servers.
2. Data Insight displays the recommendations for the patches under the **Product Updates** column. The recommendation for upgrading the product version is displayed on the footer of the page.

   Note: The **Product Updates** column is displayed only when the Data Insight Management Server is able to connect to the SORT website. When there is no connection, an error message is displayed in the footer.

3. Click the link to the latest patch that Data Insight recommends. You will be redirected to the SORT website.
4. Download the patch from the **Downloads** page on the website.
5. You can refer to the README on the page for the installation instructions and to verify the problems that have been fixed in the patch.

Deploying upgrades and patches remotely

You can remotely deploy installers and patches on the Data Insight worker nodes and Windows file server agents from the Data Insight Management Console. This simplifies the task of upgrading and configuring numerous nodes individually in a large Data Insight deployment. You can do the following:

- Install rolling patches on the Indexer nodes, the Collector nodes, and the Windows file server agents. Remote installation of rolling patches is not supported on Linux indexer nodes.

- Deploy the product installer on the Collector nodes and the Windows file server agents.

  Note: Remote upgrade of an Indexer node (neither Windows nor Linux Indexer nodes) is not supported. You must upgrade an Indexer node manually.

- Windows File Server agents can be upgraded remotely, from the **Settings** > **Data Insight Servers** page of the Data Insight Management Console.
- For Collector nodes, both upgrade and installation of rolling patches can be done remotely.

Note: If the first three version numbers for Management server and the concerned node match, then only the option to install a rolling patch is available. For example, when upgrading from 4.5.0 to 4.5.1, if the first three numbers do not match, then the option to upgrade must be used. For example, when upgrading a node from 4.0.0 to 4.5.0. However, for Windows File server nodes, a rolling patch can be installed in spite of version mismatch between the Management Server and the concerned Windows File Server nodes, because backward compatibility is supported.

You can either use existing saved credentials for upgrading a node or create new credentials.

You can also perform all the remote deployment actions using the `installcli.exe` utility from the Windows command prompt. For detailed information on `installcli.exe`, see the *Command File Reference*.

To remotely deploy upgrades and patches

1. In the Management Console, navigate to **Settings** > **Data Insight Servers** to display a list of configured product servers.
2. Select the server node for which you want to remotely deploy patch or upgrade.
3. Click the **Install** drop-down.
4. Do any of the following:
   - Click **Install Rolling Patch** for installing a rolling patch on the selected node.
   - Click **Upgrade** for installing a product version upgrade.

You can view the progress of the remote deployment operation from the **Installation Status** page.

See "Viewing the status of a remote installation" on page 104.

Using the Upload Manager utility

Use the Upload Manager to upload agent bundle zip files, var files, and patch installers on the worker nodes.

Before you can install the agent, ensure that the Windows File Server agent packages are uploaded on the relevant Collector nodes. You can use the Upload Manager utility to upload the agent packages to the Collector nodes in your Data Insight configuration.

To upload the agent packages

1. In the Console, click **Settings** > **Upload Manager**.
2. Browse to the location where the agent packages are saved.
3. Select the Collector nodes on which you want to upload the packages.
4. Click **Upload Bundle**.

The agent installation bundle is a zip file that contains the agent installer and various installation template files. There is one bundle for each processor architecture. You must upload the appropriate bundles to the Collector worker nodes based on the architecture of your file servers. The bundles are available along with the main install media and have the name, Veritas_DataInsight_windows_winnas_5_1_0_XXX_arch.zip. You can customize the agent installation by extracting the bundle in a temporary location, editing the installation templates as required, recreating the zip bundle, and then uploading the updated bundle to the appropriate Collector nodes using the Upload Manager utility.

To install a rolling patch, upload the rolling patch executable to the Management Server node. To upgrade Collector nodes, upload the newer version of the product installer to the Management Server.

Note: When a new version is installed on the Management Server, the installer automatically copies itself to the `installers` folder of the Management Server. In such cases, you do not need to separately upload the package to the Management Server for upgrading other worker nodes (except the Windows File Server agents).

Note: Remote install for Linux indexers is not supported.

About migrating storage devices across Indexers

Every storage device that is configured in Data Insight has exactly one Indexer node associated with it. The Indexer node processes the audit data from the Collector to service queries from the Management Server.

You may consider migrating a storage device to another Indexer in the following situations:

- If the existing Indexer for the storage device displays high resource utilization and the backlog of files being processed starts accumulating on the node. Migrating to another Indexer helps balance the load on the existing Indexer. See "Monitoring the performance of Data Insight servers" on page 289.
- If the existing Indexer is being decommissioned.
- If you want to migrate from a Windows server to a Linux server.

Before you decide to migrate the storage device to another Indexer, ensure the following:

- Both the source and the destination Indexers are in good health and can communicate with each other using the fully qualified domain name. Ensure that port 8383 is open between the indexers to allow the Communication Service to communicate between them.
- The destination Indexer has enough disk space, CPU, and memory capacity to take the additional load.
- The Collector nodes that are associated with the storage device have enough disk space, because the files may accumulate on the Collector nodes during the migration.
- All classification requests involving the source Indexer are in Completed state. Pending and in-progress classification requests are canceled during the migration.
- The migration should take place during quiet hours, so that additional files for consumption are not flowing in rapidly.

See "Editing filer configuration" on page 236.

See "Editing web applications" on page 260.

See "Managing cloud storage accounts" on page 274.

Viewing the status of a remote installation

You can view the progress of a remote installation of a Data Insight server or a Windows File Server agent, whether you have initiated the installation from the Management server or by using the command line utility.

To view the status of a remote installation

1. From the Data Insight Management Console, navigate to **Settings** > **Installation Status**.
2. The progress of the ongoing operations is displayed along with the following information:
   - The node for which the installation was initiated.
   - The status of the installation.
   - Time when the installation was initiated.
3. Click **View Progress** to view a more detailed status of the install operation.

Chapter 5

Configuring saved credentials

This chapter includes the following topics:

- About saved credentials
- Handling changes in account password

About saved credentials

An authentication credential can be stored as a saved credential in a central credential store. It can be defined once, and then referenced by any number of filers, shares, and Active Directory servers. Passwords are encrypted before they are stored.

The saved credential store simplifies management of user name and password changes.

You can add, delete, or edit stored credentials.

See “Managing saved credentials ” on page 105.

Managing saved credentials

You can add saved credentials to Data Insight, view details of the configured credentials and delete one or more saved credentials on the Saved Credentials details page.

You can add new credentials to the credential store. These credentials can later be referenced with the credential name.

To add a saved credential

1. In the Management Console, click **Settings** > **Saved Credentials**, and click **Create Saved Credentials**.
2. Enter the following information:

Saved Credential Name	Enter your name for this stored credential. The credential name must be unique within the credential store. The name is used only to identify the credential.
Access Username	Enter the user name for authentication.
Access Password	Enter the password for authentication.
Confirm Password	Re-enter the password.
Domain	Enter the name of the domain to which the user belongs.

3. Click **Save**.
4. You can later edit or delete credentials from the credential store.

You can delete or edit a saved credential.

To delete a saved credential

1. In the Management Console, click **Settings** > **Saved Credentials**.
2. Locate the name of the stored credential that you want to remove.
3. Click the **Delete** to the right of the name.

   A credential can be deleted only if it is not currently used for filers, shares, Active Directory, Fpolicy service, EMC Celerra service, permission remediation scripts, custom action scripts, Enterprise Vault server, and as Data Insight server credentials..

To edit a saved credential

1. Locate the name of the saved credential that you want to edit.
2. Click the **Edit** to the right of the name.
3. Update the user name or password.
4. If you change the password for a given credential, the new password is used for all subsequent scans that use that credential.
5. Click **Save**.

For the purpose of access control, only a user assigned the role of Server Administrator can add, edit, and view all saved credentials. A user assigned the Product Administrator role can add new saved credentials, but can only view and edit those credentials which the user has created.

Handling changes in account password

You use various account credentials at the time of configuring the Data Insight system. Accounts are used when configuring the following:

- Fpolicy Service
- Celerra Service
- Filers
- SharePoint web applications
- Scanner
- Active Directory

Perform the following steps to ensure that updates to account passwords are synchronized with the passwords used in Data Insight:

To handle changes in account password

1. Determine the places where the account is being used.
2. Log in to the Data Insight console and edit the saved credential password.

   For example, navigate to **Settings** > **Saved Credentials**, and edit the credential to update the password.
3. If the password of an account, which is used for Fpolicy service or Celerra service configuration, has changed, you must reconfigure the services as well.

   Navigate to Server details page for the corresponding nodes acting as Collectors, click the **Reconfigure Fpolicy** or **Reconfigure Celerra** sections on the page.

See “Managing saved credentials ” on page 105.

Chapter 6

Configuring directory service domains

This chapter includes the following topics:

- About directory domain scans
- Adding a directory service domain to Data Insight
- Managing directory service domains
- Fetching users and groups data from NIS+ scanner
- Configuring attributes for advanced analytics
- Deleting directory service domains
- Scheduling scans
- Configuring business unit mappings
- Importing additional attributes for users and user groups

About directory domain scans

Veritas Data Insight periodically scans the configured directory service domains in your organization to fetch information about users and user groups. Data Insight correlates this information with file and folder access logs to provide access and usage reports. This information is stored on the Management Server in the user database. Veritas recommends that you add each such domain to Data Insight whose users access file system and SharePoint resources of your organization. The time it takes to scan a directory service domain depends on the number of users and groups in the domain.

Data Insight supports the following implementations of a directory service:

- Microsoft Active Directory
- Network Information Service
- LDAP

By default, Data Insight also automatically scans local users of all Windows File Server agents, all NetApp and Celerra filers, and SharePoint site collections.

See "Adding a directory service domain to Data Insight" on page 109.

Adding a directory service domain to Data Insight

Data Insight disables addition of directory service domains if you do not have a valid license or if your license has expired. See "Overview of Data Insight licensing" on page 21.

You can configure Data Insight to scan one or more directory service domains.

To add a directory service domain to Data Insight

1. In the console, click **Settings** > **Directory Services** to display the configured directory service domains in Data Insight.
2. From the **Add New Directory Service** drop-down, select the type of directory service domain you want to add Active Directory, LDAP, NIS, or NIS+.
3. On the **Add New Directory Service** page, enter the server properties.
4. Click **Save**.
5. On the **Directory Services** listing page, click **Scan Now**.

   Once the initial scan is complete, the users and groups appear under the **Workspace** tab.

Add/Edit Active Directory options

Use this dialog box to add an Active Directory server to Data Insight, or edit the properties of an existing Active Directory server.

Table 6-1 Add/Edit Active Directory options

Field	Description
Domain Name	Enter the name of the domain which you want to scan. The domain name is used for display purpose only. The domain name that appears on the **Workspace** tab depends on the name set in the domain.
Domain Controller IP	Enter the hostname or IP address of the Active Directory domain controller.
Scanning Details	Do the following: **1** Select the saved credentials from the drop-down or specify new credentials. **2** Click **Test Credentials** to test the availability of network connection between the Management Server and the Active Directory Server, and also to verify that the credentials given are correct. Veritas recommends that you test the connection before proceeding to ensure that the Management Server is able to scan the Active Directory domain controller.
Bind Anonymously	Select the check box if you want to allow Data Insight to connect to the Active Directory server without a credential. Use anonymous binding only when connecting to an Active Directory server belonging to a trusted domain.
Disable scanning	Select the check box to disable the scanning of the directory server.

Add/Edit LDAP domain options

Use this dialog box to add a LDAP directory service server to Data Insight.

Table 6-2 Add/Edit LDAP properties options

Field	Description
Fully Qualified Domain Name	Enter the fully qualified name of the domain that you want to scan. Entering the FQDN will automatically populate the User and Group search Base DN fields.
LDAP server address	Enter the hostname and the port of the LDAP server. By default, the LDAP server runs on HTTPS port 389. If TLS is enabled, the LDAP server runs on port 636, by default.

Table 6-2 Add/Edit LDAP properties options *(continued)*

Field	Description
Type	The type of LDAP schema used by the directory service. Data Insight extracts the attributes from the schema attribute file when scanning the domain. Select one of the following: ■ OPENLDAP ■ Sun ONE You can also create a schema attribute file with customized attributes for each LDAP implementation that does not match the defaults. Ensure that you name the file as `ldap_<ldap_type>.conf` and save it at `$data\conf\ldap` on the Management Server.
Search base DN	The Organization Unit (OU) in which all users and groups have been defined.
This directory uses secure connection (TLS)	Select this check box if the LDAP server uses the TLS protocol.

Table 6-2 Add/Edit LDAP properties options *(continued)*

Field	Description
Scanning details	Select the saved credentials from the drop-down or specify new credentials. The credentials should belong to an LDAP user who has appropriate privileges to scan the LDAP domain. If you are specifying scanning credential other than the directory administrator, then make sure that you have specified the correct DN for that user. For example, uid=ldapuser,ou=People,dc=openldap,dc=com. You can connect to the LDAP database to verify the DN for an LDAP user. The example below shows the DN of a sample user, ldapuser, created on a Linux openLDAP server: uid=ldapuser,ou=People,dc=openldap,dc=com. The DN string may change depending upon the LDAP schema used. Refer to the LDAP schema to get correct DN for the user. If specifying a user other than the directory administrator, ensure that the following limits have been set appropriately on the LDAP server: ■ nsSizeLimit - Specifies maximum entries that are returned in response to a Data Insight scan. Set this attribute to -1 to return unlimited entries. ■ nsLookThroughLimit - Specifies the maximum number of Data Insight user entries checked for matches during a search operation. Set this attribute to -1 to indicate that there is no time limit. ■ limit.nsIdleTimeout - Specifies the time a Data Insight connection to the server can be idle before it is terminated. The value is given in seconds. Set this attribute to -1 to indicate that there is no time limit. The schema attribute names for setting these limits may vary depending upon the LDAP implementation. The above example is for Sun ONE.
Test Credentials	Click to verify that the given credentials are correct and to test the availability of network connection between the Management Server and the LDAP server. Veritas recommends that you test the connection before proceeding to ensure that the Management Server is able to scan the LDAP server.
Bind anonymously	Select the check box if you want to allow Data Insight to connect to the LDAP server without a credential. Select this option only if the LDAP server permits anonymous connections.

Table 6-2 Add/Edit LDAP properties options *(continued)*

Field	Description
Disable scanning	Select the check box to disable the scanning of the directory server.

Add/Edit NIS domain options

Use this dialog box to add a NIS directory service server to Data Insight.

Table 6-3 Add/Edit NIS properties

Field	Description
Fully Qualified Domain Name	Enter the name of the domain that you want to scan.
Hostname/IP address	Enter the hostname or IP address of the NIS server.
Scanning Details	Click **Test Credentials** to verify that the given credentials are correct and to test the availability of network connection between the Management Server and the NIS server. Veritas recommends that you test the connection before proceeding to ensure that the Management Server is able to scan the NIS server.
Disable scanning	Select the check box to disable the scanning of the directory server.

Add/Edit NIS+ domain options

Use this dialog box to add a NIS+ directory service server to Data Insight.

Table 6-4 Add/Edit NIS+ properties

Field	Description
Fully Qualified Domain Name	Enter the name of the domain that you want to scan.
Hostname/IP address	Enter the hostname or IP address of the NIS+ server.

Table 6-4 Add/Edit NIS+ properties *(continued)*

Field	Description
Configured in NIS compatibility mode	This check box is only available when adding a NIS+ server. When configuring a NIS+ server, select the **Configured in NIS compatibility mode** check box if the NIS+ server is configured in the NIS compatibility mode. In this mode, Data Insight can fetch the users and groups data from the NIS+ server remotely in most cases. In non NIS-compatible mode or when Data Insight cannot scan users and groups remotely, you must manually fetch the users and groups data from the NIS+ server. See "Fetching users and groups data from NIS+ scanner" on page 115.
Scanning Details	Click **Test Credentials** to verify that the given credentials are correct and to test the availability of network connection between the Management Server and the NIS+ server. Veritas recommends that you test the connection before proceeding to ensure that the Management Server is able to scan the NIS+ server.
Disable scanning	Select the check box to disable the scanning of the directory server.

Managing directory service domains

You can do the following on the Directory Services listing page:

- Add directory service domains to Data Insight.
- View details of the configured domains.
- Scan one or more domains.
- View the log of events for the configured directory service domains.

To manage the directory service domain servers

1. In the Console, click **Settings** > **Directory Services** to display the configured directory service domains in Data Insight.
2. Filter the list by using the **Domain Type** filter or by specifying the name of the domain in the search bar at the top of the page.
3. Review the following information about the configured domains:
   - The name of the domain.
   - The address of the domain server hosting the domain.
   - The type of directory service - Microsoft Active Directory, LDAP, or NIS+.

- The number of users and groups in the directory service domain.
- The additional attributes that Data Insight extracts for the domain.

4 To scan all domains, click **Scan Now**.

Note: Data Insight scans all domains together because dependencies might exist between the different domains.

---

5 To edit the scan schedule for the configured domains, click **Edit Schedule**.

By default, Data Insight scans all domains at 3:00 A.M. everyday.

On the **Set Directory Scan Schedule** dialog, change the schedule, and click **Update Schedule**.

The updated schedule is used for all subsequent scans of the configured domains.

6 To edit the properties of a directory service domain, from the **Select Actions** drop-down, select edit **Edit**.

On the directory service properties screen make the necessary changes, and click **Save**.

7 To delete a configured directory service domain, from the **Actions** drop-down, select **Delete**.

Select **OK** on the confirmation message.

8 To view events pertaining to configured directory services, click **Events**.

The events of all directory services are displayed. You can filter these events by data, domain server name, severity, and the type of event.

Fetching users and groups data from NIS+ scanner

If the NIS+ server is configured in a NIS non-compatible mode, you must manually fetch users and groups data from the NIS+ server. However, you must still add the NIS plus domain with correct information like domain name and IP address to Data Insight.

To get users and groups data

1. Log in as root to the NIS+ server.
2. Open the command prompt, and type the following commands:

To get users data	`niscat passwd.org_dir > users.txt`
To get the groups data	`niscat group.org_dir > groups.txt`

3. Save the `users.txt` and `groups.txt` files at `$data\users\nisplus\`*`example.com`* on the Management Server, where example.com is your domain name.
4. On the Directory Services listing page, click **Scan Now** to import the NIS+ domain data. You can also run the scan job from the command line by executing the following command on the Management Server:

   ```
   configcli execute_job ADScanJob
   ```

Configuring attributes for advanced analytics

Data Insight lists the custom attributes along with their display names or aliases that are configured for a domain. It also provides you the ability to edit configured attributes and add new custom attributes to Data Insight. Certain well-known custom attributes such as Email are added to Data Insight by default. The Email attribute is used to send emails for reporting and remediation.

Data Insight also uses the configured attributes for the following purposes:

- For advanced analytics, included in reports, and displayed on the console.
- To sort users in a Social Network Map cluster based on the attribute distribution.
- To group users who have activity on any given level in the hierarchy. This attribute-based grouping is displayed in the **Dashboard** view of the **Workspace** tab.
- To create attribute-based filters. These filters are used to drill down the Social Network Map to identify specific user(s) with the selected attribute.

Data Insight adds the following user attributes implicitly:

1. **LatestSID**: This attribute is used to keep a track of migration of users across domains. Its value is the SID of the user in the new domain after migration.

2. **Watchlisted**: This attribute indicates whether a user is watch-listed. Its value is 1 for the users that are watch-listed.

See "Configuring Watchlist settings" on page 57.

Data Insight adds the following group attributes implicitly:

1. **LatestSID**: This attribute is used to keep a track of migration of groups across domains. Its value is the SID of the group in the new domain after migration.
2. **groupType**: This is an Active Directory attribute that defines the type and scope of the group.

   Ensure that the names of configured attributes, including the ones fetched through CSV file, are not the same as any of the implicitly added user or group attributes. If the attributes are same, the process of updating user information in the GUI slows down.

To configure user attributes

1. In the Management Console, click **Settings** > **Directory Services**. The **Domains** page that lists all configured domains opens by default.
2. Click the **Attributes** sub-tab.

   The page lists certain well-known attributes, such as Email and Manager. You can edit these attributes, but you cannot delete them.
3. Click **Add New Custom Attribute** to add the additional attributes that you want Data Insight to extract from the directory service.
4. On the **Custom Attributes** pop-up, enter the following details:

   - The logical name for the attribute as displayed in Data Insight. You can enter a display name of the attribute.
   - From the **Type of custom attribute** drop-down, select whether the attribute applies to a user or a group.

   Data Insight extracts the following attributes from Active Directory for users :

   - displayName
   - distinguishedName
   - givenName
   - objectSid
   - sAMAccountName
   - memberOf
   - primaryGroupID

- userAccountControl
- sn

Whereas, for user groups,Data Insight extracts the following attributes from Active Directory:

- distinguishedName
- sAMAccountName
- memberOf
- objectSid

5 To edit a well-known custom attribute, click **Edit** and do the following:

- Enter the default name of the attribute as used by your directory service.
- Select the domain for which you want to specify an alias for the attribute.
- Enter the name of the attribute. An attribute can be referred to by many different names across domains. Specify the alias by which the attribute is named in the selected domain.

  The **Domain Specific LDAP Attibute Names** field supports the auto-complete feature. Data Insight provides suggestions for custom attributes when you enter part of an attribute name in the field.

  For example, the attribute name, Department, can map to the alias *Dept* in domain A and *Department* in domain B. In this example, select A and enter the *Dept* as the alias.

6 Select the **Do not fetch attribute** check box if the value for the attribute does not exist in the selected domain.

7 Click **Add New LDAP Attribute Name** to add other domain-specific names for the custom attribute.

8 Click **Save**.

For detailed information about using the Social Network Map to analyze collaborative activity, see the *Veritas Data Insight User's Guide*.

See "Choosing custom attributes for advanced analytics" on page 49.

Deleting directory service domains

You can delete a configured directory service domain.

To delete a directory service domain

1. In the Console, click **Settings** > **Directory Services** to display the configured directory service domains.
2. Click **Select Action** > **Delete** for the domain that you want to delete.
3. Click **OK** on the confirmation message.

Note: Users from a deleted directory domain are removed from Data Insight only after the next directory scan runs.

Scheduling scans

Veritas Data Insight scans configured domains everyday at 3:00 A.M., by default. You can, however, configure the scanning schedule, as needed.

Data Insight also scans local users of all file servers and site collections that are managed Data Insight. Information from these scans becomes visible in Data Insight after the directory scan runs.

See "About directory domain scans" on page 108.

See "Managing directory service domains " on page 114.

Configuring business unit mappings

Veritas Data Insight lets you associate a business unit name and business unit owner with each user imported from directory services. This information is later included in the report outputs and also sent to Symantec Data Loss Prevention as a part of ownership information.

To import business unit mappings

1. Create a .csv file, `bucsv.csv`, in the `users` folder in the Data Insight data directory on the Management Server. By default, the `users` directory on the Management Server is located at `C:\DataInsight\data\users`.

   The CSV file must contain the following information:

   - The name of the user in the format, user@domain name.
   - The name of the business unit.
   - The name of the business unit owner.

For example, *John_Doe@mycompany.com,Sales,Greg Smith*

2 The information is imported into the users database when the next Active Directory scan runs. To do so immediately, run the following command:

```
adcli.exe --mode importbu
```

Note: The domain name that is given in the .csv file must be among the domains that are scanned by Data Insight.

Importing additional attributes for users and user groups

Data Insight periodically scans directory domains and fetches basic information about users, groups, and their relationships. You can also import the additional attributes for users and user groups from other user management systems into Data Insight users database. This attribute information is later included in the report outputs and also sent to Symantec Data Loss Prevention as a part of ownership information.

To import additional attributes

1 Create and save a `.csv` file, `customattr.csv` in the `users` folder in the Data Insight data directory on the Management Server. By default, the `users` directory on the Management Server is located at `C:\DataInsight\data\users`.

Note: When you save a `.csv` file with multibyte characters, you must select UTF-8 encoding instead of unicode or default encoding.

The .csv file must contain the following information:

- Whether the entity is a user or group.
- The name of the user or group in the format, logon_name@domain name.
- The type of custom attribute, for example, Email or DN. Four types of custom attributes are defined, Integer, String, DNString, and Email.
- The name of the custom attribute.
- The value of the custom attribute. In case a custom attribute has multiple values, list each value separated by commas.

<user/group>,<logon_name@domain>,<value_type>,<attribute_name>,<value1> <value2>,...<value_n>.

For example, *user, John_Doe@mycompany.com, integer, age, 26*

2 The information that is contained in the `.csv` file is imported into the users database when the next Active Directory scan runs. To do so immediately, run the following command:

```
adcli.exe --mode importcustomattr --csvfile <location of csv file>
```

Note: To troubleshoot issues related to import of such attributes, check the log file, `adcli.log` in the `log` folder on the Management Server

---

Chapter 7

Configuring containers

This chapter includes the following topics:

- About containers
- Adding containers
- Managing containers

About containers

A container can consist of similar entities such as filers, shares, web applications, site collections, or DFS paths. Grouping the entities under a single container allows you to easily define the scope of a role assigned to a user.

For example, User1 is assigned the Product Administrator role. You can further define the scope of the role by selecting a container that contains only the filers that you want User1 to access.

Adding containers

You must add containers to Data Insight that group the filers, web applications, shares, site collections, or DFS paths, as required.

To add a new container

1. In the Console, click **Settings** > **Container**.
2. On the Containers page, click **Add new container**.
3. On the **Add New Container** screen, enter the container name.
4. From the **Container Type** drop-down, select the device type.
5. Select the **Physical** or **DFS** radio button. Based on your selection, the relevant configured resources get populated.

6 Select the resources that you want to add to the container. The selected data set is listed in the **Selected resources** pane.

7 Click **Add new container**.

Add new container/Edit container options

Use this dialog box to add a container to Veritas Data Insight or to edit the configuration of an existing container.

Table 7-1 Add new container/ Edit container options

Field	Description
Container Name	Enter a logical name for the container.
Container Type	From the drop-down, select **Filer/Web Application**, **Shares/Site Collection**, or **DFS paths**. Based on the selection, Data Insight filters the list of entities. Do the following to select the resources: **1** Select the **Physical Hierarchy** radio button to view the configured file servers or SharePoint web applications. Or select the **DFS Hierarchy** radio button to view the configured DFS paths in a domain. **2** The selected data set is listed in the **Selected resources** pane.

Managing containers

You can add containers to Data Insight, view details of the configured containers, edit the existing containers, and delete containers on the Containers listing page.

To edit the containers

1 In the Console, click **Settings** > **Containers** to display the Containers details page.

The list of configured containers appears.

2 Select the container and click **Edit**. You can change the container type and the data set details.

3 On the Edit Container screen, make the necessary configuration changes.

4 Click **Save**.

See "Add new container/Edit container options" on page 123.

To delete the containers

1. In the Console, click **Settings** > **Containers** to display the Containers details page.

   The list of configured containers appears.

2. Select the container and click **Delete**.

3. Click **OK** on the confirmation message.

Section 3

Configuring storage devices in Data Insight

- Chapter 8. Configuring NetApp file server monitoring
- Chapter 9. Configuring clustered NetApp file server monitoring
- Chapter 10. Configuring EMC Celerra or VNX monitoring
- Chapter 11. Configuring EMC Isilon monitoring
- Chapter 12. Configuring Hitachi NAS file server monitoring
- Chapter 13. Configuring Windows File Server monitoring
- Chapter 14. Configuring Veritas File System (VxFS) file server monitoring
- Chapter 15. Configuring monitoring of a generic device
- Chapter 16. Managing file servers
- Chapter 17. Configuring SharePoint monitoring
- Chapter 18. Configuring cloud content sources
- Chapter 19. Renaming storage devices

Chapter 8

Configuring NetApp file server monitoring

This chapter includes the following topics:

- About configuring NetApp file server monitoring
- Prerequisites for configuring NetApp file servers
- Credentials required for configuring NetApp filers
- Credentials required for configuring NetApp NFS filers
- Configuring SMB signing
- About FPolicy
- Preparing Data Insight for FPolicy
- Preparing the NetApp filer for Fpolicy
- Preparing the NetApp vfiler for Fpolicy
- Preparing a non-administrator domain user on the NetApp filer for Data Insight
- Enabling export of NFS shares on a NetApp file server
- Excluding volumes on a NetApp file server
- Handling NetApp home directories in Data Insight

About configuring NetApp file server monitoring

Data Insight supports the monitoring of Network Appliance (NetApp) 7-mode file servers. When you add a NetApp file server to the Data Insight configuration, it automatically discovers the CIFS and the NFS shares on the NetApp file server.

Note: Data Insight does not support scanning of NFS shares using a Collector node that is running Windows Server 2012 or Windows Server 2012 R2 edition.

Data Insight uses the FPolicy framework monitor the filer and collect access events from it.

To configure the monitoring of NetApp file servers in Data Insight, complete the following tasks:

- Complete the prerequisites for configuring the NetApp file servers.
  See "Prerequisites for configuring NetApp file servers" on page 127.
- Review the credentials that are required for configuring the file server in Data Insight.
  See "Credentials required for configuring NetApp filers" on page 128.
  See "Preparing a non-administrator domain user on the NetApp filer for Data Insight" on page 141.
- Configure the DataInsightFPolicy service on the Collector node that is configured to monitor the filer.
  See "Preparing Data Insight for FPolicy " on page 135.
- Enable FPolicy on the NetApp file server to let the Data Insight FPolicy server (Collector node) to register with the NetApp filer.
  See "Preparing the NetApp filer for Fpolicy" on page 136.
- To enable Data Insight to discover the NFS shares on the NetApp filer, enable the export of NFS shares on the filer.
  See "Enabling export of NFS shares on a NetApp file server" on page 144.
- Add the NetApp filer to the Data Insight configuration.
  See "Adding filers" on page 210.o
  See "Add/Edit NetApp filer options" on page 211.

Prerequisites for configuring NetApp file servers

Before you start configuring NetApp filers, verify the following:

- The filer is accessible from the collector node using the short name or IP address you plan to use when adding the filer.

- There is connectivity to the collector node from the filer using the short name and the Fully Qualified Host Name (FQHN) of the Collector node.
- The DNS lookup and reverse-lookup for hostname of the Collector node from the filer is working fine.
- The http port 80 and standard RPC ports are open in the firewall.
- The local security policies are set. The installer automatically registers the local security policies on Windows 2008 machines which are used as collector nodes. However, if the installer fails to register the security policies, you must set them manually. Click **Administrative Tools** > **Local Security Policy** > **Local Policies** > **Security Options** and change the following settings:
  - Network access: Named Pipes that can be accessed anonymously - Add NTAPFPRQ to the list.

  You must restart the machine after making these changes.

Credentials required for configuring NetApp filers

Table 8-1 lists the set of credentials that are required by Veritas Data Insight during system configuration.

Table 8-1 Credentials for configuring NetApp filers

Credential	Details
Credentials required to configure DataInsightFpolicy service. The DataInsightFpolicy service runs on the Collector and processes the events that are sent by the NetApp filer using the FPolicy RPC interface. This service must be configured on each Collector that is used to connect to NetApp filers. DataInsightFpolicy service also performs miscellaneous tasks, such as gathering storage information from the filer.	The credential should belong to a user in the domain of which the Data Insight Collector node and the NetApp filers are a part. This user should be a part of the Backup Operators group on the filer. The credential should belong to the same domain as that of the Collector node assigned to the filer. However, when you add filers, the account you specify during filer configuration must be a domain user and must have Backup Operator privileges on the filer. **Note:** LocalSystem account is not supported.

Table 8-1 Credentials for configuring NetApp filers *(continued)*

Credential	Details
Credentials required during filer configuration through the Data Insight Management Console.	Required to discover shares and enable FPolicy on the NetApp filer. This credential belongs to the NetApp ONTAP user who has administrative rights on the NetApp filer (for example, root) or a domain user who is part of the Administrators group on the filer. Or, this credential belongs to the NetApp ONTAP user or a domain user who is a non-administrator user on the filer, but has specific privileges. See "Preparing a non-administrator domain user on the NetApp filer for Data Insight" on page 141. **Note:** The domain user can be the same user account that was specified when configuring the DataInsightFpolicy service. If you use the Local System account to configure the DataInsightFpolicy service on the Collector, the user you specify here must also belong to the Backup Operators group on the filer. See "Preparing Data Insight for FPolicy " on page 135.

Table 8-1 Credentials for configuring NetApp filers *(continued)*

Credential	Details
Credentials required for scanning of shares.	

Table 8-1 Credentials for configuring NetApp filers *(continued)*

Credential	Details
	Required for scanning of shares from the NetApp filer. When scanning CIFS shares, this credential belongs to the user in the domain of which the NetApp filer is a part. This user must belong to either the Power Users or Administrator's group on the NetApp filer. If the credential is not part of one of these groups, the scanner will not be able to get share-level ACLs for shares of this filer. You do not need this privilege if you do not want to get the share-level ACLs. In this case you will only need privileges to mount the share and scan the file system hierarchy. You must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs: ■ Traverse Folder/Execute File ■ List Folder/Read Data ■ Read Attributes ■ Read Extended Attributes ■ Read Permissions For scanning NFS shares, Data Insight needs a UNIX account with at least read and execute permissions on all folders, along with at least Read permission on all files. By default, Data Insight uses User ID or Group ID 0 to scan NFS shares. You can configure an alternate User ID or Group ID from the **Settings** > **Advanced Settings** section of the Collector node. See "Configuring advanced settings" on page 78. When monitoring only NFS shares, you can specify **Use Local System account** from the scanning credentials drop-down, else you can specify credentials required to scan CIFS shares.

Table 8-1 Credentials for configuring NetApp filers *(continued)*

Credential	Details
	Note: Veritas recommends that the credentials used to discover shares on the NetApp file server must be the same as the credentials used to configure DataInsightFpolicy the service.

Credentials required for configuring NetApp NFS filers

Table 8-2 lists the set of credentials that are required by Veritas Data Insight during system configuration.

Table 8-2 Credentials for configuring NetApp NFS filers

Credential	Details
Credentials required to configure DataInsightFpolicy service. The DataInsightFpolicy service runs on the Collector and processes the events that are sent by the NetApp filer using the Fpolicy RPC interface. This service must be configured on each Collector that is used to connect to NetApp filers. DataInsightFpolicy service also performs miscellaneous tasks, such as gathering storage information from the filer.	The credential should belong to a user in the domain of which the Data Insight Collector node and the NetApp filers are a part. The credential should also be part of the local Backup Operators group of each of the NetApp filer being monitored by the Collector assigned to the filer. **Note:** LocalSystem account is not supported.

Table 8-2 Credentials for configuring NetApp NFS filers *(continued)*

Credential	Details
Credentials required during filer configuration through the Data Insight Management Console.	Required to discover shares and enable Fpolicy on the NetApp filer. This credential belongs to the NetApp ONTAP user who has administrative rights on the NetApp filer (for example, root) or a domain user who is part of the Administrators group on the filer. Or, this credential belongs to the NetApp ONTAP user or a domain user who is a non-administrator user on the filer, but has specific privileges. See "Preparing a non-administrator domain user on the NetApp filer for Data Insight" on page 141. See "Preparing Data Insight for FPolicy " on page 135.
Credentials required for scanning of shares.	Required for scanning of shares from the NetApp filer. This credential belongs to the user in the domain of which the NetApp filer and the Data Insight Collector node are a part. This user can also be the same domain user who is authorized to run the DataInsightFpolicy service. You can also use the LocalSystem account for scanning. Additionally, you must also have share-level READ permissions on the NFS export. **Note:** Veritas recommends that the credentials used to discover shares on the NetApp file server should be the same as the credentials used to configure DataInsightFpolicy the service.

Configuring SMB signing

Ensure that Server Message Block (SMB) signing is either turned on or turned off on both the Collector node and the NetApp filer. If SMB signing is turned on, all

packets of data that is sent over a network to a remote host are signed. A mismatch in the setting on the Collector node and the NetApp filer can cause the filer to drop the FPolicy connection to the Collector node.

To configure SMB signing

1. Check whether the SMB signing option on the NetApp filer, `options cifs.signing.enable` is set to off or on.
2. On the Collector node that is assigned to the NetApp filer, open the Windows' Registry Editor (**Start** > **Run** > **regedit**).
3. In Registry Editor, navigate to **HKEY_LOCAL_MACHINE** > **SYSTEM** > **CurrentControlSet** >**SERVICES** > **LanmanServer** > **Parameters**.
4. Modify the following registry entries:
   - `enablesecuritysignature` - Enter the value 0 to turn signing off and enter the value 1 to turn on signing.
   - `requiredsecuritysignature` - Enter the value 0 to turn signing off and enter the value 1 to turn on signing.

About FPolicy

Veritas Data Insight uses the FPolicy framework provided by NetApp to collect access events from the NetApp filers.

NetApp provides an interface called FPolicy which allows external applications to receive file access notifications from the NetApp Storage subsystem. FPolicy allows partner applications to perform tasks like file access screening and auditing. The FPolicy interface uses Remote Procedure Calls (RPC) and external applications can use these tools to register with the NetApp Filer as FPolicy servers. FPolicy supports both CIFS and NFS.

The unit of FPolicy configuration on the NetApp filer is called a policy, which is identified by a user specified name. You can configure a policy to monitor all or a list of volumes on the NetApp filer along with a specified set of operations. The monitored operations are open, close, read, write, create, delete, rename, and set attribute. As soon as a file operation is performed on a file or folder on the filer which is being monitored, a notification is sent to the registered FPolicy server asynchronously.

Note: The policy created by Veritas Data Insight should not be shared by any other applications or clients.

By default, Data Insight does not register for read and close events from NetApp filers. Data Insight treats an open event as a read event. This behavior reduces the load on the filer in case of peak traffic loads from third party applications like backups over CIFS. It also does not have an adverse effect for most consumer applications because consumer applications seldom write to a file before first reading it. Data Insight assumes that an open event is almost always be followed by a read event and then optionally by a write event. However, you can customize the default behavior as per your requirements.

See "Enabling export of NFS shares on a NetApp file server" on page 144.

See "Preparing the NetApp filer for Fpolicy" on page 136.

Preparing Data Insight for FPolicy

The Data Insight FPolicy server can reside on the Management Server and/or on each Collector worker node. The Management Server and/or the Collector worker node must register with the NetApp filer to receive audit information. Before you assign a Data Insight server as a collector for a NetApp filer, you must configure the FPolicy service on that server.

To set up the environment for Data Insight FPolicy service

1. Provision a Windows 2003 or 2008 server in the same Active Directory domain as the filers you want to monitor using FPolicy. This computer hosts the FPolicy server. Ensure that the filer and the FPolicy server are part of the same domain.
2. Install the Data Insight Collector worker node or the Data Insight Management Server on this server.
3. Logon to the Data Insight Management Console.
4. In the Console, click **Settings** > **Data Insight Servers** to open the listing page for the server.
5. Select the server from the server list to open the details page for the server.
6. Click the **Services** tab.
7. Select **DataInsightFPolicy** to expand the section.
8. In the **Policy Name** field, enter the policy name that will be enabled on each filer, of this node Collector. The default name is *matpol*.

9 Under **Credentials**, select the saved credentials that the service needs to run as. The credentials must belong to the Backup Operators group on the NetApp filer that is being monitored by the Collector.

See "Credentials required for configuring NetApp filers" on page 128.

10 Click **Configure** to apply these settings to the server and start the FPolicy service.

See "Configuring SMB signing" on page 133.

See "About FPolicy" on page 134.

Preparing the NetApp filer for Fpolicy

The Veritas Data Insight Fpolicy server registers with the NetApp filer and receives file access events from it. Fpolicy must be enabled and configured on that NetApp filer. Veritas recommends that you automatically configure FPolicy on the filer to enable auditing when adding filers.

See "Adding filers" on page 210.

However if you want more control on the shares you want to monitor use the manual steps. The manual steps are valid for NetApp ONTAP version 7.0 and higher.

Note: The steps below assume that the name of the policy is *matpol*.

To configure the Fpolicy on the NetApp filer using manual steps

1 Launch a Telnet session with the filer and run the following commands, as appropriate:

- To create a policy:

```
fpolicy create matpol screen
```

- To enable a policy:

```
fpolicy enable matpol -f
```

2 Use the following optional commands for monitoring:

- To set the Fpolicy for CIFS to monitor specific events:

```
fpolicy mon add matpol -p cifs -f write,
open,close,delete,rename,create
```

- To set the Fpolicy for NFS to monitor specific events:

```
fpolicy mon add matpol -p nfs -f create,delete,rename,write,
open,link,symlink,setattr
```

- To monitor specific events on NetApp filer versions 7.3 or higher:
  - Enable set attributes operation:

    ```
    fpolicy mon options cifs_setattr on - add mon
    ```

    ```
    fpolicy mon options nfs_setattr on - add
    ```

  - Add events to be monitored:

    ```
    fpolicy mon add matpol -p cifs -f write,
    open,close,delete,rename,create,setattr
    ```

    ```
    fpolicy mon add matpol -p nfs -f create,delete,rename,write,
    open,link,symlink,setattr
    ```

- To see details of a configured policy:

  ```
  fpolicy show matpol
  ```

- To disable monitoring of specific events:

  ```
  fpolicy mon remove matpol -p cifs -f write,
  open,close,delete,rename,create
  ```

  ```
  fpolicy mon remove matpol -p nfs -f create,delete,rename,write,
  open,link,symlink,setattr
  ```

- To disable use of a policy:

  ```
  fpolicy disable matpol
  ```

- To delete a policy:

  ```
  fpolicy destroy matpol
  ```

3 To add a domain user to the administrator's group:

```
useradmin domainuser add domain-username
-g Administrators
```

Note: The domain user is the user who is configured to run the Fpolicy service on the collector.

To configure a non-administrator user:

See "Preparing a non-administrator domain user on the NetApp filer for Data Insight" on page 141.

4 To display a list of users who are already configured:

```
useradmin domainuser list -g Administrators
```

A list with the SIDs of the configured domain users appears. To resolve the SIDs, run the following command:

```
cifs lookup SID
```

See "Configuring SMB signing" on page 133.

Preparing the NetApp vfiler for Fpolicy

The Data Insight Fpolicy server can register with the NetApp vfiler and receive file access events from it. Fpolicy has to be enabled and configured on that NetApp vfiler manually.

To configure the Fpolicy on the NetApp vfiler using manual steps

1 Launch a Telnet session with the filer and run the following commands, as appropriate:

- To get the vfiler name:

  ```
  vfiler status
  ```

  Choose the name of the vfiler that you want to configure and then perform the following operations for that vfiler. Ignore the name, *vfiler0*, which is the default name given to the physical filer by NetApp.

Note: Consult your system administrator to get the IP address of the vfiler. You will need this IP address while adding the vfiler from the Management Console.

See "Adding filers" on page 210.

- To create a policy:

```
vfiler run vfilername fpolicy create matpol screen
```

- To enable a policy:

```
vfiler run vfilername fpolicy enable matpol -f
```

2 Use the following optional commands for monitoring:

- To set the Fpolicy for CIFS to monitor specific events:

```
vfiler run vfilername fpolicy mon add matpol -p cifs
 -f write,open,close,delete,rename,create
```

To set Fpolicy for NFS to monitor specific events:

```
vfiler run vfilername fpolicy mon add matpol -p nfs -f create,
delete,rename,write,open,link,symlink,setattr
```

- To set the Fpolicy for CIFS to monitor specific events on NetApp filer versions 7.3 or higher:
  - Enable set attributes operation:

```
vfiler run vfilername fpolicy options cifs_setattr on
```

```
vfiler run vfilername fpolicy options nfs_setattr on
```

  - Add events to be monitored:

```
vfiler run vfilername fpolicy mon add matpol -p cifs
-f write,open,close,delete,rename,create,setattr
```

```
vfiler run vfilername fpolicy mon add matpol -p nfs -f create,
delete,rename,write,open,link,symlink,setattr
```

- To see details of a configured policy:

```
vfiler run vfilername fpolicy show matpol
```

- To disable monitoring of specific events:

```
vfiler run vfilername fpolicy mon remove matpol -p cifs
-f write,open,close,delete,rename,create
```

```
vfiler run vfilername fpolicy mon remove matpol -p nfs -f create,
delete,rename,write,open,link,symlink,setattr
```

- To disable use of a policy:

```
vfiler run vfilername fpolicy disable matpol
```

- To delete a policy:

```
vfiler run vfilername fpolicy destroy matpol
```

where, *vfilername* is the name of the vfiler you want to configure.

3 To add a domain user to the administrator's group:

```
vfiler run vfilername useradmin domainuser
add domain-username -g Administrators
```

Note: The domain user is the user who is configured to run the Fpolicy service on the collector. See "Preparing the NetApp filer for Fpolicy" on page 136.

To configure a non-administrator user:

See "Preparing the NetApp filer for Fpolicy" on page 136.

4 To display a list of users who are already configured:

```
vfiler run vfilername useradmin domainuser list
-g Administrators
```

A list with the SIDs of the configured domain users appears. To resolve the SIDs, run the following command:

```
cifs lookup SID
```

Preparing a non-administrator domain user on the NetApp filer for Data Insight

To configure a NetApp filer from the Management Console, you can use an account which is not in the administrators group on the NetApp filer, but has some specific privileges.

Perform the following steps on the NetApp filer console to add a non-administrator user, for example, *testuser*.

To create a non-administrator user

1. Create a new role, for example *testrole*, using the `useradmin` utility on the filer.
2. Add the login-* and API-* capabilities to the role.

   For example, `useradmin role add testrole -a login-*,api-*`.

   You can also choose to assign specific capabilities to the role.

   Table 8-3 provides a detailed description of each capability.
3. Create a new group, for example, *testgroup* and apply the role *testrole* to it.

   For example, `useradmin group add testgroup -r testrole`.
4. Add the user *testdomain\testuser* to *testgroup*

   For example, `useradmin domainuser add testdomain\testuser -g testgroup`.
5. Add the user *testdomain\testuser* to *Backup Operators* group.

   For example, `useradmin domainuser add testdomain\testuser -g Backup Operators`.

Note: For vfilers, append the above command-line examples with `vfiler run <vfilername>`.

Table 8-3 Additional capabilities for adding a non-administrator user account

Capability	Description
`login-http-admin`	Enables you to log into the NetApp filer and run commands. With this capability, you can get latency statistics (for scan throttling), volume size information, or discover shares.

Table 8-3 Additional capabilities for adding a non-administrator user account *(continued)*

Capability	Description

`api-system-get-ontapi-version`
 	
`api-system-get-version`
	Enables you to get the ONTAPI version number and the system version number respectively. These are required to set the login handle context properly. Data Insight reports a failure when you test the connection to the filer, if these capabilities are absent. Also, if these capabilities are absent, you cannot execute any APIs including those required to discover shares, and get latency statistics.

`api-fpolicy-set-policy-options`	Enables you to set a flag on the NetApp filer to enable ACL change notifications. If you choose not to supply this capability to Veritas, the filer administrator must set this property manually using the root telnet console (`fpolicy options <policyname> cifs_setattr on`).

`api-fpolicy-list-info`
 	
`api-fpolicy-create-policy`
 	
`api-fpolicy-enable-policy`
 	
`api-fpolicy-disable-policy`
 	
`api-fpolicy-destroy-policy`
	Required to enable Data Insight to automatically create and enable FPolicy on the NetApp filer. Optionally, the filer administrator can set up the policy using the root telnet console.
	
	Data Insight reports a failure when you test the connection to the filer, if these capabilities are absent. However, you can save a filer with the policy that is already created and enabled by the filer administrator without testing the connection to the filer.

`api-fpolicy-server-list-info`
 	
`api-fpolicy-list-info`
	Used to retrieve useful statistics from the NetApp filer, such as, total event count and event failures. These APIs are used every two hours so they do not load the system. However, absence of these capabilities does not cause any problems.

`api-license-list-info`	Used to check if this NetApp filer has CIFS license and print the appropriate diagnostic message. Data Insight reports a failure when you create or enable FPolicy on the filer or when you test the connection to the filer, if these capabilities are absent. However, you can save a filer with the policy that is already created and enabled by the filer administrator without testing the connection to the filer.

Table 8-3 Additional capabilities for adding a non-administrator user account *(continued)*

Capability	Description
`api-options-set`	Used to enable the global FPolicy flag on the NetApp filer. In the absence of this capability, FPolicy creation fails. Also, Data Insight reports a failure when you test the connection to the filer, if these capabilities are absent. However, you can save a filer with the policy that is already created and enabled by the filer administrator without testing the connection to the filer.

`api-cifs-share-list-iter-start`
`api-cifs-share-list-iter-next`
`api-cifs-share-list-iter-end`
	Used to discover shares on the NetApp filer. Absence of these capabilities can result in a failure to discover the shares. Optionally, you can add shares manually from the Data Insight console.

`api-perf-object-get-instances-iter-start`
`api-perf-object-get-instances-iter-next`
`api-perf-object-get-instances-iter-end`
	Used to get CIFS latency information from the NetApp filer, which enables the self-throttling scan feature of Data Insight. Absence of these APIs can cause scanner to fail if you enable the feature to throttle scanning.

`api-volume-list-info`	Used to periodically fetch size information for NetApp volumes.
`api-nfs-exportnfs-list-rules` `api-nfs-exportnfs-list-rules-2`	Used to discover all NFS shares that are exported from the NetApp filer. If this capability is absent, these NFS shares are not discovered.
`api-net-ping` `api-net-resolve`	Used to check network connectivity from the filer to the Data Insight Collector. These APIs are useful to run some diagnostic checks on the filer. However, such checks can also be done manually by the NetApp administrator, and hence these APIs are not mandatory.
`api-fpolicy-volume-list-set`	Used to set the volume names on the filer which are to be excluded from being monitored by FPolicy. If this capability is absent, you cannot exclude volumes from being monitored by FPolicy from the Data Insight console . However, you can exclude volumes manually using the CLI on the NetApp console. See "Excluding volumes on a NetApp file server" on page 145.

Table 8-3 Additional capabilities for adding a non-administrator user account *(continued)*

Capability	Description
`api-system-get-info`	Used to discover the NetApp system serial number. The information is used by external reporting tools like Veritas Operations Manager (VOM) to report about the NetApp filers that Data Insight monitors. This privilege is mandatory if VOM integration is required.

Enabling export of NFS shares on a NetApp file server

Before you add a NetApp filer to Data Insight, you must enable the export of NFS shares on the NetApp filer to allow Data Insight to discover the NFS shares on the filer.

Note: Data Insight does not support scanning of NFS shares using a Collector node that is running Windows Server 2012 or Windows Server 2012 R2 edition.

To enable export of NFS shares on the NetApp filer

1. On the NetApp filer web console, select **NFS** > **Manage exports**.
2. On the Export wizard, click **Add Export** or you can edit the existing exports to modify them.
3. On the first page of the wizard ensure that you have at least selected read only and root access, Other options can also be specified, as required, and click **Next.**
4. Define the export path and give read only access to the Data Insight Collector node, and click **Next**.
5. On the Read-Write Access page, enable read-write access for all clients or for specific hosts, as per your need.
6. Click **Next**.
7. On the Root Access page, define root access to the Data Insight Collector node, and click **Next**.
8. On the Security page, accept the default options, and click **Next**.
9. On the Summary page, review the configuration and click **Commit** to save the changes.

See “Adding filers” on page 210.

See “Add/Edit NetApp filer options” on page 211.

Excluding volumes on a NetApp file server

To optimize the performance of a NetApp file server. you can choose to exclude a set of volumes from being monitored by FPolicy using NetApp commands.

If a configured NetApp filer hosts applications that create a huge number of access events on certain volumes, it may cause a high I/O latency for CIFS. Veritas recommends that you exclude such volumes from being monitored by FPolicy.

To exclude volumes on the NetApp filer

1. Log in to the Net App file server.
2. Execute the following command:

   ```
   fpolicy vol exc add <name of policy> <comma separated name list
   of volumes to be excluded>
   ```

   Note that the list of volumes is not the path of the volumes, but names of the volumes.

Handling NetApp home directories in Data Insight

Data ONTAP enables you to create users’ home directories on the NetApp storage system. You can access your home directory share the same way you can access any other share on the file system. You can connect only to your home directory matching your user name without seeing other users’ home directories. The user name for matching a home directory can be a Windows user name, a domain name followed by a Windows user name, or a UNIX user name. Refer the NetApp documentation for more details about how Data ONTAP matches a home directory with a user,

Data Insight cannot parse FPolicy events coming from home directory shares because of insufficient information regarding mapping to existing shares. To monitor events on home directory shares, you must configure an artificial share in Data Insight. This share represents all home directories on a NetApp filer, and events from the users' home directories are directed to this share. Note that Data Insight does not scan the artificial share.

To configure an artificial share in Data Insight

1. Log in toData Insight Management Console.
2. Click **Settings** > **Filers**.

3. Click the filer for which you want to add a share.
4. On the filer details page, click **Monitored Shares**.
5. On the **Monitored Shares** page, click **Add New Share**.
6. On the New Share pop-up, enter the following details:
   - The name of the share that you want to add - `CIFS.HOMEDIR`.
   - The physical path of the share on the filer -`/CIFS.HOMEDIR`..
7. Select the **Define custom cron schedule** radio button, and select **Never** from the schedule drop-down.
8. To verify that the share is added, on the Management Console, navigate to **Settings** > **Filers**.
9. Click the relevant filer. On the filer details page, click the **Monitored Shares** tab to review the list of configured shares.
10. Repeat the steps for each NetApp filer for which you want to add an artificial share.

Chapter 9

Configuring clustered NetApp file server monitoring

This chapter includes the following topics:

- About configuring a clustered NetApp file server
- About configuring FPolicy in Cluster-Mode
- Pre-requisites for configuring clustered NetApp file servers
- Credentials required for configuring a clustered NetApp file server
- Preparing a non-administrator local user on the clustered NetApp filer
- Preparing a non-administrator domain user on a NetApp cluster for Data Insight
- Preparing Data Insight for FPolicy in NetApp Cluster-Mode
- Preparing the ONTAP cluster for FPolicy
- About configuring secure communication between Data Insight and cluster-mode NetApp devices
- Enabling export of NFS shares on a NetApp Cluster-Mode file server

About configuring a clustered NetApp file server

Veritas Data Insight supports the monitoring of CIFS shares on clustered NetApp 8.2.1 or later NAS devices along with standalone NetApp file servers.

On Clustered Data ONTAP (cDOT) versions 8.2.3 and 8.3.1, Data Insight also supports monitoring of NFS exports.

In Data ONTAP Cluster-Mode (C-mode), a Storage Virtual Machine (SVM) is a logical unit within an ONTAP cluster which can contain a CIFS server and an SVM with NFS protocol enabled. An SVM facilitates data access within a cluster. It contains data volumes and one or more LIFs through which they serve data to the clients. For the purpose of monitoring, a CIFS server within an SVM represents a filer for Data Insight. Every SVM can have multiple physical nodes. Data Insight uses the FPolicy interface on ONTAP to receive event notifications, which are sent in XML format over a TCP connection to the Data Insight's FPolicy server.

Data Insight automatically discovers the following entities in the cluster environment:

- The SVMs and the CIFS server underneath each SVM.
- The CIFS shares for each CIFS server.
- The NFS exports on the NetApp cluster.

To enable Data Insight to discover shares in the cluster, you must ensure the following:

- Complete all the pre-requisites.
  See "Pre-requisites for configuring clustered NetApp file servers" on page 149.
- Provide correct cluster user credentials when adding the clustered file server. You can also use the credentials of a local cluster administrator user.
  See "Credentials required for configuring a clustered NetApp file server" on page 151.
- You can also choose to configure the NetApp cluster in Data Insight by using SSL certificates.
  See "About configuring secure communication between Data Insight and cluster-mode NetApp devices" on page 159.
- Configure the DataInsightFPolicyCmod service on the Collector that is configured to monitor the filer.
  See "Preparing Data Insight for FPolicy in NetApp Cluster-Mode" on page 157.
- After you add a CIFS server or an SVM with NFS protocol enabled from the ONTAP cluster to Data Insight, Data Insight automatically enables FPolicy on the corresponding SVM on the ONTAP cluster. This operation helps the ONTAP cluster register with the Data Insight FPolicy server. Note that in ONTAP 8.2 Cluster-Mode, the connection is initiated by ONTAP.
  See "Preparing the ONTAP cluster for FPolicy" on page 158.
- Add the clustered NetApp file server to the Data Insight configuration.
- See "Adding filers" on page 210.

See "Add/Edit NetApp cluster file server options" on page 215.

Once a TCP connection is established with the Data Insight FPolicy server (Collector node), it starts receiving access event information from the ONTAP cluster.

About configuring FPolicy in Cluster-Mode

FPolicy is a file access notification framework which allows screening of NFS and CIFS accesses. See "About FPolicy" on page 134.

In case of a Cluster-Mode configuration, FPolicy requires that all the nodes in the NetApp cluster are running Data ONTAP 8.2 or later.

To receive access events from the ONTAP cluster, you must first install the DataInsightFPolicyCMod service on the Data Insight Collector node that is configured to monitor the NetApp cluster. The service communicates with the NetApp cluster and fetches information about the file operations that Data Insight monitors. Data Insight automatically configures the list of operations to monitor when you add the NetApp cluster to Data Insight. However, you can also choose to configure the operations to be monitored manually.

Note: The FPolicy server for a NetApp standalone (7-mode) configuration and C-Mode configuration can co-exist on the same Data Insight Collector node.

For detailed information about the DataInsightFPolicyCmod service, see the *Veritas Data Insight Installation Guide*.

See "Preparing Data Insight for FPolicy in NetApp Cluster-Mode" on page 157.

See "Preparing the ONTAP cluster for FPolicy" on page 158.

Pre-requisites for configuring clustered NetApp file servers

Before you can start using Data Insight to monitor the NetApp file servers operating in Cluster-Mode, verify the following:

- ONTAP version 8.2.1 or higher cluster is configured in accordance with NetApp documentation. Note that Data Insight supports the monitoring of NFS exports on Clustered Data ONTAP (cDOT) versions 8.2.3 and 8.3.1 or higher.

  Note: Data Insight does not support monitoring of NFS exports on Clustered Data ONTAP (cDOT) version 8.3.

- Veritas recommends that Windows 2008 or Windows 2008 R2 server is used as the Collector node, if you want to NFS exports on NetApp Cluster-Mode.
- The DataInsightFPolicyCmod service is installed on the Data Insight Collector that is configured to monitor the NetApp filers operating in Cluster-Mode.
  The Data Insight FPolicy server for NetApp cluster runs on a standard recommended computer for the Data Insight Collector node. However, a Windows 2008 or Windows 2012 64bit server with 8GB RAM and quad core processor is preferred.
- To discover SVMs with NFS protocol enabled, ensure that:
  - An NFS client is installed on the Collector node.
  - The QoS setting configured on the SVM.
- FPolicy is configured for every Storage Virtual Machine (SVM) in the ONTAP cluster.
  See "Preparing the ONTAP cluster for FPolicy" on page 158.
- The ONTAP Cluster Management host is accessible from the Collector using the short name or IP address.
- Data Insight Collector should be able to communicate with port 80 on the ONTAP cluster management host.
- Data Insight should be able to communicate with the CIFS server hosted within the ONTAP cluster and the SVM which has NFS protocol enabled.
- Run the following command on each VServer where you want to enable NFS monitoring:
  ```
  Vserver nfs modify -vserver <vserver name> -v3-ms-dos-client
  enabled
  ```
  See "Enabling export of NFS shares on a NetApp Cluster-Mode file server" on page 165.

See "About configuring a clustered NetApp file server" on page 147.

Credentials required for configuring a clustered NetApp file server

Table 9-1 Credentials for configuring NetApp file servers in Cluster-Mode

Credential	Details
Credentials required during filer configuration through the Data Insight Management Console. Data Insight uses these credentials to discover shares and enable FPolicy on the NetApp Cluster-Mode filer.	This credential belongs to the NetApp ONTAP cluster administrator user who is a local user on the ONTAP cluster. Or, this credential belongs to the ONTAP cluster non-administrator user with specific privileges. See "Preparing a non-administrator local user on the clustered NetApp filer " on page 153.

Table 9-1 Credentials for configuring NetApp file servers in Cluster-Mode *(continued)*

Credential	Details

Credentials required for scanning of shares.	Required for scanning of shares from the NetApp filer.
	
	When scanning CIFS shares, this credential belongs to the user in the domain of which the NetApp filer is a part. This user must belong to either the Power Users or Administrator's group on the NetApp filer. If the credential is not part of one of these groups, the scanner is not able to get share-level ACLs for shares/exports of this filer.
	
	Run the following commands on every Storage Virtual Machine (SVM) that Data Insight is monitoring:
	
	`cifs users-and-groups local-group add-members -vserver <vserver name > BUILTIN\Administrators -member-names <domain name \usename>`
	
	`cifs users-and-groups local-group add-members -vserver <vserver name > BUILTIN\Power Users -member-names <domain name \usename>`
	
	You do not need this privilege if you do not want to get the share-level ACLs. In this case you only need privileges to mount the share and scan the file system hierarchy.
	
	You must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs:
	
	■ Traverse Folder/Execute File
	■ List Folder/Read Data
	■ Read Attributes
	■ Read Extended Attributes
	■ Read Permissions

Preparing a non-administrator local user on the clustered NetApp filer

To configure a NetApp cluster file server from the Data Insight Management Console, you can use a local user account which is not in the administrators group on the NetApp cluster, but has some specific privileges.

To create a non-administrator user

1. Launch a Telnet session with the NetApp Cluster Management host.
2. Create a new role, for example *testrole*, using the `useradmin` utility on the filer.
3. Run the following commands to create the role with specific privileges:

```
security login role create -role testrole -cmddirname "version"
-access all
```

```
security login role create -role testrole -cmddirname "vserver cifs"
-access readonly
```

4 `security login role create -role` *`testrole`* `-cmddirname "vserver" -access readonly`

`security login role create -role` *`testrole`* `-cmddirname "vserver cifs share" -access readonly`

`security login role create -role` *`testrole`* `-cmddirname "vserver nfs" -access all`

`security login role create -role` *`testrole`* `-cmddirname "volume" -access readonly`

`security login role create -role` *`testrole`* `-cmddirname "vserver fpolicy policy" -access all`

`security login role create -role` *`testrole`* `-cmddirname "vserver fpolicy" -access readonly`

`security login role create -role` *`testrole`* `-cmddirname "vserver fpolicy enable" -access all`

`security login role create -role` *`testrole`* `-cmddirname "vserver fpolicy disable" -access all`

`security login role create -role` *`testrole`* `-cmddirname "statistics" -access readonly`

`security login role create -role` *`testrole`* `-cmddirname "network interface show" -access readonly`

The `network interface show` privilege automatically assigns the following privileges to the role:

`network interface create`

`network interface delete`

`network interface modify`

5 Run the following command to create a local user, for example, *testuser*, and assign the role that you created in 3 to the user:

```
security login create -username testuser
-application ontapi -authmethod password -role testrole
```

Preparing a non-administrator domain user on a NetApp cluster for Data Insight

To configure a NetApp cluster from the Management Console, you can use an account which is not in the administrators group on the NetApp filer, but has some specific privileges. You can use the credentials of a domain user to configure Data Insight to monitor a NetApp cluster. These credentials are required to discover shares and to enable FPolicy on the NetApp cluster.

Perform the following steps on the NetApp filer console to add a non-administrator user, for example, *testuser*.

To use domain user credentials to configure a NetApp cluster

1. Log on using SSH to the NetApp cluster with administrator credentials. Do one of the following:

   - If the NetApp cluster has a data SVM with a CIFS server that is already created, you can use that data SVM as an authentication tunnel. Use the following command:

     ```
     security login domain-tunnel create -vserver name of data SVM
     ```

     The following security command displays the specified authentication tunnel:

     ```
     login domain-tunnel show
     ```

   - If the cluster does not have a data Storage Virtual Machine (SVM) with a CIFS server created, you can use any data SVM in the cluster and join it to a domain by using the `vserver active-directory create` command. Set the `--vserver` parameter to the data SVM. Joining a data SVM to a domain does not create a CIFS server or require a CIFS license. However, it enables the authentication of users and groups at the SVM or cluster-level.

2. Grant a user or a group access to the cluster or SVM with the `-authmethod` parameter set to *domain*.

   Also, create a new role, for example *testrole*, using the `useradmin` utility on the filer.

   The following command enables <*testuser*> in the <*DOMAIN1*> domain to access the cluster through SSH:

   ```
   cluster1::> security login create -vserver cluster1 <-user-or-group-name
   <DOMAIN1\testuser> -application ontapi -authmethod domain -role testrole
   ```

   Where, <*cluster1*> is the name of Admin Vserver.

Note the following:

- The value of the `-user` and `-group-name` parameter must be specified in the format *domainname\username*, where <*domain name*> is the name of the CIFS domain server and user name is the name of the user or group you want to grant access to.
- The user group authentication supports only SSH and ONTAPI for the `-application` parameter.
- If the authentication tunnel is deleted, the directory service logon sessions cannot be authenticated by the cluster, and users and groups cannot access the cluster. The open sessions that were authenticated before the deletion of the authentication tunnel remain unaffected.

3 You can also choose to assign specific capabilities to the role.

```
security login role create -role testrole -cmddirname
"version"-access all
```

Enables you to log into the NetApp filer and run commands. With this capability, you can get latency statistics (for scan throttling), volume size information, or discover shares.

Run the following commands to create the role with specific privileges:

```
security login role create -role testrole -cmddirname
"version"-access all
```

```
security login role create -role testrole -cmddirname "vserver
cifs"-access readonly
```

```
security login role create -role testrole -cmddirname "vserver
nfs" -access all
```

```
security login role create -role testrole -cmddirname
"vserver"-access readonly
```

```
security login role create -role testrole -cmddirname
"volume"-access readonly
```

```
security login role create -role testrole –cmddirname "vserver
fpolicy policy" -access all
```

```
security login role create -role testrole –cmddirname "vserver
fpolicy" -access readonly
```

```
security login role create -role testrole –cmddirname "vserver
fpolicy enable" -access all
```

```
security login role create -role testrole –cmddirname "vserver
fpolicy disable" -access all
```

```
security login role create -role testrole –cmddirname "statistics"
-access readonly
```

```
security login role create -role testorle -cmddirname "network
interface show" -access readonly
```

The `network interface show` privilege automatically assigns the following privileges to the role:

- `network interface create`
- `network interface modify`
- `network interface modify`

You can optionally specify a default role such as *admin/vsadmin* which already has these privileges.

Preparing Data Insight for FPolicy in NetApp Cluster-Mode

The Veritas Data Insight FPolicy server is represented by DataInsightFpolicyCmod service which runs on each Collector worker node. Before you assign a Data Insight server as a Collector for a clustered NetApp filer, you must configure the Cluster-Mode FPolicy service on that server.

To configure the DataInsightFPolicyCmod service

1. Provision a Windows 2003 or 2008 server. Veritas recommends a minimum requirement of a Windows 2008 64-bit server with 4 to 8GB RAM and a quad core processor. A higher configuration may be required if the load on the FPolicy server is high.

   This computer hosts the FPolicy server.

2. Install the Data Insight Collector worker node or the Data Insight Management Server on this server.
3. Log in to the Data Insight Management Console.
4. In the Console, click **Settings** > **Data Insight Servers** to open the listing page for the server.
5. Select the server that is configured to monitor the NetApp clustered file server to open the details page for the server.
6. Click the **Services** tab.

7 Select **DataInsightFPolicyCmod** to expand the section.

8 To configure the service, enter the following details:

- The user-configured name to create FPolicy the ONTAP cluster.
- The IP address of the Data Insight Collector running the FPolicy server.
  The NetApp ONTAP Cluster-Mode filer connects with the DataInsightFPolicyCmod service running on the Data Insight Collector node on this IP address.
- The TCP port used on the DataInsightCollector server.
  The NetApp ONTAP Cluster-Mode filer connects on this port to the Data Insight Collector. Ensure that this port is not blocked by firewall.

9 Click **Configure**.

See "Configuring Data Insight services" on page 76.

Preparing the ONTAP cluster for FPolicy

The Veritas Data Insight FPolicy server registers with the ONTAP cluster and receives file access events from it, if FPolicy is enabled and configured on the corresponding Storage Virtual Machine (SVM) in the cluster. Veritas recommends that you automatically enable auditing when adding the clustered filers.

See "Configuring Data Insight services" on page 76.

When you enable FPolicy from the Data Insight console, Data Insight automatically does the following:

- Creates an FPolicy with a unique name.
- Creates an FPolicy engine by specifying the server IP address and the server port.
- Creates a CIFS and/or NFS event object.

Once you enable FPolicy on the SVM, it initiates a TCP connection to the Data Insight FPolicy server.

Note: You can choose to configure FPolicy on the OnTAP cluster manually. However, Veritas does not recommend using manual steps to monitor the SVMs in the cluster.

To configure FPolicy on the ONTAP cluster using manual steps

1. Launch a Telnet session with the SVM on which you want to configure FPolicy.
2. Run the following command to create an External Engine Object on the ONTAP shell:

```
diontapclust::> vserver fpolicy policy external-engine create
-vserver <Vserver name> -engine-name <choose an external engine
name> -primary-servers <IP address of Data Insight FPolicy server>
-port <port number on which Data Insight FPolicy server is
listening> -extern-engine-type asynchronous -ssl-option no-auth
```

3. Run the command to configure FPolicy to monitor specific CIFS events:

```
diontapclust::> vserver fpolicy policy event create -vserver
<Vserver name> -event-name <choose an event name> -protocol cifs
-file-operations create, create_dir, delete, delete_dir,
read,close, rename, rename_dir -filters first-read,
close-with-modification
```

4. Run the command to configure FPolicy to monitor specific NFS events:

```
diontapclust::> vserver fpolicy policy event create -vserver
<Vserver name> -event-name <choose an event name> -protocol nfsv3
-file-operations create, create_dir, delete, delete_dir, link,
read,write, rename, rename_dir, setattr
```

5. Run the command to create a resident FPolicy on the SVM:

```
diontapclust::> vserver fpolicy policy create -vserver <Vserver
name> -policy-name <choose a policy name> -events <specify list
of events> -engine <specify engine name> -is-mandatory false
```

6. Run the command to configure the scope of FPolicy:

```
diontapclust::> vserver fpolicy policy scope create -vserver
<Vserver name> -policy-name <choose a policy name>
-volumes-to-include "*" -export-policies-to-include "*"
```

About configuring secure communication between Data Insight and cluster-mode NetApp devices

You can secure the data flowing between the Management Server and the NetApp cluster using the Secure Socket Layer (SSL) protocol. Connections between Data Insight and the cluster-mode NetApp devices use a single certificate.

You can enable secure communication between Data Insight and cluster-mode NetApp devices through an SSL connection by using self-signed or CA-signed digital SSL certificate. Whether you use a CA-signed digital certificate depends on your security requirements and budget. You can obtain a self-signed digital certificate for free, but a digital certificate signed by a trusted CA can incur a considerable expense. A self-signed digital certificate is not as secure as a digital certificate signed by CA.

Follow these steps to configure the SSL communication:

- Generate the SSL certificates.
  See "Generating SSL certificates for NetApp cluster-mode authentication" on page 160.
- Upload the SSL certificates on the Data Insight Collector node configured to monitor the filer.
- Prepare the NetApp cluster for SSL authentication.
  - Install the certificates on the NetApp Admin Vserver.
  - Create security roles.
  - Create security logon.
  - Enable client authentication.

  See "Preparing the NetApp cluster for SSL authentication" on page 162.

When you configure the NetApp cluster-mode device from the Data Insight Management Server, you must select the SSL certificate which includes a `.key` file and a `.pem` file that are installed on the cluster-mode NetApp storage device.

See "Add/Edit NetApp cluster file server options" on page 215.

Generating SSL certificates for NetApp cluster-mode authentication

Generating Self-signed digital SSL certificate

Use the `openssl` command to create a self-signed certificate.

Generating Self-Signed certificate

1 To generate a Certificate Signing Request (CSR), go to the command prompt window and execute the following command:

```
openssl req -x509 -nodes -days 365 -newkey rsa:1024 -keyout
<yourKeyFileName>.key -out <yourCertName>.pem
```

The generated certificate looks as follows:


```
-----BEGIN CERTIFICATE-----
MIICwjCCAiugAwIBAgIJAJpgINzlWl06MA0GCSqGSIb3DQEBBQUAMHoxCzAJBgNV
BAYTAklOMRMwEQYDVQQIDApTb21lLVN0YXRlMSEwHwYDVQQKDBhJbnRlcm5ldCBX
aWRnaXRzIFB0eSBMdGQxEDAOBgNVBAMMB2Fhc2hyYXkxITAfBgkqhkiG9w0BCQEW
EmFhc2hyYXlAbmV0YXBwLmNvbTAeFw0xMzA3MzAxNjQ2NDRaFw0xNDA3MzAxNjQ2
NDRaMHoxCzAJBgNVBAYTAklOMRMwEQYDVQQIDApTb21lLVN0YXRlMSEwHwYDVQQK
DBhJbnRlcm5ldCBXaWRnaXRzIFB0eSBMdGQxEDAOBgNVBAMMB2Fhc2hyYXkxITAf
BgkqhkiG9w0BCQEWEmFhc2hyYXlAbmV0YXBwLmNvbTCBnzANBgkqhkiG9w0BAQEF
AAOBjQAwgYkCgYEAv8jid3ADQH/HQ05iZ6Tk0NF2cY9iiEna71PVKjM1L8GGkyWJ
kGioW2j1qoHO4kJEXUOMoX7YREOKLYbBQW5nx6rrg8Z3iFvP09YJnByonUIuN9QZ
96OHQ+ws9u6wNgM2LTJbcbOUUdJuOQNgaQ4XhzLDa6g0jEzyDBHbC05m2XUCAwEA
AaNQME4wHQYDVR0OBBYEFDdavnhJnCUHDJXgZEAovxcoYAsxMB8GA1UdIwQYMBaA
FDdavnhJnCUHDJXgZEAovxcoYAsxMAwGA1UdEwQFMAMBAf8wDQYJKoZIhvcNAQEF
BQADgYEAdnD5BzSlV2SiZJbOjzmhkYraNwG3WauDYlnzo8K0v6BFhxKEC/abjUaa
Ic/mBXEE8JqnLN7uqQf1wZtqIU60eNexMMdg+tstYe5O0Fnu27ss9HsmDD51A9LZ
kT5+XIfG21EYJMnFa1LwWTtmkla66GNhVEzzJKUtOXD23H6SyNc=
-----END CERTIFICATE-----
```


2 Save the content of the generated certificate in two separate files - <*yourCertName*>.pem and <*yourKeyFileName*>.key in the -keyout and -out files as mentioned in the command.

Generating a CA-signed digital SSL certificate

To generate a unique CA-signed certificate

1 Collect the following information to generate a certificate request:

- Common name
  The fully qualified DNS name of the NetApp Admin Vserver. This name must be the actual name of the server that is accessible by all the clients. For example: https://ditest.
- Organization Name
  For example, Veritas □¡ City □¡ State □¡ Country □¡ Expiration
- Organizational unit (optional)

- Country
  For example, US
- State
  For example, CA
- City
  For example, San Francisco
- Expiration
  Expiration time in days

2. On the cluster-mode Admin Vserver, switch to Administrator mode. From the command prompt window, run the following command:

   ```
   set -privilege admin
   ```

3. Run the command to generate the Certificate Signing request:

   ```
   security certificate generate-csr -common-name <hostname of the
   Admin Vserver> -size 2048 -country <name> -state <name of state>
   -locality <name of locality> -organization <name of organization>
   -unit <name of organizational unit> -email-addr <email address>
   ```

4. Copy the certificate and the key strings and save them in two different files with extension `.pem` and `.key` respectively.

   Paste the certificate to a text document exactly as it appears on the screen. Include the top line and bottom line (-----BEGIN CERTIFICATE REQUEST----- and -----END CERTIFICATE REQUEST-----) with the extension `.pem`. Make sure that no extra lines, spaces, trailing carriage returns, or characters have been inadvertently added.

   Paste the key strings from (-----BEGIN RSA PRIVATE KEY ----- and -----END RSA PRIVATE KEY-----) to a text document with the extension `.key`.

5. Send the CSR string (the string portion from “-----BEGIN CERTIFICATE REQUEST----- to-----END CERTIFICATE REQUEST-----“) to a Certification Authority (CA) electronically to apply for a digital certificate. After the CA sends you the signed digital certificate, you must install it with the associated private key (<*yourKeyFileName*>.key) on the Vserver.

Preparing the NetApp cluster for SSL authentication

You must install the signed digital certificate with the associated private key (<*yourKeyFileName*>.key) on the NetApp Admin Vserver

Installing the SSL certificate on cluster mode NetApp device

To install the certificate on the NetApp Admin Vserver

1. Log on to the Admin Vserver, and manually copy over the `.key` file to the server.
2. From the command prompt window, navigate to the location where you have saved the `.key` file.
3. Run the following command:

   ```
   security certificate install -type client-ca -vserver <Host name of AdminVserver>
   ```

   Note: : If you have a CA-signed SSL certificate, install the root certificate and each intermediate certificate of the CA that signed the certificate by using the command `security certificate install` command with `-type client-ca` parameter.

4. To verify whether the certificate is installed, run the following command:

   ```
   security certificate show -vserver <Host name of AdminVserver>
   ```

Note: You should keep a copy of your certificate and private key for future reference. If you revert or downgrade to an earlier release, you must first delete the certificate and private key.

Creating security roles

You must security roles for ONTAPI access with the common name that you have mentioned in the certificate.

Note: Authentication of a domain user credentials using SSL for monitoring NFS exports is not supported.

To create the security roles

1. From the command prompt on the Storage Virtual Machine (SVM), run the following commands:

   ```
   security login role create -role testrole -cmddirname "version"-access all
   security login role create -role testrole -cmddirname "vserver cifs"-access readonly
   ```

```
security login role create -role testrole -cmddirname
"vserver"-access readonly

security login role create -role testrole -cmddirname
"volume"-access readonly

security login role create -role testrole –cmddirname "vserver
fpolicy policy" -access all

security login role create -role testrole –cmddirname "vserver
fpolicy" -access readonly

security login role create -role testrole –cmddirname "vserver
fpolicy enable" -access all

security login role create -role testrole –cmddirname "vserver
fpolicy disable" -access all

security login role create -role testrole –cmddirname "statistics"
-access readonly

security login role create -role testrole -cmddirname "vserver
nfs" -access all

security login role create -role testorle -cmddirname "network
interface show" -access readonly
```

The `network interface show` privilege automatically assigns the following privileges to the role:

- `network interface create`
- `network interface modify`
- `network interface modify`

2 To view the security roles that were created, run the following command:

```
security login role show –vserver <Host name of AdminVserver>
```

Creating a security logon

You must create two security logon credentials - one with `authmethod` as *cert* and another with `authmethod` as *password*.

To create a security logon

1. Create a security logon with the common name that you have mentioned in the certificate. For the purpose of this procedure, let us assume that the common name is *ditest*. Run the following commands:

   ```
   security login create -username ditest -application ontapi
   -authmethod cert -role ditest -vserver <Host name of AdminVserver>
   ```

   ```
   security login create -username ditest -application ontapi
   -authmethod password -role ditest -vserver <Host name of
   AdminVserver>
   ```

2. To view the security logon that you created, run the following command:

   ```
   security login show –vserver Host name of AdminVserver>
   ```

Enabling client authentication on the NetApp cluster

To enable client authentication

1. Run the following command:

   ```
   security ssl modify –vserver <Host name of AdminVserver>
   ```

   For example -`security ssl modify –vserver < name of admin vserver > -ca < default ontap cert > -serial <default ontap cert> -common-name <default cert name> -server-enabled true -client-enabled true`

   Note: Default values can be obtained by pressing tab key from Ontap management interface command prompt.

2. To check the status, run the following command:

   ```
   >> security ssl show –vserver <Host name of AdminVserver>
   ```

After you install all required certificates, create the appropriate role and logon credentials, and enable client authentication, you can add the NetApp cluster to the Data Insight configuration.

Enabling export of NFS shares on a NetApp Cluster-Mode file server

Before you add a NetApp Cluster-Mode file server to Data Insight, you must identify the volumes and export policies to allow Data Insight to discover the NFS shares on the cluster.

Data Insight requires Read only access to root namespace (/) and super user (ROOT) access to volume namespace to scan NFS exports. If the root namespace and volume are configured with two different export policies, then you must provide access to the Data Insight Collector under each export policy rules to enable Data Insight to scan the NFS shares.

To enable export of NFS shares on a NetApp Cluster-Mode file server

1. To verify whether the Data Insight Collector IP address is specified in the export policy rule, execute the following command.

   ```
   export-policy rule show -vserver <name of vserver> -policyname
   <export policy> -clientmatch <IP of Data Insight Collector>
   ```

   The command must be repeated if root namespace is configured to use different export policies.

2. The following command will display a denied access message if the IP address of the Collector host is not specified in the export policy rule.

   ```
   export-policy check-access -vserver <vserver name> -volume <volume
   name> -client-ip <collector IP of DI> -authentication-method sys
   -protocol nfs3 -access-type read
   ```

3. If the Data Insight Collector is not specified in the export policy, modify the export policy rule. Run the following command:

   ```
   export-policy rule modify -vserver <name of vserver> -policyname
   <export policy> -clientmatch <IP of Data Insight Collector>
   ```

4. To validate the access given, run the following command and check read/write access is shown:

   ```
   export-policy check-access -vserver <vserver name> -volume <volume
   name> -client-ip <collector IP of DI> -authentication-method sys
   -protocol nfs3 -access-type read-write
   ```

Chapter 10

Configuring EMC Celerra or VNX monitoring

This chapter includes the following topics:

- About configuring EMC Celerra or VNX filers
- Credentials required for configuring EMC Celerra filers

About configuring EMC Celerra or VNX filers

Veritas Data Insight supports EMC Celerra, EMC VNX, and EMC Isilon file servers. Data Insight uses the EMC Common Event Enabler (CEE) framework to collect access events from the EMC filers.

As a prerequisite, you must complete the following tasks to enable Data Insight to monitor an EMC Celerra file server:

- Download the CEE framework from the EMC website. Install the framework on the same Windows server as the Data Insight Collector node or on a remote server in the same Active Directory domain.
- Obtain the necessary user credentials for accessing the EMC Celerra filer.
  See “Credentials required for configuring EMC Celerra filers” on page 173.
- Configure the EMC Celerra filer for auditing.
  See “Preparing the EMC filer for CEPA” on page 168.
- Configure Data Insight to receive event notifications from an EMC Celerra file server.
  See “Preparing Data Insight to receive event notification” on page 171.
- Add the EMC Celerra filer to Data Insight.
  See “Adding filers” on page 210.

See "Add/Edit EMC Celerra filer options" on page 220.

About EMC Common Event Enabler (CEE)

The CEE framework provides a working environment for the following mechanisms:

- VNX/Celerra Antivirus Agent (CAVA)
- VNX/Celerra Event Publishing agent (CEPA)

Veritas Data Insight uses the CEPA functionality of the CEE framework to receive event notifications. The EMC VNX/Celerra Event Publishing Agent (CEPA) is a mechanism that enables Data Insight to register with the EMC VNX or Celerra filer to receive event notifications from the filer. You can specify filters for the event type, the CIFS server, and the shares that you want to monitor during registration with the CEPA facility in the CEE framework. CEPA then sends notifications regarding the registered events to Data Insight.

Preparing the EMC filer for CEPA

The Veritas Data Insight server registers with the EMC Celerra or the EMC VNX filer through the EMC Common Event Enabler (CEE) framework. Data Insight uses this framework to receive notifications of file access events from the filer.

See "About EMC Common Event Enabler (CEE)" on page 168.

To configure the EMC Celerra or EMC VNX filer to send event information to Veritas Data Insight

1 Create a cepp.conf file on the EMC filer. The following is a sample of the code that the cepp.conf file must contain:

```
surveytime=90
pool name=matrixpool \
servers=<IP Address/Hostname of Windows server running the EMC CAVA service> \
postevents=* \
option=ignore \
reqtimeout=500 \
retrytimeout=50
```

Note: If the CEE server pool contains more than one server, you may separate each of the `server` entry by a | character. The setting ensures that the filer sends events to the CEE servers in a round robin fashion which ensures load balancing and high availability. However, the filer does not concurrently forward events to the multiple CEE servers. In case of VNX, you must modify the `cepp.conf` file so that events are simultaneously forwarded to the CEE server pool.

If there are two or more applications that make use of EMC's CAVA / CEE audit functionality, the cepp.conf file that resides on the filer(s) must be configured to include both servers. In such a case, the cepp.config contains details of each application pool that uses the CAVA/CEE auditing functionality.

```
surveytime=90
pool name=matrixpool \
servers=<IP Address/Hostname of Windows server running the EMC CAVA service> \
postevents=* \
option=ignore \
reqtimeout=500 \
retrytimeout=50
pool name=QFSpool \
servers=<IP Address/Hostname of Windows server running the EMC CAVA service> \
```

```
postevents=* \
option=ignore \
reqtimeout=500 \
retrytimeout=50
```

2 Copy the cepp.conf file to the root directory of the Data Mover. Run the following command: `server_file <datamover_name> -put cepp.conf cepp.conf`

   For example, `server_file server_2 -put /tmp/CEPA/cepp.conf cepp.conf`

3 Start the CEPP service on the filer. Run the following command:

   ```
   server_cepp <datamover_name> -service -start
   ```

   Ensure that the service has started by running the following command:

   ```
   server_cepp name of data mover -service -status
   ```

Note: For detailed information about configuring CEPA, refer to the EMC documentation.

Preparing Data Insight to receive event notification

The EMC Common Event Enabler (CEE) can be installed on the same Windows server as the Data Insight Collector node or on a remote server in the same Active Directory domain.

You must perform the following steps to route events from the Windows server on which the EMC CEE is installed to the Collector node.

To prepare Data Insight to receive event notification

1 Provision a Windows 2003 or 2008 server to run the EMC CEE framework in the same Active Directory domain as the filers you want to monitor.

2 Open Windows' Registry Editor (**Start** > **Run** > **regedit**).

3 In Registry Editor, navigate to `HKEY_LOCAL_MACHINE > SOFTWARE > EMC > CEE > CEPP > Audit > Configuration`.

4 Double-click **Endpoint**.

5 Modify the registry entry for the EMC CAVA service to allow access to the Data Insight Collector node. Depending on the type of your Data Insight deployment, there can be the following different scenarios:

   - The EMC CAVA service and the Collector node are running on the same machine, and the EMC CAVA service is only being used by Data Insight.

In this case, add the Data Insight key, `SymantecDataConnector`, to the **Endpoint** option.

- The EMC CAVA service and the Collector node are running on the same machine, and the EMC CAVA service is also being used by applications other than Data Insight. In this case, append the Data Insight key, `SymantecDataConnector`, to the **Endpoint** option. Each entry must be separated by a semi-colon.

Note: The above-mentioned scenarios are automatically configured at the time adding filers.

- The EMC CAVA service and the Collector node are running on separate machines, and the EMC CAVA service is being used only by Data Insight. In this case, add the Data Insight key in the format, `SymantecDataConnector@<IP address of the Collector>`, to the **Endpoint** option.
- The EMC CAVA service and the Collector node are running on separate machines, and the EMC CAVA service is also being used by applications other than Data Insight. In this case, append the Data Insight key in the format, `SymantecDataConnector@<IP address of the Collector>`, to the **Endpoint** option.

If the EMC CAVA service is installed on multiple machines, modify the registry entries on each of these machines.

6 Double-click the **Enabled** setting, and set its value to 1.

7 To start the EMC CAVA service, run the following command on the EMC filer to check the service status. For example,

```
Server_cepp  server_2 -pool -info
```

8 Install Data Insight Collector node.

9 Log in to the Data Insight Management Console.

10 Navigate to **Settings** > **Data Insight Servers** to open the Data Insight Servers details page for the Collector.

11 Navigate to the service configuration section on the filer, and click **Enable** to start the DataInsightCelerra service on the Collector node.

12 Under Credentials, enter the credentials that the service needs to run as. The specified credentials must be that of a domain user.

13 Click **Configure** to apply these settings to the server and start the EMC CAVA service.

See "Adding filers" on page 210.

See "Add/Edit EMC Celerra filer options" on page 220.

Credentials required for configuring EMC Celerra filers

Table 10-1 lists the set of credentials that are required by Veritas Data Insight during system configuration.

Table 10-1 Credentials for configuring EMC filers

Credential	Details
Credentials required to configure DataInsightCelerra service. The DataInsightCelerra service runs on the Collector and processes events sent by the CAVA services using the Windows RPC interface. This service must be configured on each Collector node that is used to connect to the EMC filers.	Required by the DataInsightCelerra service to run and authenticate itself with the EMC CAVA service provided by EMC, which runs on the Data Insight Collector node. The credential should belong to the user in the domain of which the Data Insight Collector node and the EMC filer are part.
Credentials required during filer configuration through the Data Insight Management Console.	Required to discover shares for EMC filer. This credential belongs to the EMC filer Control Station user who has administrative rights including XMLAPI v2 privilege (for example, nasadmin). See "Preparing Data Insight to receive event notification" on page 171.

Table 10-1 Credentials for configuring EMC filers *(continued)*

Credential	Details
Credentials required for scanning of shares.	Required for scanning of shares from the EMC filer. This credential belongs to the user in the domain of which the EMC filer is a part. Additionally, to be able to obtain share-level ACLs, the credentials must belong to the Domain Administrators group on the file server. You do not need this privilege if you do not want to get the share-level ACLs. In this case you will only need privileges to mount the share and scan the file system heirarchy. You must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs: ■ Traverse Folder/Execute File ■ List Folder/Read Data ■ Read Attributes ■ Read Extended Attributes ■ Read Permissions See "About configuring EMC Celerra or VNX filers" on page 167.

Chapter 11

Configuring EMC Isilon monitoring

This chapter includes the following topics:

- About configuring EMC Isilon filers
- Prerequisites for configuration of Isilon file server monitoring
- Credentials required for configuring an EMC Isilon cluster
- Configuring audit settings on EMC Isilon cluster using OneFS GUI console
- Configuring audit settings on EMC Isilon cluster using the OneFS CLI
- Configuring Isilon audit settings for performance improvement
- Preparing Veritas Data Insight to receive event notifications from an EMC Isilon cluster
- Creating a non-administrator user for an EMC Isilon cluster
- Using SmartConnect mapping for access zones
- Purging the audit logs in an Isilon filer

About configuring EMC Isilon filers

Veritas Data Insight supports the monitoring of EMC Isilon OneFS file servers, version 7.1 or later. Isilon OneFS is a cluster that has multiple hosts and a unified namespace-based file system across those hosts. Data Insight uses the CEE framework that is provided by EMC to fetch access events from EMC Isilon file servers. Ensure that CEE version 6.1 is used to enable Data Insight to fetch access event information from the filer.

Note: To avail the performance enhancements for auditing, Veritas recommends that the EMC Isilon cluster should be running OneFS version 7.1.0.6 or higher. Consult EMC Support to know if the OneFS version on your Isilon cluster is updated with the audit performance enhancements.

Veritas Data Insight also supports scanning and event monitoring of user's home directories on the EMC Isilon storage system.

Complete the following tasks to enable Data Insight to monitor an EMC Isilon file server:

- Complete all the prerequisites.
  See "Prerequisites for configuration of Isilon file server monitoring " on page 176.
- Obtain the necessary user credentials for accessing the EMC Isilon filer.
  See "Credentials required for configuring an EMC Isilon cluster" on page 177.
- Configure the audit settings on the EMC Isilon filer using either the GUI console or the command line.
  See "Configuring audit settings on EMC Isilon cluster using OneFS GUI console" on page 178.
  See "Configuring audit settings on EMC Isilon cluster using the OneFS CLI" on page 180.
- Perform additional auditing configuration for improved performance.
  See "Configuring Isilon audit settings for performance improvement" on page 182.
- Configure Data Insight to receive event notifications from an EMC Isilon cluster.
  See "Preparing Veritas Data Insight to receive event notifications from an EMC Isilon cluster" on page 183.
- Add the EMC Isilon filer to Data Insight.
  See "Adding filers" on page 210.
  See "Add/Edit EMC Isilon file server options" on page 223.

Once you have configured an Isilon file server, as a maintenance activity you must periodically clear the audit logs.

See "Purging the audit logs in an Isilon filer" on page 186.

Prerequisites for configuration of Isilon file server monitoring

Before you start configuring Isilon file servers, ensure the following pre-requisites are complete:

- The EMC Common Event Enabler (CEE) version 6.1 or later is installed. You can install CEE either on the same Windows server as the Data Insight Collector or on a remote server in the same directory service domain.

  Note: Veritas recommends you to use the latest version of CEE.

- The DataInsightCelerra service is installed on the Data Insight Collector.
- Microsoft .Net Framework version 4.5 is installed on your Collector node.
- A local or a domain user is configured on Isilon for use by Data Insight.
- The file-system auditing is enabled and audit settings are configured on the Isilon cluster.
  See "Configuring audit settings on EMC Isilon cluster using OneFS GUI console" on page 178.
- Note the port number from the URL used to access the Isilon OneFS Management Console. This port number is used by Data Insight for discovery purposes. Ensure that this port is not blocked by the Windows firewall in the Collector node.

Credentials required for configuring an EMC Isilon cluster

Table 11-1 lists the set of credentials that are required by Veritas Data Insight during system configuration.

Table 11-1

Credential	Details
Credentials required to configure DataInsightCelerra service.	Required by the DataInsightCelerra service to run and authenticate itself with the Isilon cluster. The service runs on the Data Insight Collector node. The credential should belong to the user in the domain of which the Data Insight Collector node and the Isilon cluster are part.
Credentials required during filer configuration through the Data Insight Management Console.	Required to discover shares for the Isilon cluster. See "Preparing Veritas Data Insight to receive event notifications from an EMC Isilon cluster" on page 183.

Table 11-1 *(continued)*

Credential	Details
Credentials required for scanning of shares	Required for scanning of shares from the Isilon cluster. This credential belongs to the user in the domain of which the Isilon is a part. Additionally, to be able to obtain share-level ACLs, the credentials must belong to the Domain Administrators group on the file server. You do not need this privilege if you do not want to get the share-level ACLs. In this case you will only need privileges to mount the share and scan the file system hierarchy. You must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs: ■ Traverse Folder/Execute File ■ List Folder/Read Data ■ Read Attributes ■ Read Extended Attributes ■ Read Permissions See "About configuring EMC Isilon filers" on page 175.

Configuring audit settings on EMC Isilon cluster using OneFS GUI console

To enable Data Insight to receive access event information from the Isilon cluster, you must configure auditing on the Isilon cluster. You can configure the audit settings either from the Isilon OneFS web-based GUI console. Alternatively, you can perform the same configuration using the Isilon command line interface.

See "Configuring audit settings on EMC Isilon cluster using the OneFS CLI" on page 180.

To configure audit settings on Isilon using OneFS console

1. On the OneFS WebUI, navigate to the **CLUSTER MANAGEMENT** > **Auditing**.
2. Under the **Edit Auditing Settings** section, select **Enable Protocol Access Auditing**.

   Note: It is recommended that you enable the OneFS auditing feature only after you install and configure Data Insight for your storage environment. Otherwise, the backlog consumed by Data Insight may be so large that results may be stale for a prolonged time.

3. Under the **Audited Zones** section, add the access zone that you want to audit. To enable auditing for the entire Isilon cluster, you can select the default System zone. For more information about access zones, see the EMC Isilon documentation.

   Note: Do not enable auditing for the access zones having file systems configured only for the NFS protocol. Data Insight currently does not support monitoring of NFS shares.

4. Under the **Event Forwarding** section, enter the uniform resource identifier for the server where the Common Event Enabler is installed. The format of the entry is: http://*<IP of the CEE server>*:port/cee.

   For example: http://10.209.302.152:12228/cee.

   Note that 12228 is the default CEE HTTP listen port. You must choose a port number that is same as the one configured in the registry on the computer where CEE is installed.

   See "Preparing Veritas Data Insight to receive event notifications from an EMC Isilon cluster" on page 183.

5. Under the section, **Event Forwarding**, add the host name of the storage cluster. You can either use the EMC Isilon SmartConnect cluster name or the DNS resolvable host name of one of the nodes of the cluster. Do not leave this field blank. This host name is later used to add the EMC Isilon cluster to the Veritas Data Insight configuration.

See "Configuring Isilon audit settings for performance improvement" on page 182.

Configuring audit settings on EMC Isilon cluster using the OneFS CLI

You can configure the audit settings on an EMC Isilon cluster using the command-line interface (CLI). You must be a root user on the EMC Isilon cluster to perform the configuration steps.

To configure and view audit settings on Isilon using the OneFS CLI

1. Log on to the Isilon OneFS cluster using the command line interface.
2. Issue the following commands:
   - To enable auditing:
     ```
     di-isilon-1# isi audit settings
     modify--protocol-auditing-enabled on
     ```

     Note: Veritas recommends that you enable the OneFS auditing feature only after you install and configure Data Insight for your storage environment. Otherwise, the backlog consumed by Data Insight may be so large that results may be stale for a prolonged time.

   - To adjust the audit log time
     Sometimes, auditing on the OneFS cluster may have been enabled before Data Insight is configured to monitor the cluster. In such cases, the cluster attempts to forward all events from the time auditing was configured. On OneFS 7.2, you can configure a setting to manually adjust time from which to begin forwarding the events to Data Insight.
     For example, use the following command to update the time to forward events newer than November 19, 2014 at 2:00 P.M.
     ```
     isi audit settings modify --cee-log-time "protocol@2014-11-19
     14:00:00"
     ```
   - To disable auditing:
     ```
     di-isilon-1# isi audit settings
     modify--protocol-auditing-enabled off
     ```
   - To configure the audit settings:
     ```
     di-isilon-1# isi audit settings modify
     --add-cee-server-uris=http://cee.example.com:12228/cee
     --audited-zones=system --hostname=di-isilon.example.com
     --config-auditing-enabled=no
     ```
   - To view the audit settings:

```
di-isilon-1# isi audit settings view
```

Depending upon the configured audit settings, you may see the following response:

```
Protocol Auditing Enabled: Yes
            Audited Zones: System
          CEE Server URIs: http://cee.example.com:12228/cee
                 Hostname: di-isilon.example.com
  Config Auditing Enabled: No
```

After you have enabled audit settings on the EMC Isilon cluster you must configure the Access Zones on the cluster.

Note: Do not enable auditing for the access zones having file systems configured only for the NFS protocol. Data Insight currently does not support monitoring of NFS shares.

To configure Access Zones using the OneFS CLI

1. Log on to the Isilon OneFS cluster using the command line interface.
2. Issue the following command:

```
isi zone zones modify <zone>

--audit-success {close | create | delete | get_ security | logoff
| logon | read | rename | set_security | tree_ connect | write |
all} | --clear-audit-success

--add-audit-success {close | create | delete | get_security |
logoff | logon | read | rename | set_security | tree_ connect |
write | all}

--remove-audit-success <string>

--audit-failure {close | create | delete | get_security | logoff
| logon | read | rename | set_security | tree_ connect | write |
all} | --clear-audit-failure

--add-audit-failure {close | create | delete | get_security |
logoff | logon | read | rename
```

Using the command line interface, you can enable specific audit events.

To enable specific audit events using the OneFS CLI

1. Log on to the Isilon OneFS cluster using the command line interface.
2. Issue the following command:

```
isi zone zones modify system --audit-success
create,delete,rename,set_security
```

Depending on your configuration, you may see the following response:

```
di-isilon-1# isi zone zones list -v
                Name: System
          Cache Size: 4.77M
       Map Untrusted:
          SMB Shares: -
      Auth Providers: -
      Local Provider: Yes
        NetBIOS Name:
      All SMB Shares: Yes
  All Auth Providers: Yes
  User Mapping Rules: -
Home Directory Umask: 0077
  Skeleton Directory: /usr/share/skel
       Audit Success: create, delete, rename, set_security
       Audit Failure: -
             Zone ID: 1
```

See “Configuring audit settings on EMC Isilon cluster using OneFS GUI console” on page 178.

See “Configuring Isilon audit settings for performance improvement” on page 182.

Configuring Isilon audit settings for performance improvement

Use the Isilon OneFS web-based GUI console to enable auditing of the following types of events by default:

- audit-success
- audit-failure
- create
- delete
- rename

- set_security

The number of different types of events that Data Insight monitors affects the system performance. To reduce performance overhead, you may disable auditing of audit-failure events. You must use the OneFS CLI to disable auditing.

To disable auditing for audit-failure events

1. Log on to the Isilon OneFS cluster using the command line interface.
2. Issue the command:

```
isi zone zones modify system --remove-audit-failure all
```

Note: To avail the performance enhancements for auditing, Veritas recommends that the EMC Isilon cluster should be running OneFS version 7.1.0.6 or higher. Consult EMC Support to know if the OneFS version on your Isilon cluster is updated with the audit performance enhancements.

Preparing Veritas Data Insight to receive event notifications from an EMC Isilon cluster

EMC Common Event Enabler (CEE) version 6.1 or later is required to receive events from EMC Isilon OneFS cluster. You can install CEE either on the same Windows server as the Data Insight Collector node or on a remote server in the same directory service domain.

To prepare Data Insight to receive event notification

1. Install EMC CEE framework on a Windows server so that it runs in the same directory service domain as the filers you want to monitor.
2. On the machine where you have installed the CEE, open the Windows Registry Editor (**Start** > **Run** > **regedit**).
3. Navigate to the location: **HKEY_LOCAL_MACHINE** > **SOFTWARE** > **EMC** > **CEE** > **CEPP** > **Audit** > **Configuration**. Depending on the type of your Data Insight deployment, perform any of the following:
   - If EMC CAVA service and the Collector node are running on the same machine, add the Data Insight key, `SymantecDataConnector`, to the **Endpoint** option.
   - If EMC CAVA service and the Collector node are running on separate machines, add the Data Insight key in the format, `SymantecDataConnector@<IP address of the Collector>`, to the **Endpoint** option.

Additionally, configure the following registry values:

- Navigate to the location: **HKEY_LOCAL_MACHINE** > **SOFTWARE** > **EMC** > **CEE** > **Configuration**. Note the key called **HttpPort**. The default key value is specified as `12228`. You can modify this key value but you must ensure that this value matches the port in the CEE URL that is specified during Isilon audit configuration.
  See "Configuring audit settings on EMC Isilon cluster using OneFS GUI console" on page 178.
- Navigate to the location: **HKEY_LOCAL_MACHINE** > **SOFTWARE** > **EMC** > **CEE** > **CEPP** > **Audit** > **Configuration**. Set the value of key **Enabled** to `1`.

4. Log on to the Data Insight Management Console.

   Navigate to **Settings** > **Data Insight Servers** to open the Data Insight Servers details page for the Collector.
5. Navigate to the Service Configuration section on the filer, and click **Enable** to start the DataInsightCelerra service on the Collector node.
6. Under **Credentials**, enter the credentials that the service needs to run as. The specified credentials must be that of a domain user.
7. Click **Configure** to apply these settings to the server.
8. Restart the EMC CAVA and EMC CEE Monitor service on the Collector machine or on the remote CAVA server, if the services are remotely installed.

Creating a non-administrator user for an EMC Isilon cluster

Data Insight requires a user account on Isilon to perform automatic discovery of CIFS shares and to list all local groups, group memberships, and local users. Data Insight can use a non-administrator account for this purpose. This account can be a local Isilon OneFS account or a domain account.

To configure a domain user for discovery and scanning of CIFS shares

1. Log on as an Isilon administrator to the Isilon cluster CLI using SSH or Telnet .
2. Run the following commands:
   - To create a role named *dirole*

     ```
     isi auth roles create --name dirole --description "Read-only
     role for DataInsight"
     ```

- To give the user the privileges to log on to the REST API platform framework to get a list of CIFS shares and to list users and groups

  ```
  isi auth roles modify dirole --add-user=username@domain
  --add-priv-ro=ISI_PRIV_SMB --add-priv-ro=ISI_PRIV_LOGIN_PAPI
  --add-priv-ro=ISI_PRIV_AUTH --add-priv-ro=ISI_PRIV_NETWORK
  ```

- To add the user to the Backup Operators group on Isilon, which enables Data Insight to scan all the CIFS shares

  ```
  isi auth groups modify "Backup Operators" --add-user
  username@domain
  ```

To configure a local user for discovery of CIFS shares

1. Log on as an Isilon administrator to the Isilon cluster CLI using SSH or Telnet.
2. Run the following commands:

   To create a new local user called diuser

   - `isi auth users create diuser --enabled yes --password xxxxxx`
   - To create a role named *dirole*

     ```
     isi auth roles create --name dirole --description "Read-only
     role for DataInsight"
     ```

   - To grant the user the privileges to log on to the REST API platform framework to get a list of CIFS shares and to list users and groups

     ```
     isi auth roles modify dirole --add-user=diuser
     --add-priv-ro=ISI_PRIV_SMB --add-priv-ro=ISI_PRIV_LOGIN_PAPI
     --add-priv-ro=ISI_PRIV_AUTH
     ```

Using SmartConnect mapping for access zones

EMC Isilon publishes shares through access zones. Access zones are virtual containers in which the Isilon cluster is partitioned. An access zone may be associated with an IP address pool. A pool is used to map an access zone with a SmartConnect zone and a range of IP addresses. The SmartConnect zone hostname and alias is typically used to discover shares instead of the access zone name. Data Insight discovers access zone to SmartConnect zone mapping and uses these names for scanning the shares. You must configure your preference when adding an Isilon filer to the Data Insight configuration. Ensure that access zone names are DNS resolvable.

See “Add/Edit EMC Isilon file server options” on page 223.

Note: Support for using SmartConnect to discover shares is available for Isilon OneFS version 8.0 and later.

You can also choose to input the alternate SmartConnect names against access zones mappings using a CSV file as input.

The CSV file must be in the following format:

filer name, access zone name, SmartConnect name.

The CSV file must be placed in the directory `DATADIR/conf`on the Management server. The name of the CSV is `isilon.csv`. When DiscoverySharesJob_mrg runs the next time, it discovers and scans shares using alternate name provided in the CSV file.

Note:

Data Insight does not support adding of two DFS shares with the same name from two different access zones to its configuration.

Purging the audit logs in an Isilon filer

By default an Isilon file server does not automatically clear the audit logs. Accumulation of audit logs over a long time can clutter disk space and as a result interrupt the auditing process itself. To avoid these problems you must periodically perform manual cleanup of the audit logs.

Note: The cleaning of audit logs may cause interruptions in the SMB client connections between the Isilon file server and the Collector node, leading to scan failures. To avoid disruption in scanning service, perform the cleaning operation during a planned maintenance window. The described procedure is applicable only to OneFS 7.1 and OneFS 7.2. For more information refer to EMC Isilon technical support.

To purge the audit logs from an Isilon filer

1. Log on to the Isilon cluster CLI using SSH or Telnet as a root user.
2. Stop the isi_audit_d and isi_audit_cee processes from automatically restarting, by executing the commands:

```
isi services -a isi_audit_d ignore

isi services -a isi_audit_cee ignore
```

3. Terminate the isi_audit_d and isi_audit_cee processes, by executing the commands:

   ```
   isi_for_array 'pkill isi_audit_d'
   ```

   ```
   isi_for_array 'pkill isi_audit_cee'
   ```

4. Verify that no isi_audit processes are running on the cluster, by executing the command:

   ```
   isi_for_array -s 'pgrep -l isi_audit'
   ```

5. Change directory to the audit directory, by executing the command:

   ```
   cd /ifs/.ifsvar/audit
   ```

6. To ensure that you are in the /ifs/.ifsvar/audit directory, execute the command:

   ```
   pwd
   ```

7. Optionally, create a backup of your audit directory if you want to preserve your old audit logs. You can move or copy them to another directory by using either `mv` or `cp` command.

8. Delete the audit directory by executing the command:

   ```
   rm -rf /ifs/.ifsvar/audit
   ```

9. Inform the Master Control Program (MCP) to resume monitoring the audit daemons, by executing the following commands:

   ```
   isi services -a isi_audit_d monitor
   ```

   ```
   isi services -a isi_audit_cee monitor
   ```

   MCP automatically restarts the audit daemons and reconstructs the audit directory on each node when the isi_audit_d process is running.

10. Check if the audit processes have restarted, by executing the command:

    ```
    isi_for_array -s 'pgrep -l isi_audit'
    ```

11. Verify that audit data was removed and reconstructed, by executing the command:

    ```
    isi_audit_viewer -t protocol
    ```

12. Verify that the audit log files are being populated after audit processes have restarted, by executing the command:

    ```
    isi_audit_viewer -t protocol
    ```

Chapter 12

Configuring Hitachi NAS file server monitoring

This chapter includes the following topics:

- About configuring Hitachi NAS
- Credentials required for configuring a Hitachi NAS EVS
- Creating a domain user on a Hitachi NAS file server for Data Insight
- Preparing a Hitachi NAS file server for file system auditing
- Advanced configuration parameters for Hitachi NAS

About configuring Hitachi NAS

Veritas Data Insight lets you monitor the storage devices running Hitachi NAS 12.x. For Hitachi NAS, you can monitor only the CIFS shares. Monitoring of NFS shares are not supported.

Each Hitachi NAS (HNAS) file server can consist of several Enterprise Virtual Servers (EVS). An EVS is a logical entity with its own IP address and file system. Data Insight monitors the configured EVS and the shares residing on these virtual servers.

Note: Veritas recommends that you do not use any version which is lower than Hitachi NAS 12.x. This may lead to serious degradation of filer performance when you enable auditing.

Complete the following tasks to enable Data Insight to monitor a Hitachi NAS file server:

- Obtain the necessary user credentials for accessing the Hitachi EVS host.
  See "Credentials required for configuring a Hitachi NAS EVS" on page 189.
- Create a domain user with necessary privileges on the Hitachi NAS EVS.
  See "Creating a domain user on a Hitachi NAS file server for Data Insight" on page 189.
- Configure the audit settings on the Hitachi NAS file server.
  See "Preparing a Hitachi NAS file server for file system auditing" on page 190.
- Add the Hitachi NAS EVS to Data Insight.
  See "Adding filers" on page 210.
  See "Add/Edit Hitachi NAS file server options" on page 233.

Credentials required for configuring a Hitachi NAS EVS

Table 12-1

Credential	Details
Credentials required during filer configuration through the Data Insight Management Console. Data Insight also uses the same credentials also for scanning of file metadata.	A domain user in the Administrators group on the Hitachi NAS EVS.

Creating a domain user on a Hitachi NAS file server for Data Insight

Data Insight needs a domain user with administrative privileges on Hitachi NAS EVS to perform the following tasks:

- To discover shares.
- To scan the shares for metadata.
- To automatically enable the Security Access Control Lists (SACLs) on each share.

To create a domain user on the Hitachi NAS file server:

1. Log in using SSH to the Hitachi NAS Admin Services EVS using the manager (administrator) credentials.
2. Execute the following command:

```
localgroup add Administrators <domain name> /<username>
```

Preparing a Hitachi NAS file server for file system auditing

To enable Data Insight to receive event information from a Hitachi NAS file server, you must complete the following tasks:

- Enable file system auditing on each EVS file system that you want to monitor.
- Configure and enable the audit log consolidated cache.

Note: The audit log consolidated cache accumulates all the individual audit logs for each EVS. Data Insight accesses the cache to monitor the audit events.

To configure file system audit settings on a Hitachi NAS EVS

1. Log in to the Hitachi NAS console using administrator credentials.
2. Navigate to **Home** > **File Services** > **File System Audit Policies**.
3. Select the EVS for which you want to enable auditing.
4. Click **add**. The **Add File System Audit Policy** page is displayed. The **Add File System Audit Policy** page displays a set of default settings for the audit policies. Retain the default settings.
5. Click **OK**.
6. Repeat the steps from 4 through 5 for the file systems in the EVS for which you want to enable auditing.

To configure and enable the audit log consolidated cache

1. Log in to the Hitachi NAS Admin Services EVS using SSH. Use an administrator credentials.
2. Execute the following command to switch from admin EVS to file services EVS:

```
console-context --evs <EVS name>
```

3. Execute the following command to configure the audit log consolidated cache:

   ```
   audit-log-consolidated-cache add -s <Size> <EVS name>
   ```

   For example:

   ```
   audit-log-consolidated-cache add -s 50MB EVS1
   ```

   Where, EVS1 is the name of file system EVS where you want to store the audit log consolidated cache file. Veritas recommends that you provision disk space of at least 50MB size for audit log consolidated cache to avoid loss of events.

4. To verify if the auditing is enabled on the EVS for the required file system, generate some activity on the shares created on the file system. Execute the following command on the Hitachi NAS Console to see if the events are generated:

   ```
   audit-log-show <Name of file system>
   ```

Advanced configuration parameters for Hitachi NAS

Sometimes the file operations per seconds (FOPS) for a Hitachi NAS filer may be affected once Data Insight starts fetching events. Data Insight introduces an idle period between two successive read calls to let the filer improve the speed of file operations.

You can configure the following parameters to tune Hitachi NAS filer:

max_events_to_pull	Specifies the maximum number of raw events fetched by Data Insight per execution of `hnas_util.exe`. Default value is 300000.
max_valid_events	Specifies the maximum number of events that can be stored in the Data Insight audit database. Default value is 100000.

min_events_per_run	For all the read calls which fetch less number of events than min_events_per_run, Data Insight sleeps for sleep_per_run microseconds, before making next read call. Default value is 100. You can set this parameter to fetch events at an optimized rate which does not degrade the performance of the filer. If the number of events to be fetched is too high, some of the events might be lost. In a single invocation of audit utility Data Insight makes read calls till it finds no new events to be fetched.
sleep_per_run (in microseconds)	Specifies the amount of time by which Data Insight waits between two successive read calls if the minimum number of events fetched per read call decreases below the value specified by the min_events_per_run parameter Default value is 500000 microseconds (0.5 seconds).
max_audit_interval (in minutes)	Specifies the approximate time interval after which audit file would be generated. Default value is 10 minutes.

To alter the configuration parameters, you can use `configdb.exe` command from the Data Insight Management Server.

For example:

```
configdb.exe -o -T filer -k 2 -J max_events_to_pull -j 50000
```

Where,

o- Object attribute

k- Filer ID

J- Attribute name

j- Attribute value

Chapter 13

Configuring Windows File Server monitoring

This chapter includes the following topics:

- About configuring Windows file server monitoring
- Credentials required for configuring Windows File Servers
- Using the installcli.exe utility to configure multiple Windows file servers
- Upgrading the Windows File Server agent

About configuring Windows file server monitoring

Data Insight uses an agent to collect access events from the Windows file server. The agent resides on the file server. The Data Insight agent consists of a filter driver that monitors the file system and records events that are relevant for Data Insight. It also consists of the Data Insight WinNAS service, which receives the event information from the filter driver and transfers it to the collector node that is configured for that filer.

Note that the Windows filter driver does not capture IP address from which accesses are made.

Complete the following tasks to enable Data Insight to monitor the Windows file server:

- Install the Data Insight agent on the Windows file server.
  You can choose to install the agent on the Windows file server automatically when adding the filer, or manually. Before you can install the agent automatically, ensure that the communication service port 8383 on the Collector node is accessible from the Windows file server.

For detailed information about installing the agent manually, see the *Veritas Data Insight Installation Guide*.

If you do not want Data Insight to access events for a Windows file server, it is possible to configure Windows file server without an agent. In this case, Data Insight scans shares of the filer from the Collector.

- Review the credentials that are required for configuring Windows file server monitoring in Data Insight.

  See "Credentials required for configuring Windows File Servers" on page 194.

- Add the Windows file server to Data Insight.

  See "Adding filers" on page 210.
  See "Add/Edit Windows File Server options " on page 225.

  You can either add a Windows file server to Data Insight through the Management Console, or if you want to add multiple filers together, you can use the `installcli.exe` utility.

  See "Using the installcli.exe utility to configure multiple Windows file servers" on page 196.

Note: All Data Insight worker nodes must be at the same level of major version as the Management Server. Windows file server agents can be one level lower than the Management Server version. Thus, Management Server 6.0 is compatible with versions 4.0.x, 4.5.x & 5.0.x as well as 6.0 of Windows file server agents. This gives you enough time to plan out the upgrade of your Windows file server agents.

You can also add a clustered Windows file server to Data Insight. Data Insight supports only a Microsoft Cluster Server (MSCS) configuration.

See "About configuring a DFS target" on page 245.

See "Using the Upload Manager utility" on page 102.

See "Adding filers" on page 210.

See "Add/Edit Windows File Server options " on page 225.

Credentials required for configuring Windows File Servers

Table 13-1 lists the set of credentials that are required by Veritas Data Insight during system configuration.

Table 13-1 Credentials for configuring Windows File Servers

Credential	Details
Credentials required to install agent on the Windows File Server.	This credential belongs to a user in the Administrators group on the Windows File Server. The credential is also used to discover shares and obtain storage utilization information from the filer.
Credentials required to discover shares and obtain storage utilization information on the filer.	Required for monitoring shares or when configuring a Windows File Server cluster. This credential belongs to a user in the Administrators group on the file server. If your configuration is not a Windows cluster or you do not want to collect storage utilization information for the filer, a credential with the privilege to list shares on the filer is sufficient.

Table 13-1 Credentials for configuring Windows File Servers *(continued)*

Credential	Details
Credentials required for scanning shares on the Windows File Server.	Required to scan a share. This credential belongs to a user with necessary share-level permissions and file system ACLs on a Windows File Server share. To be able to obtain share-level ACLs, the credentials must belong to the Power Users or Administrators group on the Windows File Server. You do not need this privilege if you do not want to get the share-level ACLs. To be able to scan a Windows File Server share successfully, you must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs: ■ Traverse Folder/Execute File ■ List Folder/Read Data ■ Read Attributes ■ Read Extended Attributes ■ Read Permissions **Note:** To enable Data Insight to successfully scan the shares on a clustered Windows File Server, you must ensure that the scanning user has domain level permissions of **Allow logon locally**.

Note: If you neither want Data Insight to install an agent automatically, nor do you want Data Insight to discover shares on the cluster or get storage utilization information, specifying the filer credentials is optional.

Using the installcli.exe utility to configure multiple Windows file servers

If you want to add multiple Windows file servers to Data Insight, you can use the `installcli.exe` utility.

The `installcli.exe` utility uses a `.csv` file with the following details as input:

- The host name or IP address of the Windows file servers that you want Data Insight to monitor.
- The host name, IP address, or ID of the Collector node that is configured to scan the filer.
- The host name, IP address, or ID of the Indexer node that is configured for the filer.
- The credentials that Data Insight should use to install the agent on the Windows file server. The credential should be in the format user@domain. `installcli.exe` also accepts LocalSystem credentials as value _LOCAL_. The same credentials must be added to Data Insight as a saved credential previously.
- True or false value indicating if the filer is clustered.
- The IP addresses of the agents. Separate multiple IP addresses with a semi-colon. If you do not want to use an agent to monitor the filer, indicate this option with a hyphen (-).
- The credentials that are required to scan the filer. The credential should be in the format user@domain. The same credentials should be added to Data Insight as a saved credential previously.
  See "Credentials required for configuring Windows File Servers" on page 194.
- True or false value indicating whether the scan should be enabled according to the specified schedule.
- In case of a Windows file server agent upgrade, RP or Full value indicating the type of upgrade you want to perform. This parameter is optional.
  Optionally, the name of the installer. If the name of the installer is not specified, an appropriate installer is picked from `installers` folder on the Collector.
- True or false value indicating whether event monitoring should be enabled. For example, *winnas.company.com,collector.company.com,indexer.company.com, Administrator@DOMAIN,FALSE,winnas.company.com, Administrator@DOMAIN,TRUE,TRUE,RP, Veritas_DataInsight_windows_winnas_6_0_0_3002_x64.exe*

To add multiple Windows file servers

1. Log in to the Data Insight Management Server.
2. Open a Windows command prompt and change to the `INSTALL_DIR\bin` directory, where `install_dir\bin` is the installation path for Veritas Data Insight.
3. Type the following command:

```
installcli -w <Path to .csv file with Windows File Server
specifications>
```

For detailed information on `installcli.exe`, see the *Command File Reference*.

Upgrading the Windows File Server agent

You can upgrade the Windows File Server agent automatically from the Data Insight Management Console.

Note: The option to upgrade the agent automatically appears only if you have configured the Windows File Server to allow Data Insight to automatically install the agent.

To upgrade Windows File Server agent automatically

1. Log on to the Management Console as Administrator.
2. Use the Upload Manager utility to upload the agent packages on Collector worker nodes corresponding to the Windows File Server agent.

   See "Using the Upload Manager utility" on page 102.
3. Select **Settings** > **Filers** to view the list of configured Windows File Servers.
4. Click the server on which you want the upgrade the agent.
5. On the configuration details page, click **Upgrade Agent**
6. Windows File Server agent upgrade window appears and displays a progress bar while upgrading.
7. Click **Finish** to exit setup.

Note: To upgrade the Windows File Server agent manually, see the *Veritas Data Insight Installation Guide*. You can upgrade multiple Windows File Server agents using the `installcli` utility. See "About configuring Windows file server monitoring" on page 193.

Note: Data Insight does not support upgrading a Windows File Server agent from 4.0 to 6.0. You must first upgrade the Windows File Server agent first to 4.5, 5.0 or 5.1, and then upgrade the agent to 6.0. Also, you can not upgrade a 32 bit 4.0.x Windows File Server agent to 6.0 as the last supported 32 bit agent is 4.5.x.

Chapter 14

Configuring Veritas File System (VxFS) file server monitoring

This chapter includes the following topics:

- About configuring Veritas File System (VxFS) file servers
- Credentials required for configuring Veritas File System (VxFS) servers
- Enabling export of UNIX/Linux NFS shares on VxFS filers

About configuring Veritas File System (VxFS) file servers

A Data Insight agent plugin, `vxdiplugind`, is used to monitor access events on the VxFs file servers. The plugin is part of the VxFS package and is automatically installed on the file server when Veritas Storage Foundation is installed. The plug-in captures events from the VxFS filer that Data Insight is monitoring, and saves it to a temporary database.The event data is then pulled by Data Insight, which fetches the access event information through Veritas Operations Manager (VOM) to gain vital insight into the user activity on the filer.

Data Insight uses NFS to scan all or a portion of VxFS shares remotely from the Collector node. Data Insight only monitors the access events on the VxFS devices exported by NFS.

Before you start configuring VxFS filers, verify the following:

- The file server must be installed with Storage Foundation 6.0.1 or higher.

- The file server must be installed with Veritas Operations Manager (VOM) 4.1 or higher.
- NFS version 3.0 is configured on the VxFS filer.
- The LDAP or NIS domains that your users are part of must be configured in Data Insight.
- The Collector node for the VxFs filer must be a Windows 2008 Enterprise server. Ensure that the Collector node monitoring the VxFS filer has services for NFS enabled as file server roles. You can install a role on Windows 2008 Enterprise server through the **Server Manager** > **Add roles** option.
- The filer is accessible from the Collector node using the host name or IP address you plan to use when adding the filer.

You can also add a clustered VxFS file server to Data Insight. Data Insight supports only a Veritas Cluster Server (VCS) configuration for VxFS file servers configured in failover mode. Parallel Clustered File System is not supported in this release.

See "Adding filers" on page 210.

See "Enabling export of UNIX/Linux NFS shares on VxFS filers" on page 202.

See "Add/Edit Veritas File System server options" on page 229.

Credentials required for configuring Veritas File System (VxFS) servers

Table 14-1 lists the set of credentials that are required by Veritas Data Insight during system configuration.

Table 14-1 Credentials required for configuring VxFS filers

Credentials	Details
Credentials required during filer configuration through the Veritas Data Insight Management Console.	Required to discover shares on the VxFs filer. This credentials belongs to a user on the UNIX server who has administrative rights on the VxFS filer (for example, root). The credential should belong to a root user on the VxFS filer. Optionally, this credential can also belong to a local user who has access to the Data Insight namespace which is already configured under `/opt/VRTSsfmh/di/web/admin`. To configure a user other than the root user, you must create or use an existing user account, which you can use to add the filer into the Data Insight namespace. To add a local user account under VOM: **1** Log in as root on the VxFs filer. **2** Change directory to `/opt/VRTSsfmh/di/web/admin`. **3** Create a `.xprtlaccess` file, and add the user to that file. For example, add *vomuser@unixpwd:user*, where vomuser is the name of the local user account.

Table 14-1 Credentials required for configuring VxFS filers *(continued)*

Credentials	Details
Credentials required for scanning on VxFS filer server	Required for scanning of shares from the VxFS filer. For scanning NFS shares, Data Insight needs a domain account or a LocalSystem account with at least read and execute permissions on all folders, alongwith at least read permission on all files. By default, Data Insight uses the User ID or Group ID 0 to scan NFS shares. You can configure an alternate User ID or Group ID from the **Data Insight Servers** > **Advanced Settings** section of the Collector node. See "Configuring advanced settings" on page 78. Additionally, you must also have share-level READ permissions on the NFS export.

Enabling export of UNIX/Linux NFS shares on VxFS filers

These instructions are for Red Hat Enterprise Linux operation system which has standalone Storage Foundation 6.0 installed and a file system created using VxFS. The steps will change depending upon other operating system flavors.

To enable export of NFS shares on VxFS filers

1. Login as root on the VxFS filer and open the `/etc/exports` file.
2. Specify the name of the share that you would like to monitor. For example, `/demoshare`, where the VxFS file system is mounted.

   Ensure that the device entries are added in `/etc/fstab` to automatically mount NFS file systems after reboot.

   Data Insight uses /etc/exports and /etc/fstab for NFS share discovery. Sample entries are shown below:

   ```
   root@RHEL5-VxFS ~]# cat /etc/fstab | grep vxfs

   /dev/vx/dsk/openldapdg/vol01 /openldaphome vxfs defaults,_netdev 0 0

   /dev/vx/dsk/openldapdg/vol02 /data vxfs defaults,_netdev 0 0

   /dev/vx/dsk/openldapdg/vol03 /didata vxfs defaults,_netdev 0 0

   [root@RHEL5-VxFS ~]# cat /etc/exports

   /openldaphome 192.168.0.10(ro,sync,no_root_squash)  192.168.0.11
   (rw,syc) /data/exportshare *(rw,sync,no_root_squash)

   /didata *(rw,sync,no_root_squash)
   ```

3. Specify the root access and read only access to Data Insight Collector node. For example,

   ```
   /demoshare <Collector node IP> (ro,sync,no_root_squash)

   ro:read only

   no_root_squash: root access.
   ```

   You can specify read `+write`, `root_squash`, `anonuid`, `anongid` or other settings, as required.
4. Run the following command to start the NFS daemon

   ```
   #service nfs start
   ```

See “Adding filers” on page 210.

Chapter 15

Configuring monitoring of a generic device

This chapter includes the following topics:

- About configuring a generic device
- Credentials required for scanning a generic device

About configuring a generic device

Data Insight supports scanning and auditing of access events on storage devices with varied file systems. You can configure Data Insight to monitor generic device filers in addtion to NetApp, EMC Celerra, Veritas File System (VxFS), Windows File Server.

Data Insight uses a web API to collect access event files from the generic device filer. The web API enables web clients to push events from the generic device filers to the DataInsightGenericCollector web service. The DataInsightGenericCollector web service collects incoming events and copies them to a specific location on the Collector worker node configured for the storage device. The events are later processed by the corresponding Indexer.

The DataInsightGenericCollector service uses one-way Secure Sockets Layer (SSL) to secure communication between the generic device filer and the Collector node. The client connects to the DataInsightGenericCollector service through a specific port, which is configurable.

You must develop a customized client using the API specification provided by Data Insight to add a filer as a generic device to Data Insight, and to start monitoring, auditing, and scanning events on the filer. For more information on the web API specification for the generic Collector service, refer to the *Data Insight Programmer's Reference Guide*.

See “Add/Edit a generic storage device options” on page 232.

Credentials required for scanning a generic device

Table 15-1 lists the set of credentials that are required by Veritas Data Insight to scan a generic storage device.

Table 15-1 Credentials for scanning a generic device

Credential	Details
Credentials required for scanning of shares.	Required for scanning of shares from the filer. When scanning CIFS shares, this credential belongs to the user in the domain of which the filer a part. While the exact set of permissions depends on the generic device being scanned, this user must generally belong to the Administrator's group on the device. If the credential is not part of Administrator's group, the scanner might not be able to get share-level ACLs for shares of this device. Typically, to scan CIFS shares, you must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs enabled for the scan credential: ■ Traverse Folder/Execute File ■ List Folder/Read Data ■ Read Attributes ■ Read Extended Attributes ■ Read Permissions For scanning NFS shares, Data Insight needs a Unix account with at least read and execute permissions on all folders, alongwith at least read permission on all files. By default, Data Insight uses User ID or Group ID 0 to scan NFS shares. You can configure an alternate User ID or Group ID from the **Settings** > **Advanced Settings** section of the Collector node. See "Configuring advanced settings" on page 78. When monitoring only NFS shares, you can specify **Use Local System account** from the scanning credentials drop-down, else you can specify credentials required to scan CIFS shares.

Chapter 16

Managing file servers

This chapter includes the following topics:

- About configuring filers
- Viewing configured filers
- Adding filers
- Custom schedule options
- Editing filer configuration
- Deleting filers
- Viewing performance statistics for file servers
- About disabled shares
- Adding shares
- Managing shares
- Editing share configuration
- Deleting shares
- About configuring a DFS target
- Adding a configuration attribute for devices
- Configuring a DFS target
- About the DFS utility
- Running the DFS utility
- Importing DFS mapping

About configuring filers

Veritas Data Insight collects and stores access events from NAS devices to service queries on user activity and data accesses. Before Data Insight can start collecting events, you must ensure that auditing is configured properly on the storage device. Data Insight collects access events using asynchronous APIs, namely, Fpolicy for NetApp filers, the CEE framework for EMC Celerra filers, and file system filter drivers for Windows File Server. Additionally, other generic devices can also publish events to Data Insight using web APIs.

Data Insight does not natively support monitoring and scanning of NFS shares for NetApp Cluster-Mode, EMC Celerra, EMC Isilon, Windows File Server, and Hitachi NAS devices. To enable Data Insight to discover NFS shares, you must add the filer as a generic device in the Data Insight configuration.

See "Managing Data Insight product servers" on page 67.

See "About configuring NetApp file server monitoring" on page 127.

See "About configuring a clustered NetApp file server" on page 147.

See "About configuring EMC Celerra or VNX filers" on page 167.

See "About configuring EMC Isilon filers" on page 175.

See "About configuring a generic device" on page 204.

See "About configuring Hitachi NAS" on page 188.

Viewing configured filers

In the Management Console, you can view all the filers that Data Insight is configured to monitor.

Additionally, you can also do the following tasks from the filer list page:

- Add multiple filers to the Data Insight configuration.
  See "Adding filers" on page 210.
- Add shares in bulk to a selected filer.
  See "Adding shares" on page 239.

Use the provided dynamic search filter to search for configured filers based on various pre-defined criteria, for example, the type of the filer. You can also use the **Filter** field at the top of the content pane to filter the list of filers based on the IP address or hostname of the filer in addition to the pre-defined filter criteria. The displayed list of filers changes automatically when you select the check box for a filter criteria. For instance, when you select NetApp in the **Device Type** category, the application displays a list of configured NetApp devices. Similarly, when you

select a Collector node in the **By Collector** filter, Data Insight displays a list of filers associated with the selected Collector node.

To view configured filers

1. In the Console, click **Settings** > **Filers**.

   The screen displays the list of configured filers

2. Review the following information about the filers:

   - The object ID of the filer. This numerical value is used to identify the filer when troubleshooting issues with the filer. This column is hidden by default. To view this column, click on the column header and select **Columns** > **ID**.
   - The name of the filer.
   - The number of shares monitored by the filer.
   - The health of the filer.
   - The type of filer -NetApp, EMC Celerra, Windows File Server, Veritas File System (VxFS) server, or a generic device.
   - Whether file system event monitoring is enabled.
   - The Collector node for the filer.
   - The Indexer node for the filer.
   - The scanning schedule for the filer. This column is hidden by default.

To review filer details

1. In the Console, click **Settings** > **Filers**.
2. Click the filer that you want to review, or click the **Select Action** drop-down and select **View**.

   The filer details screen appears.

To view filer events

1. In the Management Console, click **Settings** > **Filers**.
2. Click the **Select Action** drop-down for the corresponding server in the filers listing table, and select **Event Log**.

   The event log for that filer appears.
3. To download Data Insight logs for the filer for troubleshooting purposes, click the **Select Action** drop-down for the corresponding filer, and select **Download Logs**.

   Data Insight downloads a compressed folder containing the logs related to this filer from all relevant Data Insight servers.

   See "Downloading Data Insight logs" on page 383.

See "Adding shares" on page 239.

Adding filers

You must add filers that you want Veritas Data Insight to monitor.

To add filers

1. In the Console, click **Settings** > **Filers**.

   The Filers page displays the list of available filers.
2. On the Filers page, click the **Add New Filer** drop-down, and select the type of filer you want to add.
3. On the New Filer screen, enter the filer properties, and click **Add New Filer**.

   If you are adding a Windows File Server, Data Insight can automatically install an agent on the filer. This agent enables Data Insight to receive event notifications from the filer.

   For detailed information about installing the agent manually, see the *Veritas Data Insight Installation Guide*.

Adding filers in bulk

You can also add multiple filers at once to the Data Insight configuration.

To add filers in bulk

1. In the Console, click **Settings** > **Filers**.
2. On the Filers page, click the **Add Filers in Bulk** drop-down, and select the type of filer you want to add.

3 On the **Bulk Add <*Name of filer*>** pop-up, browse to the location of the CSV file that contains the configuration details for each filer.

Note: If you are configuring multiple filers for the first time, download a sample of the CSV file; create a CSV file with the details corresponding to the type of filer you want to configure.

4 Click **Upload**.

Note: Data Insight disables addition of filers if you do not have a valid license or if your license has expired. See "Overview of Data Insight licensing" on page 21.

Add/Edit NetApp filer options

Use this dialog box to add a new NetApp filer to Veritas Data Insight or to edit the configuration of an existing filer.

Table 16-1 Add/Edit NetApp filer options

Field	Description
Filer host name or IP address	Enter the host name or IP address of the filer that you want Data Insight to monitor. **Note:** The host name or IP address should be the same as the filer name is entered in Symantec Data Loss Prevention targets.

Table 16-1 Add/Edit NetApp filer options *(continued)*

Field	Description
Collector	Do the following: **1** Click **Select Collector**. **2** The **Select Collector** pop-up displays the configured collectors that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. You can also view the number of devices, shares or site collections assigned to each configured node. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Collector to the filer. Data Insight connects to the filer from this server. Veritas recommends that the Collector worker node share a fast network with the filer. **Note:** When monitoring NFS shares, ensure that the Collector node monitoring the filer must have services for NFS enabled as file server roles. You can install the role on Windows 2008 through the **Server Manager** > **Add roles** option.
Indexer	Do the following: **1** Click **Select Indexer**. **2** The **Select Indexer** pop-up displays the configured indexers that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Indexer to the filer. Events and metadata collected from the filer are processed and stored on the Indexer node. You can also migrate the file server to another Indexer. See “About migrating storage devices across Indexers” on page 103.
Filer administrator credentials	See “Credentials required for configuring NetApp filers” on page 128. Specifying the filer administrator credentials is optional, if you choose to not monitor events on the filer, nor enable share discovery.

Table 16-1 Add/Edit NetApp filer options *(continued)*

Field	Description
Test credentials	Click to test the availability of network connection between the Collector worker node and the filer, and to test the validity of specified credentials. Veritas recommends that you test the connection before proceeding to ensure that Data Insight is able to connect to the filer.
Filer is vFiler	Select the check box to indicate that this filer is a NetApp virtual file server.
Physical filer for vFiler	The host name or IP address of the physical NetApp file server that is associated with the virtual file server. If the Data Insight FPolicy safeguard is not enabled for the virtual file server, the field is not editable. See "Configuring scanning and event monitoring " on page 28.
Enable CIFS monitoring	Select this check box to enable monitoring of CIFS shares.
Enable NFS monitoring	Select this check box to enable monitoring of NFS shares.
Select domain	From the drop-down, select the domain to which the NetApp filer belongs. This option is enabled when you enable the monitoring of NFS shares.
Monitoring details	Select **Automatically discover and monitor shares on this filer** to allow Data Insight automatically discover shares of the filer and add them configuration. Discovery of shares takes place as soon as you add a new filer and then twice each day at 2:00 A.M. and 2:00 P.M. You can also choose to add shares manually. See "Adding shares" on page 239.
Exclude shares from discovery	Enter the details of shares which should not be included during discovery. This option is available if you select **Automatically discover all shares on this filer.** Specify comma-separated patterns that you want to ignore. Patterns can have 0 or more wildcard * characters. For example, tmp* ignores `tmp_A`, `tmp_abc`, *$ ignores shares `C$`, `EXT$`, and others.

Table 16-1 Add/Edit NetApp filer options *(continued)*

Field	Description
Enable storage utilization analytics	Select the check box to allow Data Insight to gather storage utilization information from the filer. This information is used when you generate Filer Utilization and Filer Growth Trend reports. The DataInsightFpolicy service running on the Collector node gathers information about storage utilization on the filer.
Enable file system event monitoring	Select to enable event monitoring on the filer.
Enable FPolicy automatically	Select to automatically enable FPolicy on the filer. If you clear this check box, you must manually enable FPolicy on the filer. See "Preparing the NetApp filer for Fpolicy" on page 136.
Register for explicit Read events	Select the option to register for explicit Read events. When this option is not selected, OPEN events are treated as READ events. **Note:** NFSv3 does not support OPEN events. This means that you will not see READ events for NFS shares when this check box is cleared. Veritas recommends that you do not register for explicit Read events. This can increase the load on the filer during peak traffic from third party applications such as backups over CIFS.
Enable filer scanning	Select the check box to enable filer scanning according to the specified schedule.
Scanning schedule for full scans	Select one of the following to define a scanning schedule for shares of this filer: ■ Use Collector's default scanning schedule. ■ Use custom schedule. See "Custom schedule options" on page 235. Veritas Data Insight periodically scans shares of the filer to obtain file metadata and security descriptors. Each Collector worker node by default initiates a full scan of shares at 7:00 P.M. on the last Friday of each month.
Scanner credentials	See "Credentials required for configuring NetApp filers" on page 128.

Table 16-1 Add/Edit NetApp filer options *(continued)*

Field	Description
Scan new shares immediately	Select this option to scan newly added shares immediately, instead of waiting for the normal scan schedule. Scanning proceeds only when scanning is permitted on the Collector node.

See "Enabling export of NFS shares on a NetApp file server" on page 144.

Add/Edit NetApp cluster file server options

Use this dialog box to add a new NetApp filer in cluster mode to Veritas Data Insight or to edit the configuration of an existing filer.

Table 16-2 Add/Edit NetApp cluster file server options

Field	Description
Cluster Management Host	Enter the host name or IP address NetApp Cluster Management host interface that is used to manage the nodes in the cluster.
Collector	Do the following: **1** Click **Select Collector**. **2** The **Select Collector** pop-up displays the configured Collectors that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. You can also view the number of devices, shares or site collections assigned to each configured node. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Collector to the filer. Data Insight connects to the NetApp Cluster Management host from this server. Veritas recommends that the Collector worker node share a fast network with the cluster management host.

Table 16-2 Add/Edit NetApp cluster file server options *(continued)*

Field	Description
Indexer	Do the following: **1** Click **Select Indexer**. **2** The **Select Indexer** pop-up displays the configured indexers that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Indexer to the filer. Events and metadata that are collected from the cluster are processed and stored on the Indexer node. You can also migrate the file server to another Indexer.
Cluster Management Interface credentials	Data Insight uses the credentials that you specify to discover the following: ■ The SVMs in the cluster and the CIFS server underneath each SVM. ■ The CIFS shares for each of the CIFS servers. ■ The NFS exports on the NetApp cluster. To enable Data Insight to successfully discover shares, ensure that HTTP port 80 is not blocked between Data Insight Collector and NetApp cluster management host. Specifying the filer administrator credentials is optional, if you choose to not monitor events on the filer, nor enable share discovery. See "Credentials required for configuring a clustered NetApp file server" on page 151.

Table 16-2 Add/Edit NetApp cluster file server options *(continued)*

Field	Description

Cluster Management SSL Certificate	Optionally, you can use an SSL certificate instead of the Cluster Management Interface credentials.
	■ Select the radio button if you want to authenticate the communication between Data Insight and the NetApp cluster using a digitally signed certificate.
	■ You must generate a self-signed or a CA signed certificate and install the SSL certificate on the cluster Admin Vserver to use this option.
	See "Generating SSL certificates for NetApp cluster-mode authentication" on page 160.
	See "Preparing the NetApp cluster for SSL authentication" on page 162.
	■ From the drop-down, select the a saved SSL certificate, or click **Add New**.
	■ On the Add SSL Certificate pop-up, do the following:
	  ■ In the **Saved SSL Certificate Identifier** field, enter a logical name to identify the filers for which the certificate is used.
	  ■ Browse to the location of the `.pem` and *.key* files to select each of them.
	  ■ Click **Upload** to upload the SSL certificate to the Data Insight Collector assigned to the filer.
	Note: If you choose to use SSL authentication, use local user credentials. Do not choose this option if you are using domain user credentials to configure the NetApp cluster-mode filer. The SVMs with NFS protocol enabled are not be discovered if you are using domain user credentials authenticated using SSL.

Test credentials	Click to test the availability of network connection between the Collector worker node and the filer, and to test the validity of specified credentials. Veritas recommends that you test the connection before proceeding to ensure that Data Insight is able to connect to the cluster management host. By default, Data Insight does not test credentials for the following HOMEDIR shares: ■ ~ ■ ~CIFS.HOMEDIR ■ CIFS.HOMEDIR ■ %w ■ %d

Table 16-2 Add/Edit NetApp cluster file server options *(continued)*

Field	Description
Storage Virtual Machine (SVM)	Every SVM node in the cluster has a CIFS server configured on it. The CIFS server represents a file server for Data Insight. If both CIFS and NFS protocols are enabled, Data Insight automatically discovers the CIFS server name. If only CIFs protocol is enabled, Data Insight discovers all the CIFS servers that are configured in the cluster. If only NFS protocol is enabled then Data Insight fetches the SVM name. From the drop-down, select the CIFS server/SVM that you want Data Insight to monitor. Ensure that you can resolve the CIFS server host name from the Collector node.
Enable CIFS monitoring	Select this check box to enable monitoring of CIFS shares.
Enable NFS monitoring	Select this check box to enable monitoring of NFS exports on the NetApp cluster.
Use CIFS Data LIF hostname for scanning (Optional) / Use NFS Data LIF hostname for scanning	You can use a Logical Interface (LIF) associated with the SVM to communicate with Data Insight. Data Insight uses the Data LIF to access data from the CIFS and NFS servers . If the Admin LIF and Data LIF are associated with two different networks, then you must specify the Data LIF name while scanning the CIFS shares or NFS exports that reside on configured CIFS servers. Providing a Data LIF hostname is also useful if the Admin LIF for the cluster is not configured for CIFS or NFS protocol access. Selecting the LIF hostname is optional if the name of the CIFS server and the Data LIF is the same.
Monitoring details	Select **Automatically discover and monitor shares on this filer** to allowData Insight to automatically discover shares of the filer and add them configuration. Discovery of shares takes place as soon as you add a new filer and then twice each day at 2:00 A.M. and 2:00 P.M. You can also choose to add shares manually. See “Adding shares” on page 239.

Table 16-2 Add/Edit NetApp cluster file server options *(continued)*

Field	Description
Exclude shares from discovery	Enter the details of shares which should not be included during discovery. This option is available if you select **Automatically discover all shares on this filer.** Specify comma-separated patterns that you want to ignore. Patterns can have 0 or more wildcard * characters. For example, tmp* ignores `tmp_A`, `tmp_abc`, *$ ignores shares `C$`, `EXT$`, and others.
Enable storage utilization analytics	Select the check box to allow Data Insight to gather storage utilization information from the filer. This information is used when you generate Filer Utilization and Filer Growth Trend reports. The DataInsightFpolicyCmod service running on the Collector node gathers information about storage utilization on the filer. See "About configuring FPolicy in Cluster-Mode" on page 149.
Enable file system event monitoring	Select to enable event monitoring on the filer.
Enable FPolicy automatically	Select to automatically enable FPolicy on the filer. Once you enable FPolicy, the Clustered Data ONTAP SVM node server initiates a TCP connection to the Data Insight FPolicy server. Ensure that the communication port in the firewall between the FPolicy server and the cluster management host is open. If you clear this check box, you must manually enable FPolicy on the filer. See "Preparing the NetApp filer for Fpolicy" on page 136.
Register for permission change events	Select if you want Data Insight to monitor the changes to permissions in your storage environment. By default, this option is not selected because it has a significant impact on the performace of the file server.
Enable filer scanning	Select the check box to enable filer scanning according to the specified schedule.
Scan new shares immediately	Select this option to scan newly added shares immediately, instead of waiting for the normal scan schedule. Scanning proceeds only when scanning is permitted on the Collector node.
Scanner credentials	See "Credentials required for configuring a clustered NetApp file server" on page 151.

Table 16-2 Add/Edit NetApp cluster file server options *(continued)*

Field	Description
Scanning schedule for full scans	Select one of the following to define a scanning schedule for shares of this filer: ■ Use Collector's default scanning schedule. ■ Use custom schedule. See "Custom schedule options" on page 235. Veritas Data Insight periodically scans shares of the filer to obtain file metadata and security descriptors. Each Collector worker node by default initiates a full scan of shares at 7:00 P.M. on the last Friday of each month.

Add/Edit EMC Celerra filer options

Use this dialog box to add a new EMC Celerra filer to Veritas Data Insight or to edit the configuration of an existing filer.

Table 16-3 Add/Edit EMC Celerra filer options

Field	Description
CIFS Server Name	Enter the host name of the CIFS server that is exported by the filer. Entering the IP address of the CIFS server is not permitted
Control Station Hostname/IP address	Enter the IP address of the filer's Control Station.
Collector	Do the following: **1** Click **Select Collector**. **2** The **Select Collector** pop-up displays the configured Collectors that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. You can also view the number of devices, shares or site collections assigned to each configured node. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Collector to the filer. Data Insight connects to the filer from this server. Veritas recommends that the Collector worker node share a fast network with the filer

Table 16-3 Add/Edit EMC Celerra filer options *(continued)*

Field	Description
Indexer	Do the following: **1** Click **Select Indexer**. **2** The **Select Indexer** pop-up displays the configured indexers that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Indexer to the filer. Events and metadata that are collected from the cluster are processed and stored on the Indexer node. You can also migrate the file server to another Indexer.
Control Station Credentials	Enter the credentials for the filer's Control Station. These credentials are used to discover shares on the filer and add them to the configuration. You can specify non-administrative credentials, however, Test Connection will fail. In this case, you can continue to add the filer, but you must add shares manually.
Virtual Data Mover	Select the check box if the filer is running a virtual data mover. This field is used to handle physical paths that are returned for virtual data movers.
Test credentials	Click to test the availability of network connection between the Collector worker node and the control station and the validity of the specified credentials. Veritas recommends that you test the connection before proceeding to ensure that Data Insight is able to connect to the filer. **Note:** You must skip testing of credentials if you choose to use a non-administrator user. When you click Test Credentials, Data Insight by default tries to discover shares on the filer using the Control Station credentials.

Table 16-3 Add/Edit EMC Celerra filer options *(continued)*

Field	Description
Monitoring details	Select **Automatically discover and monitor shares on this filer** to enable Data Insight to automatically discover shares of the filer and add them configuration. You can also choose to add shares manually. Clear the check box if you use Control Station credentials with insufficient privileges for share discovery. If you choose to use credentials that do not have administrator rights and XML v2 privilege, you must manually add shares to the configuration. Discovery of shares takes place as soon as you add a new filer and then twice each day at 2:00 A.M. and 2:00 P.M. See "Adding shares" on page 239.
Enable file system event monitoring	Select to enable event monitoring on the filer.
Enable filer scanning	Select the check box to enable filer scanning according to the specified schedule.
Scanning schedule for full scans	Select one of the following to define a scanning schedule for shares of this filer: ■ **Use the Collector's scanning schedule** ■ **Use custom schedule** See "Custom schedule options" on page 235. From the drop-down, select the appropriate frequency option. Veritas Data Insight periodically scans shares of the filer to obtain file metadata and security descriptors. Each Collector worker node by default initiates a full scan of shares on the last Friday of each month. **Note:** You can also customize the schedule per share using the Add/Edit Share dialog box.
Scanner credentials	See "Credentials required for configuring EMC Celerra filers" on page 173.
Scan new share immediately	Select this option to scan newly added shares immediately, instead of waiting for the normal scan schedule. Scanning will still run only when scanning is permitted on the Collector node.

Add/Edit EMC Isilon file server options

Table 16-4

Field	Description
Cluster Management Host	Enter the host name for the Isilon cluster. It can be EMC Isilon SmartConnect Cluster name or it can be the DNS resolvable host name of one of the hosts of the cluster. **Note:** The Cluster Management host name is the same host name which is entered during the configuration of audit settings on the Isilon cluster. See "Configuring audit settings on EMC Isilon cluster using OneFS GUI console" on page 178.
Cluster Management Port	Enter the port number on Isilon to be used for discovery. **Note:** This is the port which is part of the URL used to access EMC Isilon OneFS web-based management console.
Collector	Do the following: **1** Click **Select Collector**. **2** The **Select Collector** pop-up displays the configured Collectors that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. You can also view the number of devices, shares or site collections assigned to each configured node. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Collector to the filer. Data Insight connects to the filer from this server. Veritas recommends that the Collector worker node share a fast network with the filer
Indexer	Do the following: **1** Click **Select Indexer**. **2** The **Select Indexer** pop-up displays the configured indexers that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Indexer to the filer. You can also migrate the file server to another Indexer.

Table 16-4 *(continued)*

Field	Description
Cluster Management Host Credentials	Select the saved credentials from the drop-down list to access the Cluster Management host.
Test credentials	Click to test the availability of network connection between the Collector worker node and the Isilon cluster.
Monitoring Details	Select **Discover shares automatically** to enable Data Insight to discover the shares on the filer automatically and add them configuration. Indicate your preference as to the options that Data Insight should use to discover shares on the Isilon cluster. Select one of the following: ■ Access zones - Discovers shares from the configured access zones names, including the shares from the System zone. ■ SmartConnect Zone Name - Data Insight uses the SmartConnect name to discover and scan shares (including shares on system access zones) whenever SmartConnect is configured. ■ SmartConnect Zone Alias - SmartConnect aliases are used to ensure that SmartConnect names work across migration. Thus, aliases take precedence over SmartConnect names for scanning a share. If the SmartConnect Zone alias is not available, Data Insight uses SmartConnect name to discover shares across access zones. If none of these are DNS resolvable scan fails for these shares. **Note:** Support for using SmartConnect to discover shares is available for Isilon OneFS version 8.0 and later. Ensure that the host name of the access zones that you want Data Insight to discover are resolvable from the Collector node monitoring the filer. When access zone names are not used for accessing SMB shares, additional mapping can be imported into Data Insight See "Using SmartConnect mapping for access zones" on page 185.
Enable file system event monitoring	Select to enable event monitoring on the filer.
Enable filer scanning	Select the check box to enable filer scanning according to the specified schedule.

Table 16-4 *(continued)*

Field	Description
Scanning Schedule(Full Scan)	Select one of the following to define a scanning schedule for shares of this filer: ■ **Use the Collector's default scanning schedule** ■ **Use custom schedule** From the drop-down, select the appropriate frequency option.
Scanner credentials	See "Credentials required for configuring an EMC Isilon cluster" on page 177.
Scan newly added shares immediately	Select this option to scan newly added shares immediately, instead of waiting for the normal scan schedule. Note that a scan can run only when scanning is permitted on the Collector node.

Add/Edit Windows File Server options

Use this dialog box to add a new Windows File Server to Veritas Data Insight or to edit the configuration of an existing filer.

Table 16-5 Add/Edit Windows File Server options

Field	Description
Is a MSCS clustered file server	Select the check box if the Windows File Server is part of a Microsoft Cluster Server configuration.
Windows server name/Cluster name	Enter the host name or IP address of the filer that you want Data Insight to monitor. In case of a clustered Windows File Server, enter the host name or IP address of the cluster. **Note:** The hostname or IP address should be same as the filer name entered in Symantec Data Loss Prevention Discover targets.

Table 16-5 Add/Edit Windows File Server options *(continued)*

Field	Description
Select Collector node for this filer	Do the following: **1** Click **Select Collector**. **2** The **Select Collector** pop-up displays the configured Collectors that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. You can also view the number of devices, shares or site collections assigned to each configured node. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Collector to the filer. Data Insight connects to the filer from this server. It is recommended that the Collector worker node share a fast network with the filer.
Select Indexer node for this filer	Do the following: **1** Click **Select Indexer**. **2** The **Select Indexer** pop-up displays the configured indexers that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Indexer to the filer. Events and metadata that are collected from the cluster are processed and stored on the Indexer node. You can also migrate the file server to another Indexer.
Monitor mode	Select one of the following monitoring options: ■ **Monitor this filer using an agent** If you select this option, Data Insight is able to monitor all file system events on the filer and scan file system metadata. ■ **Monitor this filer without an agent** If you select this option, Data Insight scans the filer using CIFS to discover shares and obtain file metadata. However, in this case, Data Insight will not be able to monitor file system events.

Table 16-5 Add/Edit Windows File Server options *(continued)*

Field	Description
Agent names for this filer	This option is visible when adding a clustered file server that is monitored using an agent, but where the agent is installed manually. Select one or more agent nodes from the list that belong to this cluster. This option is also visible when editing a clustered file server.
Let Data Insight install the agent automatically	Select to allow Data Insight to install or upgrade the agent on the Windows File Server. Data Insight automatically installs the Windows File Server agent on the filer using the WMI interface and also registers the filer with the Management Server.
Node names to install agent	This option is only visible if you have selected **Is a MSCS clustered file server**. In the text box, enter comma-separated IP addresses or hostnames of the Windows File Server nodes, on which you want to install the agent.
Filer Administrator Credentials	Enter the credentials that Data Insightshould use to install the agent on the Windows File Server. See "Credentials required for configuring Windows File Servers" on page 194.
Test Connection	Click to test the availability of network connection between the Collector worker node and the filer, and the validity of the specified credentials. Veritas recommends that you test the connection before proceeding to ensure that Data Insight is able to connect to the filer.
Automatically discover and monitor all shares on this filer	Use this option to have Data Insight automatically discover shares of the filer and add them configuration. You can choose to exclude certain shares using the **Exclude shares** field. Discovery of shares takes place as soon as you add a new filer and then twice each day at 2:00 A.M. and 2:00 P.M.
Exclude following shares from discovery	Enter the details of shares which should not be included in share discovery. This option is available if you select **Automatically discover all shares on this filer**. Specify comma separated patterns that you want to ignore. Patterns can have 0 or more wildcard * characters. For example, `tmp*` ignores shares `tmp_A`, `tmp_abc`, *$ ignores shares `C$`, `EXT$` and others.

Table 16-5 Add/Edit Windows File Server options *(continued)*

Field	Description
Collect storage utilization information for the filer	Select to enable Data Insight to collect storage utilization information from the filer. This information is used to create Filer utilization and Filer Growth Trend reports.
Enable file system event monitoring	Select to enable event monitoring on the filer.
Enable filer scanning	Select the check box to enable filer scanning according to the specified schedule.
Scanning schedule for full scans	Select one of the following to define a scanning schedule for shares of this filer: ■ **Use the Collector's scanning schedule** ■ **Define custom schedule** Veritas Data Insight periodically scans shares of the filer to obtain file metadata and security descriptors. Each Collector worker node by default initiates a full scan of shares on the last Friday of each month. **Note:** You can also customize the schedule per share using the **Add/Edit Share** dialog box.
Scanner credentials	Select one of the following: ■ **Use LOCAL SERVICE credentials** Select to use the LOCAL SERVICE account to scan shares of the filer. This option is available only for the filers monitored using an agent. If you select this option, ensure that the LOCAL SYSTEM account has appropriate privileges to scan the shares. If the account does not have adequate privileges, the scans for such shares will fail if performed using this account. ■ **Use saved credentials** Select the saved credentials from the drop-down or specify new credentials. See "Credentials required for configuring Windows File Servers" on page 194.
Scan new shares immediately	Select this option to scan newly added shares immediately, instead of waiting for the normal scan schedule. Scanning will still take place during the hours when scanning is permitted on the Collector node.

Add/Edit Veritas File System server options

Use this dialog box to add a new Veritas File System (VxFS) filer to Veritas Data Insight or to edit the configuration of an existing filer.

Table 16-6 Add/Edit Veritas File System (VxFS) filer options

Field	Description
This is a VCS clustered file server	Select the check box if the Veritas File System server is part of a Veritas Cluster Server (VCS) configuration.
VCS cluster name	Enter the logical name of the VCS cluster. This field is available only if you select the **This is a VCS clustered file server** check box.
Cluster Node IP addresses	Enter the comma-separated list of the host names or IP addresses of the physical nodes in the VCS cluster.
Filer hostname or IP address	Enter the hostname or IP address of the filer that you want Data Insight to monitor. **Note:** The hostname or IP address should be the same as the filer name entered in Symantec Data Loss Prevention Discover targets.
Collector	Do the following: **1** Click **Select Collector**. **2** The **Select Collector** pop-up displays the configured Collectors that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. You can also view the number of devices, shares or site collections assigned to each configured node. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Collector to the filer. Data Insight connects to the filer from this server. It is recommended that the Collector worker node share a fast network with the filer. **Note:** Ensure that the Collector node monitoring the NFS must have services for NFS enabled as file server roles. You can install the role on Windows 2008 through **Server Manager** > **Add roles** option.

Table 16-6 Add/Edit Veritas File System (VxFS) filer options *(continued)*

Field	Description
Indexer	Do the following: **1** Click **Select Indexer**. **2** The **Select Indexer** pop-up displays the configured indexers that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Indexer to the filer. Events and meta-data collected from the filer is processed and stored on the Indexer node. You can also migrate the file server to another Indexer.
Login credentials	See "Credentials required for configuring Veritas File System (VxFS) servers" on page 200. Specifying filer administrator credentials is optional, if you choose not to monitor events on the filer, nor enable share discovery.
Test credentials	Click to test the availability of network connection between the Collector worker node and the filer, and to test the validity of the specified credentials. Veritas recommends that you test the connection before proceeding to ensure that Data Insight is able to connect to the filer.
Monitoring details	Select **Automatically discover and monitor shares on this filer** to have Data Insight automatically discover shares of the filer and add them to the configuration. Discovery of shares takes place as soon as you add a new filer and then twice each day at 2:00 A.M. and 2:00 P.M. You can also choose to add shares manually. See "Adding shares" on page 239.
Exclude shares from discovery	Enter the details of shares which should not be included during discovery. This option is available if you select **Automatically discover all shares on this filer.** Specify comma separated patterns that you want to ignore. Patterns can have 0 or more wildcard * characters. For example, tmp* ignores `tmp_A`, `tmp_abc`, *$ ignores shares `C$`, `EXT$` and others.

Table 16-6 Add/Edit Veritas File System (VxFS) filer options *(continued)*

Field	Description
Audit details	Select to enable file system monitoring on the filer. Enter the following details: ■ Select the **Enable file system event monitoring** check box to enable event monitoring on the VxFS filer. ■ Time to live - The value indicates the time for which the VxFS plugin will try to communicate with Data Insight. If communication fails after the specified time, the plugin will terminate and stop capturing events from the VxFS filer. The default TTL value is 24 hours. ■ Records per file - The number of records after which the events are flushed to the Collector node. You can also enable an advanced setting to flush the records to the Collector node every 10 minutes, irrespective of the number of records specified. By default, the limit is set to 100000 records per file. See "Configuring advanced settings" on page 78. ■ Domain: The name of the LDAP or NIS domain that the filer is a part of. The VxFS filer that you want to add should not be part of two domains at the same time.
Enable filer scanning	Select the checkbox to enable filer scanning according to the specified schedule.
Scanning schedule for full scans	Select one of the following to define a scanning schedule for shares of this filer: ■ **Use the Collector's scanning schedule** ■ **Define custom schedule** Veritas Data Insight periodically scans shares of the filer to obtain file metadata and security descriptors. Each Collector worker node by default initiates a full scan of shares at 7:00 P.M. on the last Friday of each month. **Note:** You can customize the schedule per share using the Add/Edit Share dialog box.
Scanner credentials	See "Credentials required for configuring Veritas File System (VxFS) servers" on page 200.
Scan newly added shares immediately	Select this option to scan newly added shares immediately, instead of waiting for the normal scan schedule.

See "Enabling export of UNIX/Linux NFS shares on VxFS filers" on page 202.

Add/Edit a generic storage device options

Use this dialog box to add a new generic storage device to Veritas Data Insight or to edit the configuration of an existing device.

Table 16-7 Add/Edit generic device options

Field	Description
Filer hostname or IP address	Enter the hostname or IP address of the device that you want Data Insight to monitor.
Collector	Do the following: **1** Click **Select Collector**. **2** The **Select Collector** pop-up displays the configured Collectors that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. You can also view the number of devices, shares or site collections assigned to each configured node. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Collector to the filer. Data Insight connects to the filer from this server. It is recommended that the Collector worker node share a fast network with the filer. **Note:** When monitoring NFS shares, ensure that the Collector node monitoring the filer must have services for NFS enabled as file server roles. You can install the role on Windows 2008 through the **Server Manager** > **Add roles** option.
Indexer	Do the following: **1** Click **Select Indexer**. **2** The **Select Indexer** pop-up displays the configured indexers that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Indexer to the filer. Events and metadata that are collected from the cluster are processed and stored on the Indexer node. You can also migrate the file server to another Indexer.

Table 16-7 Add/Edit generic device options *(continued)*

Field	Description
Domain	From the drop-down, select the domain to which the device belongs. This option is enabled when monitoring NFS shares.
Enable filer scanning	Select the check box to enable filer scanning according to the specified schedule.
Scanning schedule for full scans	Select one of the following to define a scanning schedule for shares of this filer: ■ Use Collector's default scanning schedule. ■ Use custom schedule. See "Custom schedule options" on page 235. Veritas Data Insight periodically scans shares of the filer to obtain file metadata and security descriptors. Each Collector worker node by default initiates a full scan of shares at 7:00 P.M. on the last Friday of each month.
Test credentials	Do the following: **1** Click to test whether the selected credentials have the required pemission to scan shares on the device. **2** Enter the name of a share on the device, and click **OK**.
Scanner credentials	See "Credentials required for scanning a generic device" on page 205.
Scan new shares immediately	Select this option to scan newly added shares immediately, instead of waiting for the normal scan schedule. Scanning proceeds only when scanning is permitted on the Collector node.

Add/Edit Hitachi NAS file server options

Use this dialog box to add a new Hitachi NAS file server to Veritas Data Insight or to edit the configuration of an existing file server.

Table 16-8 Add/Edit New Hitachi NAS file server

Field	Description
Hitachi EVS Hostname/IP	Enter the host name or IP address of the HNAS file system EVS that you want Data Insight to monitor.

Table 16-8 Add/Edit New Hitachi NAS file server *(continued)*

Field	Description
Collector	Do the following: **1** Click **Select Collector**. **2** The **Select Collector** pop-up displays the configured Collectors that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Collector to the filer. Veritas recommends that the Collector worker node share a fast network with the filer.
Indexer	Do the following: **1** Click **Select Indexer**. **2** The **Select Indexer** pop-up displays the configured indexers that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. You can also view the number of devices, shares or site collections assigned to each configured node. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Indexer to the filer. Events and metadata collected from the filer is processed and stored on the Indexer node. You can also migrate the file server to another Indexer. See "About migrating storage devices across Indexers" on page 103.
Hitachi EVS Credentials	Data Insight uses the credentials that you specify to discover the CIFS shares for each of the CIFS servers. See "Credentials required for configuring a Hitachi NAS EVS" on page 189.
Test Credentials	Click to test the availability of network connection between the Collector worker node and the Hitachi NAS file server.

Table 16-8 Add/Edit New Hitachi NAS file server *(continued)*

Field	Description
Monitoring Details	Select **Automatically discover and monitor shares on this filer** to allow Data Insight to automatically discover shares of the filer and add them to the configuration. Discovery of shares takes place as soon as you add a new filer and then twice each day at 2:00 A.M. and 2:00 P.M. You can also choose to add shares manually. See "Managing shares" on page 241.
Enable file system event monitoring	Select to enable event monitoring on the filer.
Enable Filer Scanning	Select the check box to enable filer scanning according to the specified schedule.
Scanning Schedule(Full Scan)	Select one of the following to define a scanning schedule for shares of this filer: ■ Use the Collector's default scanning schedule ■ Use custom schedule From the drop-down, select the appropriate frequency option.
Scan newly added shares immediately	Select this option to scan newly added shares immediately, instead of waiting for the normal scan schedule. Note that a scan can run only when scanning is permitted on the Collector node.

Custom schedule options

Table 16-9 describes the options that you can use to define the frequency of the scans.

Table 16-9 Custom schedule options

Option	Description
Never	Runs the scan as and when required.
Once	Runs the scan once at the specified time and date.
Daily	Runs the scan once every day. You must specify the time when the scan should be run.

Table 16-9 Custom schedule options *(continued)*

Option	Description
Weekly	Runs the scan once every week. You can choose to run it on every weekday, or on specific weekdays. Also, you must specify the time when the scan should be run.
Monthly	Runs the scan on the specified days of a month. You must specify the days of the month and the time when the scan should be run. Separate multiple days with a comma. For example, 2,5.
Custom Cron	Runs the scan according to a defined cron schedule. You can build strings in the cron format to specify custom schedules such as every Friday at noon, or weekdays at 10:30 a.m., or every 5 minutes between 9:00 a.m and 10 a.m. on Wednesdays and Fridays. More information about defining a cron schedule, see *http://quartz-scheduler.org/api/2.1.7*.

Editing filer configuration

After you add a filer to Data Insight, you can edit the filer's configuration. For example, you might need to edit any of the following:

- The IP address or hostname of the filer.
- The username and password of the user authorized to log in to the filer.
- The IP address or hostname of the Collector worker node configured to scan the filer.
- The Indexer node associated with the filer.
- The scanning schedule.
- The scanner credentials.
- Whether all shares are to be monitored.
- Whether new shares are to be scanned immediately.

Note: Data Insight disables addition of filers if you do not have a valid license or if your license has expired. You can edit the configuration of the filers that have already been added to Data Insight. However, Data Insight will not discover shares that have been added on the file server after the expiry of the license. See "Overview of Data Insight licensing" on page 21.

To edit filer configuration

1. In the Console, click **Settings** > **Filers**.

   This displays the list of available filers.

2. Do one of the following:

   - In the filer summary table click the **Select Action** drop-down and select **Edit**.
   - Click the filer whose configuration you want to edit. On the Filer details screen, click **Edit**.

3. On the **Edit Filer** page, make the necessary configuration changes.

   Note: The **Name** field is editable only for the filer types for which Data Insight supports renaming.

4. Click **Save** to save the changes.

See "Considerations for renaming a storage device" on page 277.

Migrating the filer to another indexer

Before you migrate a filer from one indexer to another, ensure the following:

- There is network connectivity between the old and the new Indexer.
- There is ample storage space on the new Indexer.
- Port 8383 is open between the indexers to allow the Communication Service to communicate between them.

To migrate a file server from one Indexer to another:

1. In the Data Insight Management Console, click **Settings** > **Filers**. This displays the list of available filers.

2. From the **Select Action** drop-down menu, select **Edit**.

   The Edit Filer screen is displayed. The **Connection Details** section displays the Collector and the Indexer nodes associated with the file server.

3. Click **Change Indexer**.

   The **Select Indexer** window opens displaying a list of available Indexers.

4. Select a destination Indexer.

5 Click **OK** to confirm the changes. The **Select Indexer** window closes.

The migration of Indexer takes place in the background. The time taken to complete a migration operation is proportional to the amount of data contained by the source Indexer.

6 To view the status of migration to another Indexer, navigate to the Filers list page.

7 Click the **Select Action** drop-down corresponding to the filer being migrated, and select **View Migration Status**.

8 On the **Migration Status** panel, you can view the status of each share on the filer that is being migrated. The status can either be Success or Pending.

9 You can cancel an ongoing migration by clicking the **Cancel Migration** button available in the **Migration Status** panel.

Canceling a migration reverses all the changes affected by the migration operation. You can cancel only the ongoing migrations. You cannot cancel a migration operation whose status is displayed as Success.

See "About migrating storage devices across Indexers" on page 103.

Deleting filers

You can delete a configured filer.

To delete a filer

1 In the Console, click **Settings** > **Filers** to display the configured filers.

2 Do one of the following:

- In the filer summary table, click the **Select Action** drop-down, and select **Delete**.
- Click the filer you want to delete, and on the filer details page, click **Delete** .

3 Click **OK** on the confirmation message.

Viewing performance statistics for file servers

You can view the performance graphs for the configured filers.

To view file server statistics

1 Click **Settings** > **Filers**.

2 Select the filer for which you want to view the performance statistics.

3 On the filer details page, click the **Statistics** sub-tab.

4 From the **Period** drop-down, select the duration for which you want to view the data.

5 The **Statistics** sub-tab displays the following high-level data:

- The time (in milliseconds) required to perform one CIFS I/O operation for a filer. The graph displays both average and maximum latency statistics.
- The average and maximum rate of incoming events per second for a filer.
- The count of all files and folders across all shares for a filer. The graph displays the number of paths that are scanned across all shares on a filer

Note: For EMC Celerra, VxFS, and Windows File Servers, you can only view the count of files and folders across all the shares on the filer.

About disabled shares

Data Insight automatically discovers new shares on file servers and adds them to Data Insight configuration, if automatic discovery of shares is enabled for the file server. Share discovery takes place when a new filer is added to Data Insight or when the scheduled share discovery process runs on the filer.

During the process of share discovery, Data Insight also discovers shares that have been deleted from the filer but are still present in the Data Insight configuration. Such shares are marked as disabled and Data Insight stops scanning these shares.

Adding shares

All shares on a filer are added to the Data Insight configuration when you add the filer to Data Insight. You must add shares present on the filer manually if you do not select the **Discover shares automatically** option when adding a filer. You can either add multiple shares at once using a CSV file or select individual shares on a filer that you want to add to the Data Insight configuration.

Note: Data Insight disables addition of shares if you do not have a valid license or if your license has expired. Also, Data Insight will not discover shares that have been added on configured file servers after the expiry of the license.

To add a share

1. In the Console, click **Settings** > **Filers** .
2. To add shares from a filer, do one of the following:
   - On the Filer list page, select the filer from which you want to add shares. Click **Add Shares in Bulk**. On the pop-up, browse to the location of the CSV file and select it. Click **Upload**.
     If you are adding shares in bulk for the first time, you must create a CSV file in a specific format. Download a sample CSV file to view the format.
   - Click a configured filer. On the Details page, click **Monitored Shares**.
3. On the Monitored Shares list page, click **Add New Share** or **Add Shares in Bulk**.
4. On the **Add Share** pop-up, enter the share properties, and click **Save**.

See "Add New Share/Edit Share options " on page 240.

Add New Share/Edit Share options

Use this dialog box to add a new share to Veritas Data Insightor to edit the configuration of an existing share.

Table 16-10 Field Descriptions

Field	Descriptions
Share name	Enter the name of the share you want to add. For example, *share1*. If this share belongs to a clustered filer, enter share name as *fileserver@share*, where, *fileserver* is the name of the file server within the cluster that hosts the share.
Physical path on filer	Enter the physical path of the share on the filer. For example, `F:\<Share name>`.
Scanning schedule	Select one of the following to define a scanning schedule: ■ **Use filer's scanning schedule** ■ **Define custom schedule** From the drop-down, select the appropriate frequency option.
Enable legal hold for this share	Select to preserve the activity information on the share. Selecting this option disables the deletion or archiving of access event information on the share. See "About archiving data" on page 37. See "About purging data" on page 38.

Table 16-10 Field Descriptions *(continued)*

Field	Descriptions
Parallel Scanning	■ Select the **Enable parallel scanning within this share** check box. ■ From the **Parallel Thread Count** drop-down, select the number of scan threads that you want Data Insight to run. It is recommended that you consider the number of processors while selecting the values. Based on your selection, Data Insight runs the scan threads: ■ Automatic: Data Insight computes the number of parallel scanner threads based on the number of processors and the amount of workload. ■ 1 or a higher value: Data Insight runs the specified number of scanner threads. See "Considerations for running a parallel scan" on page 31.

Managing shares

On the Monitored Shares details page you can view the detailed information about configured shares and run a customized scan on the configured shares.

Use the provided dynamic search filter to search for configured shares based on the name of the share.

To view configured shares

1. In the Console, click **Settings** > **Filers**.
2. Click the filer on which the share resides.
3. On the Filer Detail screen, click **Monitored Shares**.

   Review the following information about the shares:

   - ID of the share. The ID is required during troubleshooting. This column is hidden by default.
   - The name of the share.
     If this share belongs to a clustered filer, then the name should appear as *fileserver@share*, where, *fileserver* is the name of the file server within the cluster that hosts the share.
   - Type of this share, CIFS or NFS.
   - Enabled status of this share. This column is hidden by default.
   - Legal hold status for this share. This column is hidden by default.

- The number of parallel scan threads running on the share. This column is hidden by default.
- The scanning schedule for the share. This column is hidden by default.
- The date and time of the last full scan of the share.
- The date and time of the last incremental scan.
  Incremental scans are scans of the file system that includes only those paths that have changed since the last full scan. Incremental scans are much faster than full scans and they take place once every night at 7:00 P.M. You can configure incremental scans on the **Settings** > **Data Insight Servers** > **Advanced Settings** page.
- The time this share's index was last updated with scan information.
  After every scan, the index is updated with information about the changes to the folder hierarchy on a share. This indicates whether the last update was successful or failed. It also indicates the number of scan files pending for this share on the Indexer and the number of files that failed to be indexed. Such files are present in the `$data/indexer/err` folder on the Indexer. If there are failed files on the Indexer, you can move them from the `err` folder to `$data/inbox` folder and attempt a full scan of the share. If the information fails to be indexed again, contact Veritas Support.
- The time this share's index was last updated with access event information.
  As new access events come in, the index for the share is periodically updated with information about the new access events. This indicates whether the last update was successful or had failed. It also indicates the number of audit files pending for this share on the Indexer and the number of files that failed to be indexed. Such files are present in the `$data/indexer/err` folder on the Indexer. If there are failed files on the Indexer, you can move them to the `$data/inbox` folder on the Indexer. If they fail to be indexed again, contact Veritas Support.
- The status of event monitoring for the share, whether enabled or disabled.
- Whether a legal hold is being enforced for the share. You can choose to prevent access information for a share from being archived or deleted by putting a legal hold on the share. Data Insight preserves the access events information for such shares indefinitely.
  See "Add New Share/Edit Share options " on page 240.

4 Click the Export icon at the bottom of the page to save the data on the Monitored Shares panel to a `.csv` file.

You can also add a new share, edit the share's configuration, delete the share, start an unscheduled scan for a share, view the scan status, and download Data Insight logs from this page.

To view the scan status of a share

1. In the Console, click **Settings** > **Filers**.
2. Click the filer on which the share resides.
3. On the Filer Detail screen, click **Monitored Shares**.
4. Click the consolidated status icon to view the summary of the scan or the scan history of the share.

   Or select **Select Action** drop-down for the corresponding share, and select **Scan Status**.

   The consolidated status of the scans on the share appears. Click the **Scan History** tab on the pop-up to view the following details

   - The start and end time of the scan.
   - The time taken for the scan.
   - The type of scan, whether full or incremental.
   - The Collector node associated with the share.
   - The details of the scan. For example, if a scan has failed, the **Details** column indicates the exit code for the error message.
   - The user account that initiated the scan.

You can also view the scan status for a share from the **Scan History** sub-tab of the **Scanning** dashboard.

To view events pertaining to a share

1. In the Console, click **Settings** > **Filers**.
2. Click the filer on which the share resides.
3. On the filer details screen, click **Monitored Shares**.

4 Click the **Action** drop-down for the corresponding share, and select **Event Log**.

The event log for that share appears.

5 To download the Data Insight logs for the share, click the **Select Action** drop-down for the corresponding share, and select **Download Log**.

Data Insight downloads a compressed folder containing logs for this share from all relevant Data Insight servers.

See "Downloading Data Insight logs" on page 383.

To scan one or more shares in a batch

1 On the **Monitored Shares** tab, select a share, and select Scan from the **Select Action** drop down corresponding to a share.

Note: The **Scan** option is not available for shares that have been disabled.

---

2 To scan multiple share, select one or more shares using the check boxes.

3 Click **Scan**, and select **Scan Selected Records**.

Optionally, filter shares as needed using the filters available on the page. Click **Scan**, and select **Scan Filtered Records**.

Note: You can use a command line utility, `scancli.exe`, to further customize the scan, view the scan jobs running on a specified node, or display the scan status for specified shares. See scancli.exe on page 363. You can also use the Scanning dashboard view to scan shares and site collections based on more granular criteria.

To enable or disable shares

1 In the Management Console, click **Settings** > **Filers**.

2 Click the filer on which the share resides.

3 On the filer detail screen, click **Monitored Shares**.

4 To enable or disable a share, do one of the following:

- Select the share(s) and click **Enable Shares** or **Disable Shares**, as the case may be. You can select multiple shares and enable or disable them at once.
- Or, click the **Action** drop-down corresponding to a share, and select **Enable** or **Disable**, as the case may be.

Editing share configuration

After you add a share to Data Insight, you can edit the share's configuration. For example, you might need to edit the scanning schedule.

To edit share configuration

1. In the Console, click **Settings** > **Filers** to expand the Filer node.

   This displays the list of available filers. Click the appropriate filer to open the Filer details page.
2. Select the share whose configuration you want to edit, click the **Select Action** drop-down and select **Edit**.
3. On the Edit Share screen, make the necessary configuration changes.
4. Click **Save**.

Deleting shares

You can delete a configured share.

To delete a share

1. In the Console, click **Settings** > **Filers** to display the configured filers.
2. Click the filer, on which the share that you want to delete exists.
3. On the filer details page, under **Monitored Shares**, select the share that you want to delete.
4. Click the **Select Action** drop-down and select **Delete**.
5. Click **OK** on the confirmation message.

About configuring a DFS target

Veritas Data Insight supports Microsoft Distributed File System (DFS) targets.

DFS allows you to group shared folders located on different servers by transparently connecting them to one or more DFS namespaces. A DFS namespace is a virtual view of shared folders in an organization. DFS simplifies access to and management of shares by mapping a single logical namespace to shared folders on multiple filers. You can create folders within a DFS to create an additional level of hierarchy. For example, if you have a NetApp filer, NETAPP01, which has a shared folder called `NetAppShare1`. You can link a target, `HQ\Sales\Test`, present on a DFS server, DFSSvr01, to the subfolder named `Finance` within `NetAppShare1`.

You must first import the DFS mappings to physical shares in to Data Insight before you can view data using DFS hierarchy.

You can either map a filer using the hostname or IP address that is used to configure it in Data Insight, or use a list of aliases for the filer to map it. To use filer aliases to map filers to DFS targets, you must add a filer-level object attribute with a comma-separated list of aliases for configured filers and DFS filers.

See "Adding a configuration attribute for devices" on page 246.

When Data Insight consumes the CSV file containing the DFS mappings, it checks the object attribute to ensure that the filer is not being referred to in the Data Insight configuration by another alias. Similarly, Data Insight checks the DFS filer names in DFS filer alias list. If an alias is found for a DFS filer, then it means that the DFS filer is already configured in Data Insight and it is not added again to the Data Insight configuration.

Adding a configuration attribute for devices

To enable Data Insight to map physical and DFS filers to DFS targets, you must add a filer-level configuration attribute that lists the aliases used for the filers in the Data Insight configuration. When importing the DFS mappings, Data Insight checks the filer aliases in addition to the host name or IP address used to configure the filer. This feature ensures that DFS mapping is successful even if an alias is being used for a filer in the Data Insight configuration.

To add a configuration attribute from the Console

1. Click **Settings** > **Filers**.
2. On the **Filers** list page, click the **Select Action** drop-down for the physical or DFS filer for which you want to add the configuration attribute.
3. Select **Add Config Attribute**.
4. On the **Add Advanced Filer Attributes** pop-up, enter the following:
   - A logical name for the attribute. For example, *filer.alias*
   - A comma-separated list of aliases for the filer. For example, *FilerAlias1,Filer_Alias_2,Filer.Alias.3*
     Ensure that there is no space before and after comma in the alias list.

Configuring a DFS target

Before you can configure a DFS target you must configure all file servers which map to your DFS targets.

To set up a DFS target

1. Log in to the Management Console.
2. Create a .csv file containing the following information:
   - The name of the DFS server.
   - The DFS target.
   - The name of the filer that contains the share that you want to map to the DFS target.
   - The share on the filer.
   - Path under the physical share, if the DFS folder is mapped to a folder under physical share, else this value can be blank.

   For example, *DFSSvr01,HQ\Sales\Test,NETAPP01,NetAppShare1,\Finance*.
3. Click the **Settings** tab.
4. Click **Filers**, and select **Import DFS Mappings**.
5. In the Add new DFS mappings dialog, browse to the location of the .csv file, and click **Upload**.
6. Alternatively, open a Windows command prompt, and change to the `installdir\bin` directory, where `installdir\bin` is the installation path for Veritas Data Insight.
7. Type the following command:

   ```
   configdb -H  <name of the .csv file>
   ```

If certain DFS filer entries are dropped, you must ensure that either the device is configured or an alias for the device has been added to the Data Insight configuration. You can download the CSV file containing the dropped DFS shares, modify the DFS mappings file, and import the DFS mappings again.

About the DFS utility

The DFS utility, `mxdfsscan.exe`, maps root level DFS paths to actual storage server or share paths. It is used to export the DFS components (roots, root targets, links, and link targets) for all Windows DFS namespaces. The utility finds out physical level storage/filer link for all Domain DFS paths. It takes DFS root UNC path as input, for example`\\<DFS domain>\root`. This utility only enumerates online links and skips all offline links. It generates the output in .csv format.

The `mxdfsscan.exe` is a command line utility.

`mxdfsscan.exe` creates a DFS mapping CSV file with details of DFS servers that must be mapped in the Data Insight configuration. `mxdfsscan.exe` lets you specify only a single namespace at a time. You can run the utility twice to get information from two different DFS namespaces and create two different files from the output, for example, `test1.csv` and `test2.csv`. You can then import settings from `test1.csv` and `test2.csv` from the Data Insight Management Console

When you import a new DFS mapping file to Data Insight, the old mappings are maintained in Data Insight. For example, if you import mappings from `test1.csv` and then from `test2.csv`, the mappings from both files are displayed in Data Insight. However, if there are some duplicate mappings (the same DFS link appears twice – whether mapped to the same physical path or a different path), these mappings are not imported. A message is displayed indicating that there are duplicate mappings and hence one or more mappings cannot be imported.

Running the DFS utility

Ensure that the DFS servers are accessible from the machine you use to run the DFS utility.

To run the DFS utility

1. From the Windows Start menu, select **Run** and type `cmd` in the dialog box to open a command prompt window.
2. Change to the `installdir\bin` directory, where `installdir\bin` is the installation path for Veritas Data Insight.
3. Type the following command:

   ```
   mxdfsscan -n \\<DFS domain>\root -f dfsmap.csv
   ```

   where,

   `-n` is the name of the DFS root

   `-f` is the file name to which the DFS mappings have to be exported.

   `-e` is the option to exclude domain DFS paths. For example,

   ```
   mxdfsscan.exe -n \\MSAD\newroot -f dfsmap.csv -e exclude.txt
   ```

   An exclude list can have max 128 exclude entries. For example,

   ```
   \\DFS\root\AP\NUY
   ```

   ```
   \\DFS\root\users
   ```

   `-c` is the option to traverse a specified number of intermediate DFS servers to find a physical storage path. If the -c option is not specified then the utility takes the default value 5. This option helps avoid circular links in a DFS environment. If there are more hops then it logs all such links into `dfs_log_links.txt`.

Importing DFS mapping

You can import DFS mappings to Data Insight from the Management Console. To import DFS mappings, complete the following steps.

To import DFS mappings

1. Create a .csv file that contains information about the DFS mappings.

   See "Running the DFS utility" on page 248.
2. In the Console, click **Settings** > **Filers** to display the list of available filers.
3. Click **Import DFS mappings**.
4. On the Import DFS mappings window, browse to the location of the `.csv` file that contains information about the mapped DFS namespaces.
5. Click **OK**.

Chapter 17

Configuring SharePoint monitoring

This chapter includes the following topics:

- About SharePoint server monitoring
- Credentials required for configuring SharePoint servers
- Configuring a web application policy
- About the Data Insight web service for SharePoint
- Viewing configured SharePoint web applications
- Adding web applications
- Editing web applications
- Deleting web applications
- Adding site collections
- Managing site collections
- Removing a configured web application

About SharePoint server monitoring

You can use Data Insight to monitor unstructured data residing on servers running any of the following:

- Microsoft SharePoint™ 2016
- Microsoft SharePoint™ 2013

- Microsoft SharePoint™ 2010
- Microsoft Office SharePoint™ Server 2007 (MOSS 2007)

Data Insight monitors accesses to the data in the following SharePoint library types:

- Document library - Stores documents in the .pdf, .doc, .xls, .txt and other such file extensions.
- Picture library - Stores images.

Before you use Data Insight to scan a SharePoint server, you must complete the following tasks:

- To be able to configure Data Insight, you must know the URLs of the target SharePoint web applications.
- Ensure that .NET Framework 4.5 is installed on the Collector node that is responsible for the discovering site collections and collecting audit logs.
- Configure a policy for each web application.
  See "Configuring a web application policy" on page 252.
- Install and configure the Data Insight web service on the SharePoint server.
- Enable auditing on the SharePoint server. You can enable auditing from the Management Console when you add web applications, or directly from the SharePoint server.
  See " Add/Edit web application options" on page 257.

Credentials required for configuring SharePoint servers

Table 17-1 Credentials required for configuring SharePoint servers

Credential	Details
Credentials required to install the Data Insight web service on the SharePoint Server.	This credential belongs to a user in the Administrators group on the SharePoint server.

Table 17-1 Credentials required for configuring SharePoint servers *(continued)*

Credential	Details
Credentials required to discover web applications or site collections, and to collect scan information and audit data	This credential belongs to a site collection administrator for the configured sites and it must be in the same domain as the SharePoint server. It must have full control permissions not only on the configured web applications, but also on the web applications that are added to SharePoint subsequently. See "Configuring a web application policy" on page 252. for details on adding such a user credential.

Configuring a web application policy

When configuring SharePoint from the Data Insight console, you must specify an account for monitoring the configured site collections. This account must be a site collection administrator for the configured sites and it must be in the same domain as the SharePoint server. It must have full control permissions not only on the configured web applications, but also on the web applications that are added to SharePoint subsequently. The account should have the necessary privileges to set the appropriate audit flags, gather metadata about site collection content, and gather audit data from SQL Server databases for SharePoint

To enable Data Insight to gather audit and metadata from multiple site collections using a single user account, you must configure a policy for each web application from the SharePoint Central Administration Console.

To configure a policy for web application in SharePoint 2007

1. In the Central Administration website, click **Application Management**.
2. Under the Application Security section, click **Policy for Web application**.
3. Click **Add Users**.
4. In the **Web Application** drop-down list, select the web application that contains the site collections that you want Data Insight to monitor.
5. Select the appropriate zone. You can select **(All Zones)** if you want the user to be given permissions on all zones for the web application.
6. Click **Next**.

7 Choose the user account that has Full Control permissions.

   In the **Choose Permissions** section, select **Full Control - Has full control**

8 Specify whether this account operates as SharePoint System account. If you select the **Account operates as System** check box, all accesses made by this user account are recorded with the user name, *SharePoint System*.

9 Click **Finish**.

To configure a policy for web application in SharePoint 2010, SharePoint 2013, and SharePoint 2016

1 In the Central Administration Console, click **Application Management**.

2 Under the web applications section, click **Manage Web Applications**.

3 In the table displaying web application details, select the appropriate web application.

4 Click **User Policy**.

5 In the Policy for web application popup, click **Add Users**.

6 Select the appropriate zone. You can select **(All Zones)** if you want the user to be given permissions on all zones for the web application.

7 Click **Next**.

8 Choose the user account that will have Full Control.

   In the **Choose Permissions** section, select **Full Control - Has full control**

9 Specify whether this account operates as SharePoint System account. If you select the **Account operates as System** check box, all accesses made by this user account are recorded with the user name, *SharePoint System*.

10 Click **Finish**.

About the Data Insight web service for SharePoint

Before you can configure SharePoint monitoring in Data Insight, you must install the Data Insight web service on the SharePoint server. The Data Insight web service performs the following functions:

- Enables or disables auditing on the SharePoint server.
  You can enable or disable auditing of access events at the site collection level, either manually or from the Data Insight Management Console when adding web applications.
- Discovers site collections within a web application.
- Discovers all web sites and lists or libraries within a site collection

- Retrieves access event data from the SharePoint server. Data Insight uses this data to service queries on user activity and data access.
- Deletes audit logs from a site collection.

Installing the Data Insight web service for SharePoint

To enable Data Insight to collect access events, you must install the Data Insight web service, on a SharePoint server. When installing in a SharePoint farm, you must ensure that the web service is configured on all front-end web servers in the farm. Perform the following steps on any one front-end web server in your SharePoint farm. The web service installer automatically deploys the web service to all front-end web servers in the farm.

To install the Data Insight web service

1. Log on to the SharePoint server with an account that has SharePoint administrator privileges.
2. Load the Data Insight media on your SharePoint server.
3. Navigate to the folder where you have extracted or copied the installers.
4. To start the installation, double-click `Veritas_DataInsight_sharepoint_6.0_N.exe`, where, N is the build number.
5. Work through the installation wizard.
6. Click **Finish** to complete the installation process.
7. Verify whether the web service is deployed as expected.

After installing the Data Insight web service, you must verify whether it is successfully deployed on all front-end web servers in the SharePoint farm.

To verify the deployment of the web service in SharePoint 2007

1. In the Central Administration console, click the **Operations** tab.
2. Under **Global Configurations** section, click **Solution Management**.
3. Verify that the status for Data Insight solution for SharePoint is set to **Deployed**.
4. Click the link for the solution. Verify that the solution is deployed to all the front-end web servers in the farm by checking the value of **Deployed To** field.

To verify the deployment of the web service in SharePoint 2010, SharePoint 2013, and SharePoint 2016

1. In the Central Administration console, click the **System Settings**.
2. Under the Farm Management section, click **Manage Farm Solutions**.

3 Verify that the status for Data Insight solution for SharePoint, *sdispwebsvc.wsp*, is set to **Deployed**.

4 Click the link for the solution. Verify that the solution is deployed to all the front-end web servers in the farm by checking the value of **Deployed To** field.

Viewing configured SharePoint web applications

In the Management Console, you can view all the SharePoint web applications that Data Insight is configured to monitor.

To view configured web applications

1 In the Console, click **Settings** > **SharePoint Web Applications**.

The screen displays the list of configured web applications.

2 Click on a configured web application to view its detailed information, or click the **Select Actions** drop-down and select **View**.

The web application details page appears.

3 Review the following information about the web applications:

- The URL of the web application.
- The status of the web application — whether scanning and event monitoring are enabled for this web application.
- The Collector node for the web application.
- The Indexer node for the web application.

4 You can change the Indexer for the web application.

See "About migrating storage devices across Indexers" on page 103.

To migrate to another Indexer, do the following:

- Click **Select Indexer**.
- The **Select Indexer** pop-up displays the configured indexers that are available for selection.
- Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. You can also view the number of devices, shares or site collections assigned to each configured node. It is recommended that you assign a node on which the resource utilization is low.
- Click **Select** to assign the Indexer to the filer.

When the migration starts, click **View Details** to view the progress of the migration.

5 To view events pertaining to the web application, click **Event Log**.

The events for the selected web application are displayed. You can filter these events by Data Insight node, type of event, severity, and the service or the process that is generating the event.

Adding web applications

You must install the Data Insight web service on the SharePoint server, before you can add the web applications that you want Data Insight to monitor. In case the web service is not installed, Data Insight prompts you to install it before you can proceed with adding web applications.

Note: Data Insight disables addition of web applications if you do not have a valid license or if your license has expired.

To add web applications

1 In the Console, click **Settings** > **SharePoint Web Applications**.

The SharePoint page displays the list of configured web applications.

2 Click **Add SharePoint Web Application**.

3 On the Add web application screen, enter the URL of the web application you want to add and enter the properties.

4 Click **Save**.

Adding web applications in bulk

You can also add multiple web applications at once to the Data Insight configuration.

To add SharePoint web applications in bulk

1 In the Console, click **Settings** > **SharePoint Web Applications** .

2 On the web applications page, click the **Add SharePoint Web Applications in Bulk** drop-down, and select the type of filer you want to add.

3 On the **Bulk Add SharePoint Web Applications** pop-up, browse to the location of the CSV file that contains the configuration details for each web application.

Note: If you are configuring multiple web applications for the first time, download a sample of the CSV file; create a CSV file with the details corresponding to each web application that you want to configure.

4 Click **Upload**.

Add/Edit web application options

Use this dialog box to add a new web application to Veritas Data Insight or to edit the configuration of an existing web application.

Table 17-2 Add/Edit web application options

Field	Description
Web application URL	Enter the URL of the web application that you want Data Insight to monitor.
Collector for this Web application	Do the following: **1** Click **Select Collector**. **2** The **Select Collector** pop-up displays the configured Collectors that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Collector to the filer. Data Insight connects to the SharePoint server from Collector node. It is recommended that the Collector worker node share a fast network with the SharePoint server.

Table 17-2 Add/Edit web application options *(continued)*

Field	Description
Indexer for this Web application	Do the following: **1** Click **Select Indexer**. **2** The **Select Indexer** pop-up displays the configured indexers that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. You can also view the number of devices, shares or site collections assigned to each configured node. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Indexer to the filer. Events and metadata that are collected from the cluster are processed and stored on the Indexer node.
Default Site Collection Administrator	Enter the credentials that Data Insight should use to provide authenticated access to the Data Insight web service on the SharePoint server. See "Configuring a web application policy" on page 252.
Verify credential	Click to test the availability of network connection between the Collector worker node and the SharePoint server, and to test the validity of specified credentials. You must first ensure that the Data Insight web service is already installed on the SharePoint server. Veritas recommends that you verify the credentials before proceeding to ensure that Data Insight is able to connect to the SharePoint server.
Automatically discover and add all site collections in the selected Web applications to Data Insight	This checkbox is selected by default. This option allows you to automatically include all site collections in the selected web application for the purpose of monitoring. Clear the check box to add site collections manually. You can do this from the web application details page.
Exclude following site collections from discovery	Enter the details of the site collections which should not be included during discovery. This option is available if you select **Automatically discover and add site collections in the added SharePoint Web Applications**. Specify comma separated patterns that you want to ignore. Patterns can have 0 or more wildcard * characters. For example, https://webapp1/sites/test* ignores site collections https://webapp1/sites/testsite1 and https://webapp1/sites/testsite2.

Table 17-2 Add/Edit web application options *(continued)*

Field	Description
Monitor access for this Web application	Select to enable monitoring of access events for the web application.
Automatically enable auditing for site collections of this Web application	Select to automatically enable event monitoring for all site collections of this web application. You can also choose to manually enable auditing by logging in to the SharePoint server. For this purpose, you must have site collection administrator privileges on the SharePoint server.
Delete audit logs from SharePoint database after importing in Data Insight.	Select to delete audit logs from SharePoint to prevent the web application database from growing too large. By default, Data Insight deletes audit logs that are older than two days from the SharePoint server once every 12 hours. You can configure this interval from the **Advanced Settings** page for the corresponding Collector. You can choose to customize how often Data Insight should delete old audit logs from the Data Insight Servers node on the Management Console. See "Configuring advanced settings" on page 78.
Enable scanning for this Web application	Select the checkbox to enable SharePoint scanning according to the specified schedule.
Scanning schedule for full scans	Select one of the following to define a scanning schedule for the SharePoint servers in this farm: ■ **Use the Collector's scanning schedule.** ■ **Define custom schedule for farm.** From the drop-down, select the appropriate frequency. See "Custom schedule options" on page 235. Veritas Data Insight periodically scans site collections to obtain file metadata. Each Collector worker node by default initiates a full scan the SharePoint servers at 11:00 P.M. each night. **Note:** You can also customize the schedule for each site collection using the Add/Edit Site Collection dialog box.
Scan newly added site collections immediately	Select this option to scan newly added site collections immediately, instead of waiting for the normal scan schedule. Scanning will still proceed only when scanning is permitted on the Collector node.

Editing web applications

After you add a web application to Data Insight, you can edit its configuration. For example, you might need to edit any of the following:

- The user who is authorized to log in to the SharePoint server.
- The Collector worker node that is configured to scan the SharePoint server.
- The Indexer node that is associated with the web application.
- Enable or disable web application scanning or audit.
- The scanning schedule.

Note: Data Insight disables addition of web applications if you do not have a valid license or if your license has expired. You can edit the configuration of the web applications that have already been added to Data Insight. However, Data Insight will not discover site collections that have been added to the configured web applications after the expiry of the license.

To edit web applications

1. In the Console, click **Settings** > **SharePoint web applications**.
2. Do one of the following:
   - In the web application summary table, click the **Select Actions** drop-down and select **Edit**.
   - Click the web application whose configuration you want to edit. On the web application detail screen, click **Edit**.
3. On the Edit Web Application screen, do the following:
   - Make the necessary configuration changes.
   - To rename a web application, edit the name field.
   - To migrate a web application to another Indexer, click **Change Indexer**.
   - On the **Select Indexer** panel, select the new Indexer that you want to associate with the web application.
   - Click **Select**.
   - Click **Save** to save the changes.

4 To view the status of migration to another Indexer, navigate to the Web Application list page.

5 Click the **Select Action** drop-down corresponding to the web application being migrated, and select **View Migration Status**.

On the **Migration Status** panel, you can view the status of each site collection of the web application that is being migrated. The status can either be Success or Pending.

See "About migrating storage devices across Indexers" on page 103.

See " Add/Edit web application options" on page 257.

See "Considerations for renaming a storage device" on page 277.

Deleting web applications

You can delete a configured web application.

To delete a web application

1 In the Console, click **Settings** > **SharePoint Web Applications**.

2 Do one of the following:

- In the web application summary table, click the **Select Actions** drop-down and select **Delete**.
- Click the web application that you want to delete. On the web application detail screen, click **Delete**.

3 Select the check box to disable auditing on the web application.

4 Click **Yes** on the confirmation dialog box.

Adding site collections

All site collections in a web application are added to the Data Insight configuration when you add the web application to Data Insight. You must add the site collections manually if you do not select the **Automatically discover and add site collections in this web application** option when adding a web application. You can either add multiple site collections at once using a CSV file or select individual site collections that you want to add to the Data Insight configuration.

Note: Data Insight disables addition of site collections if you do not have a valid license or if your license has expired. Also, Data Insight will not discover site collections that have been added to configured web applications after the expiry of the license.

To add a site collection

1. In the Console, click **Settings** > **Share Point Web Applications** .
2. To add site collections do one of the following:
   - On the web applications list page, select the web application from which you want to add the site collections. Click **Add Site Collections in Bulk**. On the pop-up, browse to the location of the CSV file and select it. Click **Upload**.
     If you are adding site collections in bulk for the first time, you must create a CSV file in a specific format. Download a sample CSV file to view the format.
   - Click a configured web application. On the Details screen, click **Monitored Site Collections**.
3. On the Monitored Site Collection list page, click **Add Site Collection** or **Add Site Collections in Bulk**.
4. On the **Add New Site Collection** pop-up, enter the site collection properties, and click **Save**.

Add/Edit site collection options

Use this dialog box to add a new site collection to a configured web application or to edit the configuration of an existing site collection.

Table 17-3 Add/Edit site collection options

Field	Description
Site Collection URL	Enter the URL of the site collection that you want to add.
Site Collection Title	Enter a logical name for the site collection.

Table 17-3 Add/Edit site collection options *(continued)*

Field	Description
Scanning schedule	Select one of the following to define a scanning schedule for the site collection: ■ **Use Collector's scanning schedule** ■ **Define custom scanning schedule** From the drop-down, select the appropriate frequency. ■ See "Custom schedule options" on page 235. Veritas Data Insight periodically scans site collections to obtain metadata. Each Collector worker node by default initiates a full scan of the SharePoint servers at 11:00 P.M. each night.
Enable legal hold for this site collection	Select to preserve the activity information on the site collection. Selecting this option disables the deletion or archiving of access event information on the site collection. See "About archiving data" on page 37. See "About purging data" on page 38.

Managing site collections

On the site collections details page, you can view detailed information about a site collection, and edit or delete the site collection. You can also run a scan on the site collection from this page.

Use the provided dynamic search filter to search for configured sites based on the name of the site.

To view configured site collections

1. In the Console, click **Settings** > **SharePoint Web Application**.
2. In the web application summary table, click the web application that the site collections are a part of.

   The screen displays a list of all configured site collections.
3. On the web application details page, click **Monitored Site Collections**.
4. On the Site Collections listing page, review the following information about the configured site collections.
   - The name of the site collection.
   - The URL of the site collection
   - The scanning schedule for the site collection.

- The date and time of the last full scan of the site collection.
- The time this site collection's index was last updated with scan information.

  After every scan, the index is updated with information about the changes to the folder hierarchy on a site collection. This column indicates whether the last update was successful or has failed. It also indicates number of scan files pending for this site collection on the Indexer and the number of files that failed to be indexed. The files that failed to be indexed, are present in the `$data/indexer/err` folder on the Indexer. If you do have failed files on the indexer, you can move them from the `err` folder to the `$data/inbox` folder and attempt a full scan of the site collection.

  If the scan information again fails to be indexed, contact Veritas support.
- The time this site collection's index was last updated with access event information.

  As new access events come in, the index for the site collection is periodically updated with information about these events. This indicates whether the last update was successful or has failed. It also indicates number of audit files pending for this site collection at the Indexer and the number of files that failed to be indexed. Audit files are present in the `$data/indexer/err` folder on the Indexer. If you do have failed files on the indexer, you can try moving them back to `$data/inbox` folder on the Indexer.

  If the new audit information again fails to be indexed, contact Veritas support.
- The status of event monitoring for the site collection, whether enabled or disabled.
- Whether a legal hold is being enforced for the site collection. You can choose to prevent access information for a site collection from being archived or deleted by putting a legal hold on the site collection. Data Insight preserves the access events information for such site collections indefinitely.

  See "Add/Edit site collection options" on page 262.

5 Click the Export icon at the bottom of the page to save the data on the **Monitored Site Collections** panel to a .csv file.

You can also edit the properties of the site collection, start an unscheduled scan of the site collection, delete the site collection, view the event log or scan history of the site collection, or download logs for troubleshooting purposes.

To edit a site collection

1 On the web application details page, click **Monitored Site Collections**.

2 Select the site collection that you want to edit, and from the **Select Action** drop-down, select **Edit**.

3 On the Edit site collection screen, make the necessary configuration changes.

4 Click **Save**.

To delete a site collection

1 On the web application details page, click **Monitored Site Collections**.

2 Select the site collection that you want to delete, and from the **Select Action** drop-down, select **Delete**.

3 Click **OK** on the confirmation message.

To view the scan history of a site collection

1 On the web application details page, click **Monitored Site Collections**.

2 Select the site collection for which you want to view the scan history, and from the **Select Action** drop-down, select **Scan History**.

The scan history for the site collection appears. You can view the details in a tabular format or in a Timeline view. The tabular view displays the following details of a scan:

- The start and end time of the scan.
- The time taken for the scan.
- The type of scan, whether full or incremental.
- The Collector node associated with the site collection.
- The details of the scan. For example, if a scan has failed, the **Details** column indicates the exit code for the error message.
- The user account that initiated the scan.

The Timeline view displays an hourly and daily overview of the scans run on the site collection, including information about the status of the scan, whether failed, successful, partially successful or aborted.

You can also view the scan history of a site collection from the **Scan History** sub-tab of the **Scanning** dashboard.

To view events pertaining to a site collection

1 In the Console, click **Settings** > **SharePoint Web Applications**.

2 On the web application details screen, click **Monitored Site Collections**.

3 Click the **Select Action** drop-down for the corresponding site collection, and select **Event Log**.

The event log for that site collection appears.

4 To download the logs for the site collection, click the **Select Action** drop-down for the corresponding site collection, and select **Download Logs**.

Data Insight downloads a compressed folder containing the logs for this site collection from all relevant Data Insight servers.

See "Downloading Data Insight logs" on page 383.

To scan site collections in a batch

1 On the Monitored site collections tab, click the **Scan** button.

Note: The **Scan** option is not available for site collections that have been disabled.

---

2 On the Scan Site Collections pop-up, select one of the following:

- **Scan all** - To scan all the configured site collections immediately.
- **Scan with last scan status** - To scan site collections based on the following criteria:
  - Site collections on which the last scan has failed completely.
  - Site collections that have never been scanned before.
  - Site collections on which the last scan has failed on certain paths.

3 Select one or more of the following conditions:

- Scan the site collections that have not been scanned for *n* number of days. Enter the interval in the field.
- Include site collections matching specified patterns. You can enter multiple patterns separated by a comma. You can also specify one or more wildcards in the pattern. For example vol*, *$.
- Exclude site collections matching specified patterns. You can enter multiple patterns separated by a comma. You can also specify one or more wildcards in the pattern. For example vol*, *$.
- Select **Add to the top of scan queue** to add the scans to the top of the scan queue.

4 Select **Start scanning**.

Note: You can use a command line utility, `scancli.exe`, to further customize the scan, view the scan jobs running on a specified node, or display the scan status for specified site collections. For details, See scancli.exe on page 363.

To enable or disable site collections

1. In the Management Console, click **Settings** > **SharePoint Web applications**.
2. Click the web application that the site collections are a part of.
3. On the filer detail screen, click **Monitored Site Collections**.
4. Click the **Action** drop-down corresponding to a share, and select **Enable** to enable a site collection that has been disabled.
5. Click the **Action** drop-down corresponding to a share, and select **Disable** to disable a site collection.

See "Adding site collections" on page 261.

See "Add/Edit site collection options" on page 262.

Removing a configured web application

If you want to remove an existing SharePoint web application from Data Insight you must complete the steps in the correct order.

To remove a web application from Data Insight

1. Delete the configured web applications from the Data Insight console.

   Deleting the web application enables you to disable auditing for the monitored SharePoint web applications

   See "Deleting web applications" on page 261.
2. On the SharePoint server, disable auditing of the web applications that are deleted from Data Insight.
3. Uninstall the Data InsightWeb service from the SharePoint server.

Chapter 18

Configuring cloud content sources

This chapter includes the following topics:

- About configuring Box monitoring
- Using a co-admin account to monitor Box resources
- Viewing configured cloud sources
- Configuring Box monitoring in Data Insight
- Managing cloud storage accounts

About configuring Box monitoring

Data Insight monitors the data stored on cloud-based content sharing and management services, such as Box, to provide you access and ownership information.

Box uses the Open Authorization 2 (OAuth2) protocol to permit access to a third-party application. Data Insight uses Box Enterprise administrator credentials to gain access to a Box Administrator account, get information about all configured users, and to get information about events performed by them. The Administrator credentials authorize Data Insight to access the Box Enterprise account. The authorization received from Box is encrypted and saved in the Data Insight configuration so that only the Collector and Indexer specific to the device have a copy of it. The credentials are used by Data Insight to impersonate a user account to query Box for metadata.

In the Data Insight configuration, the Box account corresponds to a share on a file server. Since Data Insight uses the administrator account credentials to impersonate

other users, you will always see just one share for every Box account. Whereas, different user accounts correspond to folders on a file share.

Data Insight scans the Box account for metadata such as path of the file or folder, modified by, modified at, created by, created at, size, file or folder type, and whether more than one user collaborate on the resource.

Data Insight maps the user name for every user account attribute to the users' attributes in the directory service for the purpose of ascertaining ownership on folders.

See "Using a co-admin account to monitor Box resources" on page 269.

Using a co-admin account to monitor Box resources

It is recommended that you use administrator credentials to configure Box resources in Data Insight. However, if you prefer using any other user credentials to allow Data Insight to access the Box Administrator account, you must provide that user account specific privileges.

To configure a co-admin account

1. Log on to Box Administration console using Administrator account credentials.
2. Select the user account that you want to provide co-admin privileges. Select the following check boxes to edit the user's access permission settings:
   - **Allow this user to see all managed users.**
   - **Allow this user to synchronize files between Box and the user's desktop.**
   - **User is granted the following privileges:**

Users and Groups	**Manage users**
	Manage groups
Files and Folders	**View users' content**
	Log in to users' accounts
Reports and Settings	**View settings for your company**
	Run new reports and access existing reports

The privileges that are provided in Step 2 do not allow the co-admin account access to the files and folders owned by the Administrator and by other co-admin users.

3 To provide preview access to all the top- level folders of Administrator and other co-admin accounts, on the **Change User Settings** page, select the folders that you want to grant the co-admin account access to.

4 The setting in Step 3 provides access only to the shared folders of the Administrator and co-admin accounts. It does not provide access to their private folders.

To provide access to private folders of Administrator & other co-admin accounts, use the Content Manager view. Select all such folders that you want the co-admin account to access.

5 Alternatively, you can create a Box group and add to this group all such users who you want to grant preview access of all folders.

6 Use the co-admin user credentials to configure the Box account in Data Insight.

Viewing configured cloud sources

In the Management Console, you can view all the cloud storage accounts that Data Insight is configured to monitor.

Use the provided dynamic search filter to search for configured cloud storage accounts based on various pre-defined criteria, for example, the type of the cloud storage. You can also use the **Filter** field at the top of the content pane to filter the list of cloud storage based on the name of the account in addition to the pre-defined filter criteria. The displayed list of configured accounts changes automatically when you select the check box for a filter criteria. For instance, when you select a Collector node in the **By Collector** filter, Data Insight displays a list of accounts associated with the selected Collector node.

To view configured cloud storage accounts

1 In the Console, click **Settings** > **Cloud Sources**.

The screen displays the list of configured cloud storage accounts.

2 Review the following information about the cloud sources:

- The object ID of the cloud source. This numerical value is used to identify the cloud account when troubleshooting issues with it. This column is hidden by default. To view this column, click on the column header and select **Columns** > **ID**.

- The name of the cloud storage account. In the case of a Box account, the name represents a share in the Box account.
- The health of the cloud storage account.
- Whether file system event monitoring is enabled.
- The Collector node for the cloud storage account.
- The Indexer node associated with the cloud source.
- The scanning schedule for the cloud storage account. This column is hidden by default.

3 You can change Indexer for the cloud storage account.

See "About migrating storage devices across Indexers" on page 103.

To migrate to another Indexer, do the following:

- Click **Select Indexer**.
- The **Select Indexer** pop-up displays the configured indexers that are available for selection.
- Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. You can also view the number of devices, shares or site collections assigned to each configured node. It is recommended that you assign a node on which the resource utilization is low.
- Click **Select** to assign the Indexer to the filer.
  When the migration starts, click **View Details** to view the progress of the migration.

To review cloud storage account details

1 In the Console, click **Settings** > **Cloud Sources**.

2 Click the configured account that you want to review, or click the **Select Action** drop-down and select **View**.

The details screen appears.

To view events for the cloud storage account

1. In the Management Console, click **Settings** > **Cloud Source**.
2. Click the **Select Action** drop-down for the corresponding cloud source in the listing table, and select **Event Log**.

   The event log for that cloud account appears.
3. To download Data Insight logs for the cloud storage account for troubleshooting purposes, click the **Select Action** drop-down for the corresponding account, and select **Download Logs**.

   Data Insight downloads a compressed folder containing the logs related to this account from all relevant Data Insight servers.

   See "Downloading Data Insight logs" on page 383.

See "Managing cloud storage accounts" on page 274.

Configuring Box monitoring in Data Insight

Note: Data Insight requires you to install a separate add-on cloud license to monitor cloud storage resources. It disables addition of cloud storage sources if you do not have a valid base and cloud license or if your license has expired.

You must add the Box account to Data Insight configuration for monitoring.Veritas Data Insight treats every Box account as a separate device.

To add a Box account

1. In the Console, click **Settings** > **Cloud Sources**.

   The **Cloud Sources** list page displays the list of configured cloud storage accounts.
2. On the list page, click the **Add New Cloud Source** drop-down, and select **Box Account**.
3. On the **Add New Box Account** page, enter the properties for the Box account that you want to monitor, and click **Save**.

Add/Edit Box account

Use this dialog box to add a new Box account to Data Insight or to edit the configuration of an existing Box account.

Table 18-1 Add/Edit Box account options

Field	Description
Box account name	This is a free-form field. Enter a name that Data Insight uses to identify your box account. The name that you enter in this field represents a file share.
Collector	Do the following: **1** Click **Select Collector**. **2** The **Select Collector** pop-up displays the configured collectors that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Collector to the Box account. Data Insight connects to the Box application from this server.
Indexer	Do the following: **1** Click **Select Indexer**. **2** The **Select Indexer** pop-up displays the configured indexers that are available for selection. **3** Review the performance parameters such as disk and CPU utilization and backlog of files accumulating on the configured nodes. You can also view the number of devices, shares or site collections assigned to each configured node. It is recommended that you assign a node on which the resource utilization is low. **4** Click **Select** to assign the Indexer to the Box account. Events and metadata collected from the Box account are processed and stored on the Indexer node.

Table 18-1 Add/Edit Box account options *(continued)*

Field	Description
Authorize access	Do the following: **1** Click **Authorize** . This displays the Box authorization page. **2** Specify the Box administrator credentials that Data Insight uses to gain access to the Box account, and click **Authorize**. You can also use a co-admin account to allow Data Insight to access the Box account. See "Using a co-admin account to monitor Box resources" on page 269. **3** Click **Grant Access to Box** . This step creates an authorization token that is stored as a named credential in the Data Insightconfiguration. Data Insight can now access the user, folder and file metadata and the information about the activities performed on these file and folders. .
Enable Box account event monitoring	Select this check box to enable monitoring on the Box account.
Enable Box account scanning	Select the check box to enable Data Insight to scan the Box account using the Box Administrator credentials according to the specified schedule.
Scanning schedule for full scans	Select one of the following to define a scanning schedule for shares of this Box account: ■ Use Collector's default scanning schedule. ■ Use custom schedule. See "Custom schedule options" on page 235. Veritas Data Insight periodically scans the Box account to obtain file metadata and security descriptors. Each Collector worker node by default initiates a full scan of shares at 7:00 P.M. on the last Friday of each month.

Managing cloud storage accounts

On the cloud sources details page, you can view information about a configured storage account, edit the account's configuration, delete the account, and disable the account..

The **Configuration** tab on the filer details page displays the following information about a filer:

- The type of cloud source. At this time, Data Insight only supports monitoring of Box accounts.
- The name of the cloud storage account.
- The account credentials of the user authorized to log in to the filer.
- The IP address or hostname of the Collector worker node configured to scan the storage account.
- Whether event monitoring is enabled.
- The scanning schedule.

To review the account details

1. In the Console, click **Settings** > **Cloud Sources**.
2. Click the cloud storage account that you want to review.

   The details page for the cloud account appears.

You can edit the configuration of the filer.

To edit a cloud storage account

1. To edit the account configuration, on the account details page, click **Edit**.
2. On the **Edit Box** account page, make the necessary configuration changes.
3. Click **Save**.

To delete or disable a cloud storage account

1. To delete a cloud storage account, on the account details page click **Delete**.
2. Click **OK** on the confirmation message.
3. To disable a cloud storage account, click **Disable**, and click **OK** on the confirmation message.

Chapter 19

Renaming storage devices

This chapter includes the following topics:

- About renaming a storage device
- Viewing the device rename status
- Considerations for renaming a storage device

About renaming a storage device

Data Insight provides an option to conveniently change the name of a storage device such as a filer or a SharePoint web application from the Management Console without interfering with the underlying configurations.

Users with the role Server Administrator or Product Administrator with the administrative rights to the filer have the permission to rename a device. For other users, the device name field is not editable. Before renaming a device, it is recommended that you complete the ongoing migration, scan, or reports, and workflow processes configured for the device to be renamed. This ensures that the rename operation is completed without any disruption.

Once the device is renamed, the new or changed device name is automatically updated in the related databases, event logs, scan status, policy details, alerts, and reports. If the device is part of an Exclude Rule list, then the device name in the list is updated with the new name.

You can rename the filer and the SharePoint web application from the **Edit** page for the device.

See "Editing filer configuration" on page 236.

See "Editing web applications" on page 260.

Viewing the device rename status

When a device is renamed, Data Insight logs the status of the rename operation in the Data Insight event log and the webserver log file.

To access the event log, on the **Console** > **Settings** > **Inventory** > **SharePoint Web Application** or **Filers** > **Select Action** > **Event Log**.

You can filter the logs related to renaming operation, by setting the **Severity** to **Config**.

See "Location of Data Insight logs" on page 381. for information about the `webserver.log` file.

Considerations for renaming a storage device

The following considerations apply when renaming a storage device in the Console.

- The rename operation is supported for generic devices, EMC Celerra, NetApp 7-Mode, single node Veritas File System Server, Hitachi NAS, and SharePoint web applications.
- The rename operation is not supported on Box, EMC Isilon, Windows File Server, NetApp Cluster-Mode, and clustered Veritas File System Server.
- Once a device is renamed, the device references in the following are automatically updated:
  - Reports
  - Audit logs
  - Event log and event notification
  - Scan history and scan status
  - Policies
  - Alerts
  - Exclude rule list
  - Dashboard
  - Data Insight configurations
  - Enterprise Vault filer mapping configurations
- Certain special characters are supported while renaming SharePoint web applications. These consist of:

Special character	Symbol
Round brackets	( )
At sign	@
Slash	/
Colon	:
Hyphen	-
Dash	_

- Certain special characters are not supported while renaming a storage device. These consist of:

Special character	Symbol
Pipe	|
Apostrophe	'
Double quotes	" "
Caret	^
Angular brackets	< >
Comma	,
Curly brackets	{ }
Square brackets	[ ]

Section 4

Health and monitoring

- Chapter 20. Using Veritas Data Insight dashboards
- Chapter 21. Monitoring Data Insight

Chapter 20

Using Veritas Data Insight dashboards

This chapter includes the following topics:

- Viewing the system health overview
- Viewing the scanning overview
- Viewing the scan status of storage devices
- Viewing the scan history of storage devices

Viewing the system health overview

The **System Overview** dashboard is available under **Settings** > **Health & Monitoring** > **System Overview** provides a snapshot of the health of your entire environment.

Use the **System Overview** dashboard to do the following:

- Identify the issues that need immediate attention.
- Identify any global setting that is yet to be configured. For example, Data Insight flags if settings such as permission remediaton, Exclude Rules, SMTP, or Scanning and Event Monitoring are not configured.
- Review the count of inventory objects and their states.
- Check the scan status of all configured objects in your environment.
- Review the count of issues or potential issues with the configured product servers, directory servers, and storage devices.

The **System Overview** dashboard flags all potential issues. Use the link next to a notification or click the arrow icon next to a notification to navigate the relevant

object's list page. This enables you to view further details and to troubleshoot the issue.

The dashboard contains the following widgets:

Data Insight servers The Servers pie chart displays the graphical representation of the total number of Data Insight servers that are in the Faulted, At Risk, Unknown, and Healthy state.

The widget also provides an overview of the following:

- Inventory count of the number of Collectors, Indexers, Self-Service Portal nodes, and Windows File Server agents in your environment.
  Click a server type to view the filtered list for the selected server type on the **Settings** > **Data Insight Servers** page.
  Click **More details** to view the entire list of configured Data Insight servers.
- The notifications and warnings that indicate the reasons for the health status of the configured nodes.
  See "Managing Data Insight product servers" on page 67. for more information about the factors that affect the health of a node.
- Click **Add Server** to navigate to the **Add New Server** page.
  See "Adding a new Data Insight server" on page 66.

Scanning

The Scanning graph displays a color bar chart representing the number of scans during the last 7 days from the current date.

The color bar chart represents the different states of the scans- Failed [Red]; Successful [Green]; and Partially successful [Yellow].

On the Scanning chart, you can also view the following data:

- An overview of the consolidated scan status of all configured shares or site collections based on the latest scans that have taken place on these shares and site collections.
  Click on a consolidated status to view its details on the **Settings** > **Scan Status** > **Overview** tab.
- Notifications pertaining to the state of scanning and related alerts and warnings.
  For example, if scanning for a device is disabled, the dashboard displays the notification. It also provides you a link to navigate to the **Scanning and Event Monitoring** page to change the state of scanning.

Click **Global Scan Settings** to configure the scan settings.

See "Configuring scanning and event monitoring " on page 28.

See "Viewing the scanning overview" on page 283.

See "Viewing the scan status of storage devices" on page 285.

See "Viewing the scan history of storage devices" on page 286.

See "Viewing in-progress scans" on page 75.

Devices

The **Devices** pie chart provides a snapshot view of the number of storage devices in the Faulted, At Risk, Unknown, and Healthy states.

To the right of the chart, you can view the following data:

- Inventory of filers, web applications, and cloud sources that are being monitored by Data Insight.
  Click a device type to view the filtered list of configured devices of that type on the **Settings** > **Filers**, the **Settings** > **SharePoint Web applications**, or **Settings** > **Cloud Sources** page.
- Alerts associated with the storage devices that indicate the status, error conditions, and warnings that can help when troubleshooting issues with the device.
- Click **Add Device**, and select the device type to navigate to the configuration page for that device.
  See "Adding filers" on page 210.
  See "Adding web applications" on page 256.
  See "Configuring Box monitoring in Data Insight" on page 272.

Directory services

The directory services widget provides an inventory count of the configured directory services. The widget also displays information and alert notifications associated with the directory services.

Click **Add Directory Service** , and select the directory service type to navigate to the corresponding configuration page.

See "Adding a directory service domain to Data Insight" on page 109.

See "Viewing events " on page 288.

Viewing the scanning overview

The **Scan Status** dashboard displays the scanning statistics for the configured storage devices. It enables you to quickly and visually evaluate the status of the scans running in your environment. The **Scanning** dashboard displays the following charts that show the various aspects of the scan statistics:

Scan Status (Consolidated)

The pie chart provides a summary of the consolidated status of all scans on configured shares and site collections in percentage terms.

The consolidated status represents the combined status of full and incremental scans since the last full scan.

For example, a full scan successfully completes on a share. If there is subsequent failure of incremental scan on the same share, the consolidated status of scans for that share is displayed as **Failed**. Similarly, if a full scan on a share was partially successful, and all incremental scans succeed, the consolidated status stays partial. The consolidated status essentially indicates if any recent scans have failed for a share or a site collection.

Below the pie chart, you can view the summary of the total failed, successful, and partially successful scans on configured shares or site collections. Additionally, you may get a summary of scans with the warning status **Needs Attention**. You can view the detailed explanation for such warnings from the **Scan status** page for that share. The summary also displays the number of shares or site collections that have never been scanned till date.

Click on an area of the graph to view the filtered list of paths which have the selected consolidated status.

Scan History

The chart provides a graphical representation of the number of scans during a specified duration. Use the drop-down to the right of the graph to select the time period for which you want to view the summary.

The color bar chart represents the different states of the scans - Failed [Red]; Successful [Green]; and Partially successful [Orange]; for a more visual comparison of the data. In each bar, the number of failed, successful, and partially successful states are indicated by the length of the corresponding colors.

Click an area on the bar graph to view detailed information of the scans for the selected scan state on the **Scan History** sub-tab. For example, to display all successful scans, click on the green area of the bar graph. You can further narrow the list by using other filters on the **Scan History** sub-tab.

Age of Last Successful Scan

The pie chart provides a high-level overview of the age of last successful scan for a share or a site collection.

The age of last successful scan represents a consistent state of the scan metadata and comprises of a successful full scan followed by zero or more successful incremental scans. This state helps you decide if the scan metadata can be used for reporting purposes. For example, if the last successful scan for a share is 6 months old, it is recommended that you run a full scan for that share before you generate a report.

In summary, consolidated status indicates if there have been failed scans on a share, which helps you troubleshoot environmental issues. Whereas, the last successful scan indicates how recent and consistent the scan data is in the Data Insight index.

The chart is rendered for one of the following durations:

- More than 3 months old
- 3 months old
- One month old
- One week old
- One day old
- None

Clicking an area on the pie chart displays the **Scan Status** sub-tab with detailed information of the scans for the selected age. For example, clicking on the green area of the pie chart displays the shares or site collections which have a good scan state that is less than a day old.

Failed Scans in Last 24 Hours	Shows all the shares and site collections on which a scan has failed during the last 24 hours. You may view additional details for the cause of the scan failures such as type of scan that has failed, exit codes etc.

Viewing the scan status of storage devices

The **Scan Status** sub-tab of the **Settings** > **Scan Status** tab provides an overview of the status of the scans for all storage devices that are being monitored by Data Insight. The scans include the scans running on the configured file servers and SharePoint web applications.

Use the **Scan Status** page to do the following:

- Review status of the scans on configured storage devices. You can review the status of the last full and incremental scans and also list of paths for which the last scan failed.
- View the scan history.
- View the list of paths for which the last scan failed.
- Filter the list of scans.
- Start full scans.

Use the provided search filter, to search for scans based on various criteria, for example, the type of the device. Since it is a dynamic search, the displayed list of scans changes automatically as you select the check box for a filter criteria. For instance, when you select NetApp in the **Device Type** category, and Successful in the **Scan Status (Consolidated)** filter, the application displays a list of NetApp shares which have Successful as the consolidated status. Similarly, when you select a Collector node in the **By Collector** filter, Data Insight displays a list of scans running on devices associated with the selected Collector node.

To view the scan status

1. Click **Settings** > **Scan Status**.
2. The **Overview** sub- tab displays by default. Do one of the following:
   - Click on the appropriate region of the **Scan Status (Consolidated)** pie chart to filter the scans based on the status of the scans. Or, click on a region of the **Age of Last Successful Scan** pie chart to display the list of shares or site collections with the selected last successful scan.
   - Click the **Scan Status** sub-tab.

3 The **Scan Status** sub-tab displays the list of all the shares and site collections with colored icons to indicate their consolidated scan status. Use the following icons to understand the status:

- Green icon - scan successful.
- Orange icon - scan partial.
- Red icon - scan failed.
- Yellow warning icon - scan needs your attention.
- Gray icon - scan never happened.

4 Click on a status icon to know more details about the scan status. The scan status details pop-up window opens.

The **Scan Summary** tab displays possible causes, impacts, and solutions in case of scans whose status is either failed, partial, or needs attention.

To start a scan

◆ On the **Scan Status** page, click **Scan**.

Do one of the following:

- Click **Scan selected records** to scan the selected objects.
- Click **Scan all filtered records** to scan all objects that are displayed after applying a filter.

See "Viewing the scanning overview" on page 283.

See "Viewing the scan history of storage devices" on page 286.

Viewing the scan history of storage devices

Use the Scan History page to view the history of full and incremental scans of all storage devices that are being monitored by Data Insight.

Use the provided search filter in the left-side panel, to view the scans based on various criteria, for example, the type of entity or the status of the scans. Since it is a dynamic search, the displayed list of scans changes automatically as you select the check box for a filter criteria. For example, when you select CIFS in the **By Type** category and Failed in the **By Scan Status** category, the application displays a list of failed scans on all configured CIFS shares. You can also use the free-form Filter text box to enter the search criteria.

To view the scan history

1. Click **Settings** > **Scan Status**.
2. The **Overview** sub-tab displays by default. Do one of the following:
   - Click on the appropriate region of the **Scan History** chart to filter the scans based on the status of the scans.
   - On the **Scan Status** page, click the scan status icon to display the **Scan Summary** pop-up. Click the **Scan History** sub-tab
   - From the **Scan Status** page, click **Select Action** > **Scan Status**. Click the **Scan History** sub-tab.
   - Navigate to the share or site collection for which you want to view the scan history. See "Managing shares" on page 241. or See "Managing site collections" on page 263.

See "Viewing the scanning overview" on page 283.

See "Viewing the scan status of storage devices" on page 285.

Chapter 21

Monitoring Data Insight

This chapter includes the following topics:

- Viewing events
- About high availability notifications
- Monitoring the performance of Data Insight servers
- Configuring email notifications
- Enabling Windows event logging
- Viewing scan errors

Viewing events

You can monitor Veritas Data Insight recent system events on the **Events** page. The report displays entries for all system events. These events include the following information about an event:

- Time
- Severity
- Event summary
- Veritas Data Insight server where the event originated
- The user if any performing the action
- The object for which the event originated

To view system events

1. A list of recent system events appears.
2. You can choose to filter the events further using one or all of the following criteria:

- By time
- By any text appearing in the event summary
- By severity
- By the product server on which the event originates
  Enter the filter criteria in the relevant fields and click **Go**.

3 Click the Export icon at the bottom of the page to save the data to a `.csv` file.

About high availability notifications

Data Insight raises events for various conditions that might result in a loss of availability of a Data Insight system or component. Events are raised for the following conditions:

- Changes in the state of various essential services
- Saturation of the data volume
- Worker node misses heartbeat with the Management Server
- Accumulation of excessive files on the worker node
- Loss of connection between the filers and the Collector
- Excessive usage of CPU, memory, or disk space for extended period

See "Configuring advanced settings" on page 78.

Monitoring the performance of Data Insight servers

Data Insight enables you to review the performance of the remote Collector and Indexer nodes from the Management Console. You can analyze various performance parameters, such as the disk and CPU utilization and backlog of files accumulating on individual server nodes.

Processes running on the Collector node collect information about file and folder hierarchy and access events on storage devices, This information is stored in database files with appropriate timestamps, and is periodically processed and sent to the Management Server for the storage device.

Performance alerts are triggered when the resource utilization exceeds the threshold values for multiple scan cycles, and the backlog of files being processed on a node starts accumulating. The view enables you to decide whether to reassign or add a

Collector, or to make appropriate configuration changes on the **Advanced Settings** page for the server.

See "Configuring advanced settings" on page 78.

To view the performance of the Data Insight servers

1. On the console, click **Settings** > **Performance**.
2. The overview list page displays the configured Data Insight servers.
3. From the **Processing Backlog** drop-down, select the node or the folder for which you want to view the performance statistics.
4. Select the period for which you want to view the data.

   Data Insight displays the following statistics:

   - The host name of the server node.
   - The size and number of files accumulated on the server.
   - The CPU and memory consumption on the server.

Table 21-1 lists the type of files and their location on server nodes.

Table 21-1 Files and their location on Data Insight servers

File type	Server	Location	Cause
Event files	Collector	`<datadir>\collector`	Check the throughput rate of the `collector.exe` process and the number of collector threads.
Snapshot and audit files	Collector, Windows File Server	`outbox` folder on the Collector node or the Windows File Server	`FileTransferJob` is lagging. The `FileTransferJob` transfers the files from the `outbox`of folder on the Collector to the `inbox`o folder on the Indexer. Check the throughput statistics of `FileTransferJob.`

Table 21-1 Files and their location on Data Insight servers *(continued)*

File type	Server	Location	Cause
Snapshot and audit files	Indexer	`inbox` folder on the Indexer node	Check the throughput rate of `idxwriter.exe` process and the number of indexer threads. The throughput of the `idxwriter.exe` process in turn depends on the size of the batch being processed.

See "About the Health Audit report" on page 381.

Configuring email notifications

Data Insight provides email notifications for important events happening in the product. For example, CIFS scan failure or a directory scan failure. Notifications are sent out every 15 minutes, if new events are available. Email notifications are not enabled by default.

Note: Before you enable email notifications, you must enable configure the SMTP settings.

See "Configuring SMTP server settings " on page 23.

To configure email notifications

1. In the Management Console, click **Settings** > **Global Settings** > **Event Notifications**
2. On the Event Notifications page, select **Enable event notifications** checkbox.
3. In the Email recipients field, enter a comma separated list of email addresses to be notified.
4. Select the severity of events for which the email notifications must be sent.
5. Click **Save**.

Enabling Windows event logging

Veritas Data Insight can publish events to the Windows Event log. Events are published on the same machine where they originate. Event logging is enabled by default.

To configure Windows event logging

1. In the Management Console, click **Settings** > **Global Settings** > **Event Notifications**.
2. Select the **Enable Windows logging** check box.
3. Select the severity of events for which you want to enable Windows logging.
4. Click **Save**.

The events published in `DataInsightAlerts` log adhere to LEEF v 1.0 (Log Event Extended Format. LEEF is a customized event format for IBM QRadar. Here is a brief description of the event format.

Format:

```
{date} {hostname} LEEF:1.0|Veritas|DataInsight|{product_version}|{eventId}|
attr1=val1\tattr2=val2…
```

Where,

date : Alert time formatted as MMM dd HH:mm:ss

eventID : *<event ID>*

The following attributes are populated for each event:

cat	Data Insight policy type (PT_USERCENTRIC / PT_DATACENTRIC / PT_DATACENTRICWL / PT_DATACENTRICBL / PT_RT_DATACENTRIC)
policy	Data Insight policy name
devTime	Format of devTime (MMM dd yyyy HH:mm:ss)
sev	A number based on Data Insight policy severity (10: HIGH, 5: MEDIUM, 1: LOW).
usrName	User who violated the policy (FirstName LastName / user@domain / SID)
domain	Domain to which the violating user belongs.

The following attributes are specific to the type of Data Insight policy for which the event is generated:

User Activity Deviation Policy

access_count_day	Today's access count
access_count_week	This week's access count
access_count_avg_week	Avg. weekly access count
allowed_stddev	Allowed (standard) deviation
allowed_count_day	Allowed daily access count

Data Centric Policies (Data Activity, White List, Black List)

access_count	Today's access count.
access_count_details	Detailed access count.
allowed_count_day	Allowed daily access count.

Real Time Sensitive Data Policy

dlp_policies	Data Loss Prevention (DLP) policies violated by the path that was accessed.
access_types	Type of access made.
access_count	Access count
resource	Access path
src	IP addresses

If you do not want certain attributes to be reported into the Event log for privacy or security reasons, the following global attribute can be defined with comma separated list of attributes to be excluded:

```
leef.exclude.attrs=comma separated attr names Global attribute can
be set using following command on MS: configdb.exe -O -J
leef.exclude.attrs -j <attr1,attt2…>
```

Viewing scan errors

You can view a list of all the paths on which a scan has failed. In the **Workspace** tab tree-view panel, the folder icon displays a red cross mark for the paths on which a scan has failed. The scan errors displayed are from the latest scan completed on the share.

To view scan errors

1. Do one of the following to view the details of the scan errors on paths:
   - In the Management Console, click **Settings** > **Scanning**.
     The **Overview** sub-tab of the Scanning dashboard displays the paths on which scans have failed in the last 24 hours.
   - On the **Scan Status** sub-tab of the **Scanning** dashboard, click the **Select Action** drop-down corresponding to a path and select **Scan Errors**.
   - Navigate to a share or a site collection on the **Monitored Shares**/**Monitored Site Collections** sub-tab on the filer or SharePoint web applications details page., click the **Select Action** drop-down corresponding to a path, and select **Scan Errors** to view the failed scans on that path.
2. On the Scan Errors page, review the time of the error, the error code, and the possible cause of the error.

Section 5

Alerts and policies

- Chapter 22. Configuring policies

Chapter 22

Configuring policies

This chapter includes the following topics:

- About Data Insight policies
- Managing policies
- Managing alerts

About Data Insight policies

A policy is a set of conditions that you configure to monitor access events on files and folders stored on various repositories. Veritas Data Insight policies help you detect the sources of threat, access patterns on sensitive data, and anomalous user behavior. Data Insight receives information about sensitive files from Symantec Data Loss Protection (DLP). You can also import sensitive file information in to Data Insight using a CSV file.

Policies must include at least one condition that is configured to detect abnormal access patterns or user behavior. Data Insight generates an alert whenever it detects any violation of a condition in a configured policy. Additionally, you can also configure Data Insight to send email alerts to a select group of recipients whenever any policy violation is reported. The Real-time Sensitive Data Activity Policy is evaluated instantaneously whenever there is a violation. All other policies are evaluated by default at 12:00 A.M. every night.

Policies can be configured with three severities, namely, high, medium, and low. You can assign the severity level to a policy based on your organizational needs. For example, the Information Security team can define policies to monitor accesses on the share `\Finance`. For this purpose, they can configure a policy with a medium severity to monitor accesses on folders containing Finance policies and guidelines files. Whereas, they can configure a policy with a high severity to monitor accesses on files containing payroll information. When an alert is generated for a policy violation, the severity of the policy is associated with the alert.

Data Insight comes packaged with the following out-of-the-box policies that you can configure according to your needs:

- Data Activity Trigger policy
  Use this policy to define the maximum cumulative count of the meta operations on the selected paths. For example, if you have defined the maximum accesses per day as 500 on the share `\\netapp1\finshare`, and the total access count by the active set of users exceeds 500, then Data Insight generates an alert.
- User Activity Deviation policy
  Use this policy to define the threshold of deviation from the baseline activity. The baseline activity on a file or folder is the average number of accesses that are considered normal based on past access counts. If the activity, by the selected users, on the selected data exceeds the specified threshold of the baseline (for example, three standard deviations above the baseline activity), and the maximum accesses allowed per day, Data Insight generates an alert.
  You can configure how many standard deviations a user is allowed to deviate from the defined baseline.
- Data Activity User Whitelist-based policy
  Use this policy to define a whitelist of users based on the Active Directory custom attributes, who can access selected shares or paths. Also, you can create such a policy with multiple conditions with multiple values for the same custom attributes.
  If users, other than those defined in the whitelist, access selected data, Data Insight generates an alert.
- Data Activity User Blacklist-based policy
  Use this policy to define a blacklist of users based on the Active Directory custom attributes, who should not be accessing selected shares or paths. Also, you can create such a policy with multiple conditions with multiple values for the same custom attributes.
  If users, who are included in the blacklist, access selected data, Data Insight generates an alert.
- Real-time Sensitive Data Activity Policy
  Use this policy to trigger real-time alerts when a selected set of users perform any access events on the paths that violate configured DLP policies. Whenever the events that violate Real-time Sensitive Data Activity Policy are processed by a Collector node, alerts are generated and an email is sent to a configured list of recipients. The policy violations are also published in the Windows Applications and Services Logs as DataInsightAlerts events.

Managing policies

You can view, edit and delete configured policies, and add new policies to Data Insight from the **Policies** tab.

To manage policies

1. In the Console, click the **Policies** tab.

   The left pane displays the default policy groups.

2. Click a policy group.

   The policy listing page displays the configured policies for that policy group.

3. To edit an existing policy, from the Actions drop-down, click **Edit**.
4. To delete a policy, select the corresponding check box and click **Delete**.

To add a new policy

1. In the Console, click the **Policies** tab.

   The left pane displays the default policy groups.

2. Click the policy group that you want to base your policy on.
3. On the policy listing page, click **Add new policy**. Or in the tree-view panel, right-click the policy type, and select **Add**.
4. On the **Add new policy** page, click each tab and enter the relevant information.
5. Click **Save**.

See "Create Data Activity Trigger policy options" on page 299.

See "Create User Activity Deviation policy options" on page 301.

See "Create Data Activity User Whitelist-based policy options" on page 303.

See "Create Data Activity User Blacklist-based policy options" on page 305.

See "Create Real-time Sensitive Data Activity policy options" on page 308.

By default, all policies except Real-time Sensitive Data Activity policy, are evaluated at 12:00 A.M. every night. You can schedule policies to be evaluated more frequently for proof-of-concept (POC) setups. Note that a schedule that is too aggressive can put excessive load on the Indexer.

You can set a custom schedule to evaluate policies from the **Settings** tab. The schedule must be specified in the cron format.

To set a custom schedule for policies

1. Click **Settings** > **Data Insight Servers**.
2. Click the entry for the Management Server.

3 On the page for the Management Server node, click **Advanced Settings**.

4 Click **Edit**.

5 Scroll to bottom of the page and expand the **Set custom properties** section. Specify property name to be job.PolicyJob.cron and property value to be the new schedule. Schedule needs to be specified in cron format

6 In the Property name field, enter job.PolicyJob.cron.

7 In the Property value fields, enter the values as follows:

To evaluate values every N minutes, specify value as 0 0/N * * * ? *.	For example, to evaluate policies every 10 minutes, specify value as 0 0/10 * * * ? *.
To evaluate policies every N hours, specify value as 0 0 0/N * * ? *.	For example, to evaluate policies every two hours, specify value as 0 0 0/2 * * ? *.

Create Data Activity Trigger policy options

Use this dialog to create a new Data Activity Trigger policy. Options that are selected in the respective tabs are displayed in the **Summary** panel on the right of the page.

Table 22-1 Create Data Activity Trigger policy options

Option	Description
Policy Information	Enter information in the following fields: ■ **Name** - The name of the policy. ■ **Description** - A short description of the policy. ■ **Policy Type** - Data Activity Trigger is selected by default. ■ **Severity** - The severity of the policy. From the drop-down, select High, Medium, or Low. The severity level associated with the policy helps you decide the possible course of action when an event that matches the policy occurs. Select the **Enable Policy** check box to enforce the policy. The policy is not evaluated if the check box is not selected.

Table 22-1 Create Data Activity Trigger policy options *(continued)*

Option	Description
Configure Policy	Select the following conditions to configure the policy: ■ **Select Activity** - Select the type of accesses to be monitored on the selected data set. Select the **Meta Access** radio button to monitor only the high-level access events that Data Insight maps from the detailed file system and SharePoint access events. Select the **Detailed Access** radio button to monitor specific file system and SharePoint access events. ■ **Additional Condition** - From the **Minimum accesses per day for alerts** drop-down, select the minimum number of accesses on the selected data set on that day that are required to trigger an alert.
Data Selection	Do the following to select the resources: **1** Select the **Physical Hierarchy** radio button to view the configured file servers or SharePoint web applications. Or select the **DFS Hierarchy** radio button to view the configured DFS paths in a domain. **2** Click a path to select it. Or you can use a `.csv` file with information about the paths that you want to apply the policy to. Click **Browse** to navigate to the location of the `.csv` file, and click **Upload**. **3** To limit the scope of the files to be monitored, select **Select all files in folder** or **Select only sensitive files**. If you select the **Select all files in folder** option, accesses on all files in the folder are evaluated for determining any violation of the policy. If you select the **Select only sensitive files** option, accesses on only the sensitive files in the folder are evaluated for determining any violation of the policy. **Note:** The list of sensitive files is obtained from Symantec Data Loss Prevention or is imported into Data Insight using a CSV file. See "Configuring Symantec Data Loss Prevention settings" on page 44. **4** The selected data set is listed in the **Selected resources** pane.

Table 22-1 Create Data Activity Trigger policy options *(continued)*

Option	Description
Notification	Enter one or more specific email addresses for people to whom you want to send alerts that are generated for the policy. Select the **Notify custodians** check box to send email alerts to the custodians of the selected paths along with other email addresses on the notification list.

Create User Activity Deviation policy options

Use this dialog to create a new User Activity Deviation policy. Options that are selected in the respective tabs are displayed in the **Summary** panel on the right of the page.

Table 22-2 Create User Activity Deviation policy options

Option	Description
Policy Information	Enter the following information: ■ **Name** - The name of the policy. ■ **Description** - A short description of the policy. ■ **Type** - User Activity Deviation is selected by default. ■ **Severity** - The severity of the policy. The severity level associated with the policy helps you decide the possible course of action when an event that matches the policy occurs. Select the **Enable Policy** check box to enforce the policy. The policy is not evaluated if the check box is not selected.

Table 22-2 Create User Activity Deviation policy options *(continued)*

Option	Description
Configure Policy	Do the following: **1** From the drop-down, select the time range for the baseline activity. Baseline activity is then computed as the average access in that time range. From the **Threshold Configuration** drop-down, select the threshold of normal activity. The threshold is the acceptable number of standard deviations that a user is allowed to deviate. Accesses above the defined threshold trigger an alert. **2** **Additional Condition** - From the drop-down, select the minimum accesses per day per user. Alerts are raised only if the total accesses exceed the minimum value specified. This prevents Data Insight from raising too many alerts when baselines are very low.
Data Selection	Do the following to select the resources: **1** Select the **Physical Hierarchy** radio button to view the configured file servers or SharePoint web applications. **2** Click a path to select it. Or you can use a `.csv` file with information about the paths that you want to apply the policy to. Click **Browse** to navigate to the location of the `.csv` file, and click **Upload**. The selected data set is listed in the **Selected resources** pane.
User Selection	Do the following: **1** Select the **Users** or **Group** radio button to view configured users or groups respectively. The list of users and groups is sorted alphabetically. Click the star icon to display all the configured users or groups. You can use the Domain filter search bar to filter users or groups according to domains. You can also filter the users according to their Active Directory custom attributes. **2** Click a user or group to select it.

Table 22-2 Create User Activity Deviation policy options *(continued)*

Option	Description
Notification	Enter one or more specific email addresses for people to whom you want to send the alerts that are generated for the policy.

Create Data Activity User Whitelist-based policy options

Use this dialog to create a new Data Activity User Whitelist-based policy. Options selected in the respective tabs are displayed in the **Summary** panel on the right of the page.

Table 22-3 Create Data Activity User Whitelist-based policy options

Option	Description
Policy Information	Enter the following information: ■ **Name** - The name of the policy. ■ **Description** - A short description of the policy. ■ **Type** - Data Activity User Whitelist-based is selected by default. ■ **Severity** - The severity of the policy. The severity level associated with the policy helps you decide the possible course of action when an event that matches the policy occurs. Select the **Enable Policy** check box to enforce the policy. The policy is not evaluated if the check box is not selected.
Configure Policy	**Select Activity** - Select the type of accesses to be monitored on the selected data set. Select the **Meta Access** radio button to monitor only the high-level access events that Data Insight maps from the detailed file system and SharePoint access events. Select the **Detailed Access** radio button to monitor specific file system and SharePoint access events.

Table 22-3 Create Data Activity User Whitelist-based policy options *(continued)*

Option	Description
Data Selection	Do the following to select the resources: **1** Select the **Physical Hierarchy** radio button to view the configured file servers or SharePoint web applications. Or select the **DFS Hierarchy** radio button to view the configured DFS paths in a domain. 2 **2** Click a path to select it. Or you can use a `.csv` file with information about the paths that you want to apply the policy to. Click **Browse** to navigate to the location of the `.csv` file, and click **Upload**. **3** To limit the scope of the files to be monitored, select **Select all files in folder** or **Select only sensitive files** If you select the **Select all files in folder** option, accesses on all files in the folder are evaluated for determining any violation of the policy. If you select the **Select only sensitive files** option, accesses on only the sensitive files in the folder are evaluated for determining any violation of the policy. Alternatively, select **All Physical Resources**, to select all data resources, except DFS paths configured in Data Insight. All Physical Resources evaluates all the sensitive files, not all the files. The selected data set is listed in the **Selected resources** pane. **Note:** Data Insight obtains information about sensitive files from Symantec Data Loss Prevention (DLP) or is imported into Data Insight using a CSV file. See "Configuring Symantec Data Loss Prevention settings" on page 44.

Table 22-3 Create Data Activity User Whitelist-based policy options *(continued)*

Option	Description
User attribute query	Do the following: **1** Click **Add Condition**. **2** From each of the drop-down menu, select the criteria to build the query. The query is used to select users based on their Active Directory custom attributes. You can add multiple conditions to a Data Activity User Whitelist-based policy. You can add multiple conditions to a Data Activity User Blacklist-based policy. In case of multiple conditions with same attributes, the conditions are evaluated with a logical OR operation. In case of multiple conditions with different attributes, the conditions are evaluated with a logical AND operation.
Notifications	Enter one or more specific email addresses for people to whom you want to send alerts that are generated for the policy. Select the **Notify custodians** check box to send email alerts to the custodians of the selected paths along with other email addresses on the notification list.

See "Configuring Symantec Data Loss Prevention settings" on page 44.

Create Data Activity User Blacklist-based policy options

Use this dialog to create a new Data Activity User Blacklist-based policy. Options selected in the respective tabs are displayed in the **Summary** panel on the right of the page.

Table 22-4 Create Data Activity User Blacklist-based policy options

Option	Description
Policy Information	Enter the following information: ■ **Name** - The name of the policy. ■ **Description** - A short description of the policy. ■ **Type** - Data Activity User Blacklist-based is selected by default. ■ **Severity** - The severity of the policy. The severity level associated with the policy helps you decide the possible course of action when an event that matches the policy occurs. Select the **Enable Policy** check box to enforce the policy. The policy is not evaluated if the check box is not selected.
Configure Policy	**Select Activity** - Select the type of accesses to be monitored on the selected data set. Select the **Meta Access** radio button to monitor only the high-level access events that Data Insight maps from the detailed file system and SharePoint access events. Select the **Detailed Access** radio button to monitor specific file system and SharePoint access events.

Table 22-4 Create Data Activity User Blacklist-based policy options *(continued)*

Option	Description
Data Selection	Do the following to select the resources: **1** Select the **Physical Hierarchy** radio button to view the configured file servers or SharePoint Web applications. Or select the **DFS Hierarchy** radio button to view the configured DFS paths in a domain. 2 **2** Click a path to select it. Or you can use a `.csv` file with information about the paths that you want to apply the policy to. Click **Browse** to navigate to the location of the `.csv` file, and click **Upload**. **3** To limit the scope of the files to be monitored, select **Select all files in folder** or **Select only sensitive files** If you select the **Select all files in folder** option, accesses on all files in the folder are evaluated for determining any violation of the policy. If you select the **Select only sensitive files** option, accesses on only the sensitive files in the folder are evaluated for determining any violation of the policy. Alternatively, select **All Physical Resources** to select all data resources, except DFS paths configured in Data Insight. All Physical Resources evaluates all the sensitive files, not all the files. The selected data set is listed in the **Selected resources** pane. **Note:** Data Insight obtains information about sensitive files from Symantec Data Loss Prevention (DLP) or is imported into Data Insight using a CSV file. See "Configuring Symantec Data Loss Prevention settings" on page 44.

Table 22-4 Create Data Activity User Blacklist-based policy options *(continued)*

Option	Description
User attribute query	Do the following: **1** Click **Add Condition**. **2** From each of the drop-down menu, select the criteria to build the query. The query is used to select users based on their Active Directory custom attributes. You can add multiple conditions to a Data Activity User Blacklist-based policy. In case of multiple conditions with same attributes, the conditions are evaluated with a logical OR operation. In case of multiple conditions with different attributes, the conditions are evaluated with a logical AND operation.
Notifications	Enter one or more specific email addresses for people to whom you want to send alerts that are generated for the policy. Select the **Notify custodians** check box to send email alerts to the custodians of the selected paths along with other email addresses on the notification list.

See "Configuring Symantec Data Loss Prevention settings" on page 44.

Create Real-time Sensitive Data Activity policy options

Use this dialog to create a new Real-time Sensitive Data Activity policy. Options selected in the respective tabs are displayed in the **Summary** panel on the right of the page.

Table 22-5 Create Real-time Sensitive Data Activity policy options

Option	Description
Policy Information	Enter the following information: ■ **Name** - The name of the policy. ■ **Description** - A short description of the policy. ■ **Type** - Real-time Sensitive Data Activity policy is selected by default. ■ **Severity** - The severity of the policy. The severity level associated with the policy helps you decide the possible course of action when an event that matches the policy occurs. Select the **Enable Policy** check box to enforce the policy. The policy is not evaluated if the check box is not selected.
Configure Policy	**Select Activity** - Select the type of accesses to be monitored on the selected data set. Select the **Meta Access** radio button to monitor only the high-level access events that Data Insight maps from the detailed file system and SharePoint access events. Select the **Detailed Access** radio button to monitor specific file system and SharePoint access events. From the **DLP Policies violated** drop-down list, select a set of DLP policies to be considered for alerting.
User Selection	Use this tab to directly select the users. Do the following: **1** Select the **Users** or **Group** radio button to view configured users or groups respectively. The list of users and groups is sorted alphabetically. You can use the Domain filter search bar to filter users or groups according to domains. You can also filter the users according to their Active Directory custom attributes. **2** Click a user or group to select it.

Table 22-5 Create Real-time Sensitive Data Activity policy options *(continued)*

Option	Description
User Selection Using Attributes	You can also select users by using the attribute query. Do the following: **1** Click **Add Condition**. **2** From each drop-down menu, select the criteria to build the query. The query is used to select users based on their Active Directory custom attributes. You can add multiple conditions. For evaluating a query, Data Insight uses the logical AND operation between multiple conditions.
User Exclusion	Select the users you want to exclude from the policy. Any activity performed by these users will be ignored and will not be used to trigger an alert. Use this tab to directly select the excluded users. Do the following: **1** Select the **Users** or **Group** radio button to view configured users or groups respectively. The list of users and groups is sorted alphabetically. You can use the Domain filter search bar to filter users or groups according to domains. You can also filter the users according to their Active Directory custom attributes. **2** Click a user or group to select it.
User Exclusion Using Attributes	You can also exclude users by using the attribute query. Do the following: **1** Click **Add Condition**. **2** From each of the drop-down menu, select the criteria to build the query. The query is used to select users based on their Active Directory custom attributes. You can add multiple conditions. For evaluating a query, Data Insight uses the logical AND operation between multiple conditions.
Notification	Enter one or more specific email addresses for people to whom you want to send the alerts that are generated for the policy.

Managing alerts

An alert is a signal generated by a policy when the condition specified in the policy is violated.

You can view alerts on the **Alerts** tab on the Management Console.

To manage alerts

1. In the Console, click the **Policies** tab.

   The **Alerts** tab display by default. On the tab, you can view all the alerts that were generated by Data Insight.

2. In the Alerts Summary, click the drop-down arrow on any column header and select Columns. Then, select the parameters you want to show or hide. You can sort by:

   - The name of the policy.
   - The severity of the alert.
   - The type of policy associated with the alert - Data Activity Trigger, User Activity Deviation, Data Activity User Whitelist-based, Data Activity User Blacklist-based, or Real-time Sensitive Data Activity.
   - The name of the user account that violated the policy.
   - The date on which the alert was generated.
   - The resolution, if any, taken in response to the alert.

3. To send alerts in email, select the alerts and click **Send Email**.
4. Enter the email addresses and click **Send**.
5. To enter the resolution for an alert, select the alert, click in the Resolution column for the alert and type in the resolution.

   To update the resolution for multiple alerts, select the alerts and click **Update Resolution** at the top of the summary table.

The events published in the `DataInsightAlerts` log adhere to LEEF v 1.0 (Log Event Extended Format. LEEF is a customized event format for IBM QRadar.

See “Enabling Windows event logging” on page 292.

To delete alerts

- To delete an alert, select an alert and click **Delete**.

  To delete alerts by severity, click Delete and select the severity. This deletes all alerts that match the selected severity.

  To delete alerts older than a certain date, click Delete and select the date at the top of the table.

  Note: You can configure automatic deletion of alerts older than the specified interval on the Data Retention screen. However, you cannot restore the alerts once they are deleted. Alerts are also automatically published to the Windows event log.

  See "Configuring data retention settings" on page 38.

Review the Risk Dossier to know the alerts raised in response to the policies violated by a specific user.

Section 6

Remediation

- Chapter 23. Configuring remediation settings

Chapter 23

Configuring remediation settings

This chapter includes the following topics:

- About configuring permission remediation
- About managing data
- About configuring archive options for Enterprise Vault
- Viewing and managing the status of an operation

About configuring permission remediation

Data Insight provides permission recommendations on paths based on activity on the paths. To fine tune these recommendations and take action on the recommendations, the Data Insight administrator must enable and configure permission remediation. Depending on your organization process, you can configure Data Insight in any of the following two ways:

- Raise a remediation ticket.
  Configure when you have a ticketing system, which can assign tickets to relevant stakeholders. The recipients are responsible for the actual implementation of the recommendation. You can create remediation tickets by using one of the following two ways:
  - By sending an email to the ticketing system.
    Use this option if you have a ticketing system which can create a ticket by reading an email.
  - By executing a custom script.
    Use this option if you have an alternate mechanism of creating a ticket.

The scripts can be created in the *.exe*, *.bat*, *.pl*, or *.vbs* formats.

- Apply changes by invoking custom scripts.
  Configure custom scripts to enable Data Insight to directly apply the recommended changes.
  You can use the following types as custom scripts,*.exe*, *.bat*, *.pl*, or *.vbs*.

You can view the status of remediation actions on the **Settings** > **Action Status** tab of the Data Insight Management Console.

For information about custom scripts, see the *Veritas Data Insight Programmer's Reference Guide*.

See "Viewing and managing the status of an operation" on page 323.

Managing and configuring permission remediation

You can configure Data Insight to handle the way it implements the recommended permission changes. You must have a Data Insight Server Administrator role to be able to configure the permission settings.

To enable permission remediation

1. From the Data Insight Management Console, click **Settings** > **Permissions**. The **Remediation** sub-tab opens by default.
2. Click **Edit**. The page expands to display the configuration for permission remediation.
3. Select **Enable Permission Remediation**.

To configure remediation for raising a ticket

1. From the Data Insight Management Console, click **Settings** > **Permissions**. The Remediation sub-tab opens by default.
2. Click **Edit**. The page expands to display the configuration for permission remediation.

   Select **Enable Permission Remediation**, if it is not already enabled.
3. Select **Raise a ticket**. The panel expands to display the configuration details.
4. Select either of the two options:
   - **Send email** - Select to configure settings for an email-based ticketing system.
   - **Use custom action** - Select to configure settings for a non-email based ticketing system.
5. If you selected the **Send email** option, provide the relevant information in the email template:

- The email ID of the sender
- The email IDs of the recipients
- The email IDs of other recipients
- The subject line
- The header line showing priority and the queue status. The headers can be set to have custom information sent to the ticketing/request systems. For example, you can set priority=high, assign_to=permission_queue etc.
- The body of the email. You can use the default variables to enter relevant text. The variables are evaluated during run-time and are replaced by their corresponding values. Currently Data Insight provides the following dynamic variables:
  - Recomendation_text
    The value of this variable is the recommendations generated by Data Insight.
    For information about reviewing permission recommendations, see the Veritas Data Insight User's Guide.
  - Requester_name
    The value of this variable is the user who accepted the Data Insight recommended changes.
  - Action_id
    Data Insight generates the value of this variable. It is a unique identifier for the operation.
  - Todays_date
    The value of this variable is the system date.

6 If you selected **Use custom action**, do the following:
  - Create a custom script by following the guidelines documented in the *Veritas Data Insight Programmer's Reference Guide*.
  - Save the script at the location:
    `$datadir\conf\workflow\steps\permission_remediation\ticketing`.

7 In the **Enter the command to be executed** field, provide the file name of the saved script.

8 Select the relevant saved credential if your system needs to run the script using the specified credentials. The script runs with the Local System account credentials, however network calls made by the script will impersonate the specified user credential.

9 Click **Save**.

To configure the process of applying recommendations

1. Write the relevant scripts to handle changes to the following:
   - The Active Directory.
   - CIFS permissions.

   For more information about the custom scripts refer to the *Veritas Data Insight Programmer's Reference Guide*.
2. Save the scripts in the following locations:
   - For changes to Active Directory - `$DATADIR\conf\workflow\steps\permission_remediation\AD`
   - For changes to CIFS permissions - `$DATADIR\conf\workflow\steps\permission_remediation\CIFS`
3. From the Data Insight Management Console, click **Settings** > **Permissions**. The **Remediation** sub-tab opens by default.
4. Click **Edit**. The page expands to display the configuration for permission remediation.
5. Select **Enable Permission Remediation** if it is not already enabled.
6. Select **Remediate using custom scripts**. The panel expands to show you the configuration details.
7. In the **Enter the command to be executed** field, specify the file name of the custom script(s) that you have created in step 1
8. Click **Save**.

   The saved scripts are used to handle the permission remediation actions after you accept the permissions recommendations displayed on the **Workspace** tab.

For information on reviewing recommendations and initiating the process of applying them, see the *Data Insight User's Guide*.

Configuring exclusions for permission recommendation

You can specify the users and groups that you want to exclude from the purview of the permission recommendation. Once you exclude a user or group, Data Insight does not consider that group for presenting a recommendation for permission changes. You can exclude a groups by directly selecting it. You can also exclude large groups which have more than the specified number of users.

To exclude a user or group from remediation

1. From the Data Insight Management Console, click **Settings** > **>Permissions**. The **Remediation** tab opens by default.
2. Click **Recommendation**.
3. Do the following:
   - To exclude groups with more than a certain number of users, specify the value in the space provided.
     To exclude specific groups, in the **Exclude following groups from recommendations** pane, click the group's name which you want to exclude.
   - To exclude a specific user, in the **Available Members** pane, select the user.

   You can use the name filters and domain filters to view and sort the available user groups.

   The users and groups that you select are displayed in the **Exclusion List** pane.
4. Click **Save**.

About managing data

The storage devices in your environment may accumulate data that is orphan or has not been accessed for a long time. A large amount of such data on your storage devices can consume valuable storage space. Data Insight enables you to reclaim storage space occupied by inactive data. You can manage the inactive data directly from the Data Insight Management Console in the following ways:

- Archive data using Veritas Enterprise Vault.
  See "About configuring archive options for Enterprise Vault" on page 318.
- Manage data by invoking custom scripts.

About configuring archive options for Enterprise Vault

You can handle archiving and maintenance of the data that is stored on network shares using Veritas Enterprise Vault™ . This feature enables you to identify old and inactive data residing on storage devices and to archive it directly from the Data Insight Management Console. Data Insight uses the Enterprise Vault File System Archiving (FSA) API to archive. When Data Insight sends the command to Enterprise Vault to archive a file, Enterprise Vault uses the account associated with

the FSA task to archive the file. To perform an archive operation on a file, the account must have relevant read-write permissions on the file being archived.

To ensure proper functioning of the archive operations, you must have Microsoft .NET Framework 4.5 running on the Data Insight Management Server. Ensure that you restart the Management Server after you have finished installing the Microsoft .NET Framework 4.5.

Enterprise Vault (EV) must be successfully configured before you can archive data from the Data Insight Management Console. To configure Enterprise Vault complete the following tasks:

- Run the Enterprise Vault Configuration Wizard to create Enterprise Vault Directory Database and Enterprise Vault site.
- Add indexing locations.
- Add Vault Store Groups and Vault Store Partitions to store archived items. When configuring a Vault Store, if it makes sense for your organization, set the **Remove safety copies** option to **Immediately after archive**. Selecting this option ensures that the post-processing actions, such as creation of a placeholder or deletion of file are performed immediately after the file is archived. This is especially useful during Testing and Proof of Concept deployments.
- Modify the File System Archiving (FSA) task properties, if necessary.
- Add the necessary retention categories.
- Create archiving policies or modify the Default FSA Volume Policy and Default FSA Folder Policy, if necessary.
- Add archiving targets for NetApp filers and Windows file servers. When adding a Windows File Server as a target, select the option to install the placeholder service on the file server.
- Add volumes and folders for the targets, and create the necessary archive points.

For detailed instructions on completing the Enterprise Vault configuration tasks, see the Veritas Enterprise Vault documentation.

Note: If there are multiple Enterprise Vault servers in an Enterprise Vault site, then only one of the Enterprise Vault servers in the site must be added to the Data Insight configuration. If any of the Enterprise Vault servers is down, archiving of files from the shares that use vault stores that are managed by that Enterprise Vault server fails. For information about Enterprise Vault sites and vault stores, see the *Veritas Enterprise Vault™ Introduction and Planning Guide*.

For instructions on initiating archive requests, see the *Veritas Data Insight User's Guide*.

Adding new Enterprise Vault servers

You can configure Enterprise Vault servers to archive inactive data directly from the Data Insight Management Console.

To add a new Enterprise Vault server

1. In the Data Insight Management Console, click **Settings** > **Data Management**.

   The **Archiving (Enterprise Vault Configuration)** tab opens by default. It displays a list of servers which are already configured.

2. Click **Add New EV Server**.

3. In the **New EV Server** window, provide the following information:

   - The IP address of the Enterprise Vault server.
   - The port at which the Enterprise Vault server runs.
   - The relevant login credentials for the Enterprise Vault server. From the **Login Credentials** drop-down, select the credentials which are saved in the system.See "Managing saved credentials " on page 105.

   Note:

   Data Insight does not support addition of an Enterprise Vault server using the host name. You can only add an Enterprise Vault server using the IP address or the alias name. Ensure that an entry for the server's alias name is added to the hosts file on the Management Server.

4. Click **Test Credentials** to verify that Data Insight can connect to the server using the saved credentials.

5. Click **Save**.

Managing Enterprise Vault servers

You can do the following tasks on the **Data Management** page:

- Review the configured Enterprise Vault servers
- Add a new Enterprise Vault server.
  See "Adding new Enterprise Vault servers" on page 320.
- Edit the configuration of an Enterprise Vault server.
- Configure the archive options.
- Configure the schedule to pause an archive operation.

To view and manage the existing Enterprise Vault servers

1. From the Data Insight Management Console, click **Settings** > **Data Management**.

   The **Archiving (Enterprise Vault Configuration)** tab opens by default. It displays a list of configured servers.

2. Select the server for which you want to edit the configuration, and from the **Actions** drop-down, select **Edit**.

3. In the **Edit EV Server** dialog, make necessary changes.

4. Click **Save**.

You can configure additional options, such as the total size and number of the files and folders that can be archived in one archive request.

To configure archive options for Enterprise Vault servers

1. From the Data Insight Management Console, click **Settings** > **Data Management**.

   The **Archiving (Enterprise Vault Configuration)** tab opens by default. It displays a list of configured servers.

2. On the **Archive Options** panel, specify the preferred batch size in MB(s). When archiving files to Enterprise Vault, the batch of files sent to Enterprise Vault in one call does not exceed the given size.

3. Enter the number of files that you want to archive in one operation. By default, you can archive 50 files in one archive request.

4. Click **Save**.

Note: The batch size has a higher priority than the file count for deciding the list of files in an archive operation. Thus, Data Insight limits files in the archive operation after the batch size limit is reached, even if the file count does not exceed the specified limit.

You can configure a pause window for the Enterprise Vault operations by scheduling Data Insight to pause all the archive activities during a specific duration of time. When the pause occurs, Data Insight submits no more new archive requests. It places all new requests in a queue and executes them after the pause window.

To configure a pause schedule for the Enterprise Vault operations

1. From the Data Insight Management Console, click **Settings** > **Data Management**.
2. On the **Pause Workflow Schedule** panel, select **Pause workflow for specific times**. Data Insight displays a list for all the previously configured pause schedules.
3. Do any of the following:
   - To change an existing pause schedule, click the name of the schedule and click **Edit**
   - To add a new pause schedule, click **Add**.
4. Provide the following information:
   - The start time of the pause.
   - The end time of the pause.
   - The days of the week for which you want to schedule the pause.
5. Click **Save**.

Mapping file server host names

A filer that is assigned a specific host name in Data Insight could be assigned a different host name in an Enterprise Vault server. To resolve this conflict, you must map the host name assigned to a file server in Data Insight to the host name assigned by the Enterprise Vault server.

To map the host names of a file server

1. In the Data Insight Management Console, click **Settings** > **Data Management**.

   The **Archiving (Enterprise Vault Configuration)** tab opens by default.
2. Click **Add Filer Mappings**. The **Filer Mappings** page displays a detailed list of all the file servers that are configured in Data Insight. The list also displays some already mapped file servers.

   For each filer configured in Data Insight, the **Enterprise Vault Filer** field displays the corresponding mapped filer in Enterprise Vault. For a filer that has no mapping, Data Insight attempts to automatically map it with the corresponding filer in Enterprise Vault using host name matching. If Data Insight does not find an automatic match, it does not display an entry. In this case, you must manually map the filers.
3. For each file server displayed in the **Filer** column, verify if its Data Insight host name is correctly mapped to its Enterprise Vault host name.

4. If an existing mapping is incorrect, then enter the correct value for the host name in the **Enterprise Vault Filer** field. Data Insight populates the **Enterprise Vault server** field based on the Enterprise Vault Filer selected.
5. Click **Save**.

Note: Clustered Windows File Servers are added to Data Insight using the name of the cluster. Enterprise Vault requires that virtual file servers configured in the cluster must be added to the Enterprise Vault configuration. To enable Enterprise Vault to archive paths on the virtual file servers, Data Insight automatically maps the virtual file servers in the Windows cluster to the configured Enterprise Vault server.

Viewing and managing the status of an operation

You can track and manage the progress of the operations that are initiated from the Data Insight Management Console. You can perform the following tasks from the **Action Status** page.

- View the progress of an operation.
- Cancel an ongoing operation.
- Re-run a completed or canceled operation.
- Delete a completed or canceled operation.

You can view the progress of the following types of operations:

- Archive operations using Enterprise Vault.
- Permission remediation operations.
- Remediation operation using custom actions.

The **Action Status** page displays the following information:

- The unique identification number of the operation.
- The system-generated name for the triggered operation indicating its origin.
- The type of the operation. For Enterprise Vault operations, the type is specified as *EV*. For permission remediation operations, the type is *PR*. For custom action operations, the type is *CUSTOM*.
- The time when the user triggered the operation from the Management Console.
- The user who triggered the operation.

- The time when the destination server starts processing the triggered request. The destination server is an external server which is responsible for the actual execution of an operation that is triggered from the Data Insight Management Console. For example, in the case of an archive operation, the destination server is the Enterprise Vault server.
- The time when the destination server completes processing the request.
- The time it takes to complete the operation.
- The status of the operation.
- In case of archive operation using Enterprise Vault, the custom index property, including the classification tags applied to the files that are being archived.

Note: Only some columns are displayed in the default view. You can view the other columns by selecting them from the column header drop-down.

The **Details for Action** panel shows you the step-by-step break-down of the selected operation.

To view the status of an operation

1. In the Management Console navigate to **Settings** > **Action Status**. The **Action Status** page displays the details of recently triggered operations.
2. Use the check box filter to display the operations based on their **Type** or **Status**. Additionally, you can use the search facility to display the operations based on their attributes such as **Origin**, **Type**, or **Status**.
3. Click the **Origin** of the selected operation to view granular details of an operation. Alternatively, click the **Select Action** drop-down, and select **View**.
4. The details of the selected operation are displayed in the **Details for Action** panel.
5. Use the check box filter to display the operations based on the attributes such as: **Status** or **Filer**. Additionally, you can use the search facility to display the details based on attributes such as **Path** or **Status**.

You can cancel an operation that is in progress. Cancelling an operation, pauses all the activities of the operation. You can re-run a canceled operation later.

To cancel an ongoing operation

1. In the Management Console, navigate to **Settings** > **Action Status**. The **Action Status** page displays the details of recently triggered operations.
2. Use the check box filter to display the operations based on their **Type** or **Status**. Additionally, you can use the dynamic filter to display the operations based on their attributes such as **Origin** , **Type**, or **Status**.
3. Click **Select Action** for the operation you want to cancel.
4. Click **Cancel**.

You can re-run a canceled or a completed operation.

To re-run a canceled or a completed operation

1. In the Management Console navigate to **Setting** > **Action Status**. The **Action Status** page displays the details of recently triggered operations.
2. Use the check box filter to display the operations based on their **Type** or **Status**.
3. Click **Select Action** for the operation you want to re-run.
4. Select **Run Again**.
5. Select any of the following:
   - **All** - To run all the sub-steps for the operation.
   - **Unsuccessful** - To run all the failed sub-steps for the operation.

   Note: For a permission remediation or custom action-related operation that is canceled, the option to run the unsuccessful steps again is not available.

You can delete an operation that is canceled or completed.

To delete a canceled or a completed operation

1. In the Management Console navigate to **Setting** > **Action Status**. The **Action Status** page displays the details of recently triggered operations.
2. Use the check box filter to display the operations based on their **Type** or **Status**.
3. Click **Select Action** for the operation you want to delete.
4. Select **Delete**.

Section 7

Reference

- Appendix A. Backing up and restoring data
- Appendix B. Data Insight health checks
- Appendix C. Command File Reference
- Appendix D. Data Insight jobs
- Appendix E. Troubleshooting

Appendix A

Backing up and restoring data

This appendix includes the following topics:

- Selecting the backup and restore order
- Backing up and restoring the Data Insight Management Server
- Backing up and restoring the Indexer node

Selecting the backup and restore order

To maintain consistency in the configuration data during backup, the backup Data Insight components in the following order:

- Nodes with Indexer role
- Management Server

Restore the Data Insight components in the following order:

- Management Server
- Nodes with Indexer role

Backing up and restoring the Data Insight Management Server

It is mandatory to backup the Management Server.

To backup Management Server configuration files

1. Log in to Data Insight Management Server.
2. Backup the entire `$data` folder using backup tools such as Veritas NetBackup. The backup software should be capable of taking Volume Shadow Copy/Snapshot based backups. If your backup software does not have such a capability, you must stop all Data Insight services before backup, to avoid incomplete backup due to locked files.

To restore the Management Server files

1. Install the operating system. Use the same version, host name (recommended for ease of configuration), and architecture as was installed before the backup.
2. Install the same version of Data Insight that was installed before the backup.
3. Select the option to install the Management Server role while installing Data Insight .
4. Specify the original location of the `$data` directory as the previous install. By default, the `$data` directory is located at `C:\DataInsight\data`.
5. Complete the installation. Do not start the services at this time; clear the **Start services now** option when the installer prompts for it.
6. Delete the `$data` folder that is created as a part of the new installation and copy the backed up data to this location.
7. Start the Data Insight services, which include DataInsightComm, DataInsightWatchdog, DataInsightHttpd, DataInsightweb, and DataInsightConfig.
8. Check the status of the services and ensure that they come to running state.

   Successful start of all services indicates that the Management Server is successfully restored.

To restore the Management Server with a different host name or IP address

1. Repeat steps 1 through 6 as described in the section, *Restoring the Management Server files*.
2. Edit `$data/conf/<nodename>.conf` and enter the new server name.

   Open the file, `$data/conf/config.db.<N>` (N being the latest version of config.db) in an SQLITE editor.

   Update the node_name and node_ip columns in node table with the host name and IP address of the new server.

3 Run the following SQL updates:

```
update node set node_name='<new node name>'
where node_name='<prev node name>';

update node set node_ip='<new node IP>'
where node_ip='<prev node IP>';
```

4 Open a Windows command prompt and run the following command to increment the version of the config.db file that was changed in Step 3:

```
<INSTALL DIR>\DataInsight\bin\configdb -O -J dummy -j dummy
```

5 Start all Data Insight services.

6 On each worker node, except the Windows File Server agents, stop DataInsightComm and DataInsightConfig services.

7 Perform steps 2and 3 on the worker node's config.db.N

8 Start the DataInsightComm and DataInsightConfig services. Ensure that the worker nodes show online on the Data Insight Management Console.

Backing up and restoring the Indexer node

You must mandatorily backup the nodes that serve as the Indexer roles in a Data Insight deployment.

To back up the Data Insight server with Indexer role

1 Log in to the server with the Indexer role.

2 Backup the entire `$data` folder using backup tools such as Veritas NetBackup. The backup software should be capable of taking Volume Shadow Copy/Snapshot based backups. If your backup software does not have such a capability, you must stop all Data Insight services before backup, to avoid incomplete backup due to locked files.

To restore the Indexer node

1 Install the operating system. Use the same version, host name (recommended for ease of configuration), and architecture as was installed before the backup.

2 Install the same version of Data Insight that was installed before the backup.

3 Select the option to install the Indexer and Collector role while installing Data Insight. If installing Data Insight on a Linux server, select the option to install Indexer.

4. Specify the original location of the `$data` directory as the previous install. By default, the `$data` directory is located at `C:\DataInsight\data`.
5. Clear the **Launch worker node registration wizard after exit** checkbox. You do not need to register the worker node at this time as the registration information is already present in the data that you have backed up.
6. Complete the installation. Do not start the services at this time; clear the **Start services now** option when the installer prompts for it.
7. Delete the `$data` folder that is created as a part of the new installation and copy the backed up data to this location.
8. Start the Data Insight services, which include DataInsightComm, DataInsightWatchdog, and DataInsightConfig.
9. Check the status of the services and ensure that they come to running state.

   Successful start of all services indicates that the Indexer node is successfully restored.

To restore the Indexer node with a different host name or IP address

1. Repeat steps 1 through 6 as described in the section, *Restoring the Indexer node.*.
2. Edit `$data/conf/<nodename>.conf` and enter the new server name.
3. Open the file, `$data/conf/config.db.<N>` (N being the latest version of config.db) in an SQLITE editor.

   Update the node_name and node_ip columns in node table with the host name and IP address of the new server.
4. Run the following SQL updates:

```
update node set node_name='<new node name>'
 where node_name='<prev node name>';

update node set node_ip='<new node name>';
where node_name='<prev node name >';
```

5. Log in to the Management Server and stop the DataInsightComm, DataInsightWeb, and DataInsightConfig services.
6. Perform 3 on the Management Server.
7. Open a Windows command prompt and run the following command to increment the version of the config.db file that was changed in 2

```
<INSTALL DIR>\DataInsight\bin\configdb -O -J dummy -j dummy
```

8. Start all Data Insight services on the Management Server.

9 On each worker node, except the Windows File Server agents, stop DataInsightComm and DataInsightConfig services.

10 If this node is a Collector for one or more Windows File Server agents, log in to each Windows File Server, stop the DataInsightComm and DataInsightConfig services.

   Perform step 3 on the worker node's config.db.N

11 Start the DataInsightComm and DataInsightConfig services on the Indexer and all other worker nodes where configdb.N was changed. Ensure that the worker nodes show online on the Data Insight Management Console.

Appendix B

Data Insight health checks

This appendix includes the following topics:

- About Data Insight health checks
- Understanding Data Insight best practices

About Data Insight health checks

You may want to run Veritas Data Insight health checks in your environment in the following scenarios:

- Routine check of the Data Insight environment.
- Troubleshooting Data Insight performance and stability issues
- Upgrading Data Insight
- Resolving scalability issues and bottlenecks in your environment
- Checking for additional resources, for example, while installing the Self-Service portal.
- Migrating resources to new hardware.

You can use the checks for the following objectives:

- To determine if your configuration, sizing and indexing schedules meet Data Insight requirements.
- To fine-tune Data Insight efficiency.
- To free up the unwanted storage that is used by Data Insight.
- To troubleshoot common issues and collect logs and other evidence for the same.

Before you run the Data Insight health checks, you must be familiar with the following concepts:

- How Data Insight works
- Data Insight architecture
- Data Insight data flow

For more information, see the *Veritas Data Insight Installation Guide*.

To understand best practices and sizing guidelines for deploying, see the following topic:

See "Understanding Data Insight best practices" on page 347.

Services checks

Verify if the services listed in the following figure are running:

Figure B-1 Services running on a Data Insight node

DataInsightCelerra	Veritas Data Insight EMC Celerra Service	Running	Automatic	tulip0\Administrator
DataInsightComm	Veritas Data Insight Communication Server	Running	Automatic	Local System
DataInsightConfig	Veritas Data Insight Configuration Service	Running	Automatic	Local System
DataInsightFpolicy	Veritas Data Insight Fpolicy Service	Running	Automatic	tulip0\Administrator
DataInsightWatchdog	Veritas Data Insight Watchdog Service	Running	Automatic	Local System
DataInsightWeb	Veritas Data Insight Web Server	Running	Automatic	Local System
DataInsightWorkflow	Veritas Data Insight Workflow Service	Running	Automatic	Local System

Deployment details checks

Table B-1 lists the various deployment details that you must gather. The table also lists the related steps.

Table B-1 Deployment details

Deployment details	Steps
■ Number of product servers ■ Server configuration details (operating system,, physical or virtual, 32-bit or 64-bit, CPU, RAM, disk capacity, disk usage)	■ Check device manager. ■ Execute the `systeminfo` command on each server.
■ Number of NAS devices and configured shares ■ Vendor name and model numbers	■ List the details.

Table B-1 Deployment details *(continued)*

Deployment details	Steps
Review the `FILER SUMMARY:` section of the idxcheck.0.log files (in each indexer) and note the following details: ■ Overall statistics as well as per indexer and per filer statistics for events, capacity, file, and folder counts. ■ The sizes of the indexer database and the segment files.	■ To retrieve these details, on the Management Console, use the following log utility: https://*hostname*/logs OR Navigate to the following location to view required file.: `Installation directory>\log`
■ Network topology (proximity)	■ Determine the number of data centers (mention primary). ■ Determine the bandwidth between data centers. ■ Ensure that the Management Server is located close to the indexers. ■ Look for proximity of filers and collectors; gather subnet information for the collectors and filers. ■ Determine the number of configured directory services (Active Directory, LDAP, NIS) and the number of users that are configured by navigating to the following location on the Data Insight Management Console: **Settings**>**Inventory**> **Directory Services**
■ Number of the events that are generated and processed per Collector (Only in Data Insight 4.0 or later.)	■ Determine the number of events received weekly by navigating to https://*localhost*/datadir In the list of folders, select **stats** and then select **statistics.db**. In the query drop down, select "weekly_stats", enter **statid=252842000** in the **where** text box, and then click **Go**. ■ Determine number of events processed weekly. Navigate as described above but enter **statid=251793411** in the **where** text box and then click **Go**.

Generic checks

Generic checks for all servers include checks for the following information:

■ Configuration version

- Component version
- Resource usage statistics
- `Windbg` check
- Error files count
- Error events check
- Attic check
- Crash dump check

Configuration version

Table B-2 lists three configuration scenarios that you must review, and their related steps:

Table B-2 Configuration version checks

Configuration versioning	Steps
Are the versions on the Management Server and the other nodes consistent for configuration, users, and database? Note that the policy versions may vary.	Run the following command on Management Server/node: # `configdb -V`
Are the versions on the Management Server lower than other Data Insight nodes? This behavior occurs in restore cases.	Contact Veritas Support for more information.
Are the versions considerably higher on Management Server than the other Data Insight nodes?	■ Ensure that connectivity exists between Management Server and worker nodes. ■ Ensure that connectivity exists between Collector nodes and Windows File Server agents. (if applicable).

Component versioning

Perform one of the following steps to gather component versions for all nodes:

- Navigate to **Settings** > **Servers**.
- Or, to get the node number-wise breakdown of component versions, run the following command:

  # `configdb -p -T objattr_node| findstrnode_info_product.version`

Note: Data Insight Management Server, Indexer nodes, and Collector nodes must run the same version of Data Insight. Windows File Server agent version may be different from the Data Insight version on the head servers. Also note that Data Insight 5.0 supports Windows File Server agents later than version 4.0 or later.

Resource usage statistics

To view the resource consumption statistics per server, on the Management Console, navigate to **Settings** > **Servers** > **Statistics**.

Verify that the average resource consumption is consistently high.

Figure B-2 Performance chart

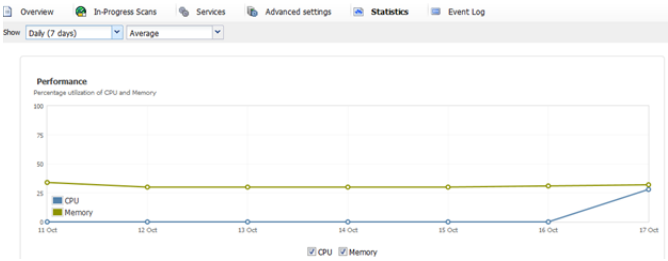

Table B-3 describes the various parameters you must verify in the Data Insight performance chart and the remedial steps required.

Table B-3 Performance checks

Performance parameter	Steps
CPU usage is consistently high.	Look for processes clogging the server.
CPU usage is consistently low.	Look for cores and threads settings in server-specific checks. Verify if a server is unable to process more load.
Memory consumption is consistently high.	Look for processes clogging the server.
Spike patterns exist in hourly or weekly charts, or both.	If the spikes occur in a weekly chart or hourly chart form a pattern, verify if they coincide with the Data Insight schedule. Also investigate other schedules on the servers.

Figure B-3 illustrates a stable disk usage pattern.

Figure B-3 Disk usage chart

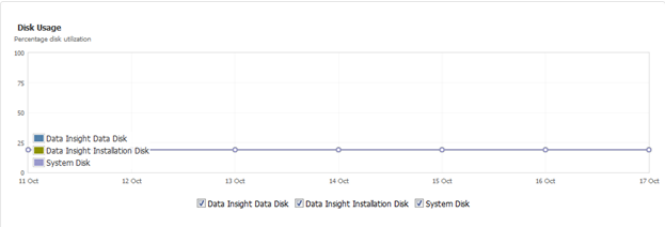

The following table describes the remedial steps that you must perform, if disk usage increases to high levels.

Table B-4 Disk usage checks

Disk usage	Steps
Disk usage is consistently high; that is, consistently reaching 100 per cent.	■ Use stale data statistics and purge data OR ■ Add new storage.

Windbg check

If the `windbg` utility is installed, check for any processes that are captured and are running in windbg. For example, if you run a report that fails, say `report.exe`, Windbg captures the process. Investigate if any process appears to be running, but is not functional.

Error files count

To view the error file count, navigate to **Settings** > **Servers** > **Statistics**

Figure B-4 Error files chart

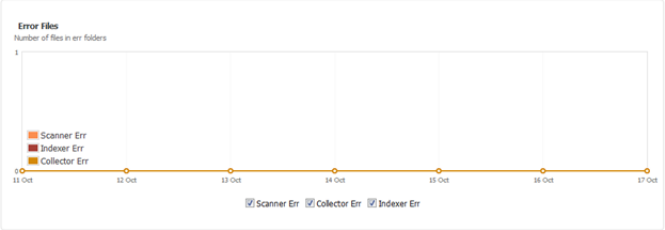

Ideally, there should be no error files. View errors at the following location: *datadir*`\indexer\err folder`.

Where *datadir* is the path of the `data` directory

For example: `C:\DataInsight\data\indexer\err` folder.

Error events check

To verify if any error events are logged, navigate to:

Settings > **Servers** > **Event Log**.

Attic check

Ensure that the attic feature is turned off on all Data Insight nodes.

That is, ensure that there are no intermediate files in the following directory: `DATADIR/attic`.

Crash dump check

Ensure there are no dump files in the following directory: `INSTALLDIR/dumps`.

Data Insight Management Server checks

Table B-5 lists Data Insight Management Server checks and related steps.

Table B-5 Management Server checks

Checks	Steps
`DATADIR/events` folder size	■ Over a period of six months, in a moderately stable environment, the folder size reaches about 200MB. Set the data retention settings accordingly. ■ Purge very old Data Insight events with the approval of the event owners.
Accumulation of backlog in `DATADIR/inbox`	■ Verify the date of the oldest file (backlog). If the file is very old, it means that the audit or the scan files are unable to keep up with the collection. ■ Use the `configcli list_jobs` command to collect the schedule and last run of the ADScanJob, ProcessEventsJob, and IndexWriterJob processes. ■ Collect Event Viewer logs. ■ Collect the output of the `indexcli -j` command.

Table B-5 Management Server checks *(continued)*

Checks	Steps
Presence of old or stale reports in the following folder: `DATADIR/console/reports/reportruns`	■ If the reports are no longer required, archive or delete them. ■ For each report, set a reasonable value in the `Maximum reports to preserve` field.
SMTP and events notification check	Ensure that these settings are configured.

Data Insight Indexer checks

Table B-6 lists various Data Insight Indexer checks that you must perform and the related remedial steps:

Table B-6 Indexer checks

Checks	Steps
Number of filers and the shares that are serviced by each Indexer	■ Navigate to **Settings** > **Filers**, and filter on each Indexer. Or ■ Check the `idxcheck` log for each Indexer.
Index integrity and space usage per Indexer	■ Check `Idxcheck0.log` under *installdir*/log folder checks for database integrity and index. ■ Go to the end of the log file to review the overall integrity check. ■ If the integrity check fails, review the index check for each share in the same file. ■ Review two key metrics: *IndexDB Size* and *Segments Size*. Others metrics indicate stale hash sizes. IndexDB Size is a sum of all index databases. It stores index information such as scan information (metadata and ACLs). Segment size is a sum of all segment sizes. It stores the audit events. Hashes are indexed against the two databases - indexDB and segments. *ActIndex Size* was introduced in Data Insight 4.0 to indicate activity in a social network map.

Table B-6 Indexer checks *(continued)*

Checks	Steps
Inbox and Outbox chart Figure B-5 illustrates the chart	■ Navigate to **Settings** > **Servers**> **Statistics**. ■ Ensure that the Inbox plot follows a smooth saw-tooth pattern. ■ If the number of files in inbox consistently rises, and so does CPU and memory usage, investigate if there is an Indexer issue. ■ Check schedule settings and number of threads to fine-tune the processing of incoming files.
Basic Settings	In Data Insight UI, navigate to **Data Insight** > ***Indexer*** > **Advanced Settings**, and fine-tune the following parameters as necessary: ■ Total Indexer threads (*): This number should not exceed the total CPU core count on the server. For a 4-CPU Indexer, ideally set this value to 'total cpu core count minus 1'. Allowing one core's worth of processor ensures that index writer processes cannot overcome a server and keeps resources available for query response and product communications. If the Indexer has more than 4 cores, leaving more cores available may be desirable if many reports are run. Disk subsystems start to lose the ability to keep up if there are more than 10 threads. If the server has more than 12 cores, do not set value to more than 10 on index writer. Specify a higher value only if the storage is very high performance. ■ Limit the maximum events that are processed in memory: Keeping this setting unlimited may lead to processing issues. Most of these issues should be resolved, but if the filer in question has excessively large number of events. (Application-generated workloads sometimes cause this type of load.) A limit of 1,000,000 is recommended. ■ Reconfirm deleted paths when reconciling full scan information: This Indexer-side setting enables Data Insight to reconfirm if paths now missing from full scan output are deleted. ■ Indexer schedule ■ Indexer integrity checking schedule

Table B-6 Indexer checks *(continued)*

Checks	Steps
Audit Events Pre-Processor Settings	In the Data Insight UI, navigate to Data Insight *Indexer*>**Advanced Settings** ■ Audit events preprocessor schedule ■ Batch size
Accumulation of backlog in `DATADIR/inbox`	■ Look for the date of oldest files (backlog): If files are very old, it means that audit or scan file processing is not able to keep up with collection. ■ Execute the `configcli list_jobs` command to gather *ADScanJob*, *ProcessEventsJob*, and *IndexWriterJob*. ■ Collect Event Viewer logs. Collect the output of the `indexcli -j` command.

Audit/scan files in `DATADIR/indexer/err`	■ Review `ldxcheck.log` for index corruption.
	■ Move files to `Inbox` and verify if they are consumed in the next run. If not, contact Veritas Support with the following information:
	  ■ Event viewer logs
	  ■ Index directory for the share or the site collection
	  ■ *installdir*`\log\commd.0.0.log`
	  ■ *installdir*`\log\indexer\index-`*msu_id*`.log`
	  ■ ldxcheck.log (summary and integrity information for all shares)
	  ■ *datadir*`\indexer\err folder`
	  ■ Output of `indexcli.exe -j` command (To view the files queued for indexing.)

The following figure illustrates a smooth saw-tooth pattern in an **Inbox and Outbox** chart:

Figure B-5 Inbox and Outbox chart

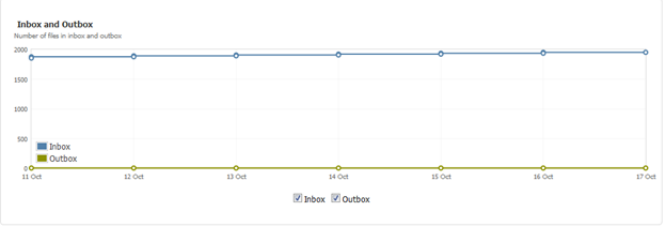

Data Insight Collector checks

The following table lists various Data Insight Collector checks that you must perform and the related steps:

Table B-7

Checks	Steps
Number of filers (and shares) serviced by each collector	If applicable, also specify the number of Windows File Servers with the Windows File Server agents that are configured.
Ping time from Collector to the filers	This value should not be more than 1 microsecond. If the value is greater than 1 microsecond, you must investigate network latency issues.
Server statistics ■ Inbox and Outbox chart Figure B-6 illustrates Inbox and Outbox chart ■ Event pre-processor throughput chart Pre-processor throughput provides an idea about the processing of collected audit events. Pre-processor collector threads govern the optimal throughput and CPU utilization. Figure B-7 illustrates the Event pre-processor throughput chart.	To collect server statistics, navigate as follows (for each server): **Server** > **Settings** > **Statistics**. Review the following: ■ In the **Inbox and Outbox** chart, ensure that the Outbox plot follows a smooth saw-tooth pattern. If the number of files in the `Outbox` folder consistently rises, investigate if there is a communication issue or network issue. Check the schedule settings and number of threads to fine-tune the processing of incoming files. ■ If the Event pre-processor throughput drops, investigate if events received are few or if pre-processing is not able to catch up.

Table B-7 *(continued)*

Checks	Steps
Basic settings for full and incremental scans under Filesystem scanner settings	In the Data Insight UI, navigate to **Data Insight** > *Collector* > **Advanced Settings**, and fine-tune the following parameters as necessary. ■ Maximum scans to run in parallel on this server (*): You may increase these, as these threads are not CPU intensive. This value can be higher than the total CPU core count on the server. Network latency and bandwidth put limits on scanners. Multiple threads use available bandwidth better, during the latency for other thread queries. However, increasing the thread count does slow down individual scans, due to saturation of network performance. Begin with a default value of 4 and monitor the network performance. If there is plenty of bandwidth available, increase thread count until the network utilization increases, but doesn't saturate. Ensure that you allocate bandwidth and CPU resources for audit events and inter-product communication. Start with two threads per core, beyond the default configuration. If network resources or CPU resource usage averages over 70 per cent, reduce the thread count. ■ Maximum shares per filer to scan in parallel: Ensure that this value is not more than 4 per filer. ■ Scan settings: Specify if you want to override a scheduled scan or pause a scheduled scan.
Audit Events Pre-Processor Settings	In the Data Insight UI, navigate to **Data Insight Servers** > *Collector* > **Advanced Settings**, and change the following settings: ■ Total collector threads: To improve CPU utilization, increase this value to half the CPU core count. In case of 4 cores, you can keep this value as 3. Default value is 1. ■ Schedule: Schedule pre-processing twice per every indexing job (as per the indexing schedule). For example, if the indexing job runs every 4 hours, pre-processing should run every 2 hours.

Table B-7 *(continued)*

Checks	Steps

Accumulation of backlog and errors in DATADIR/collector (error and staging)	■ Look for the date of oldest files (backlog): If files are very old, it means pre-processing is not able to keep up with audit events collection.
	■ Verify pre-processing settings. You can set this value to more than 2 GB. You can increase the number of threads.
	■ Collect the following information:
	  ■ To see schedule and last run status of the *CollectorJob* process, execute the command: `configcli list_jobs`
	  ■ Collect the following log: *installdir*`\log\collector_`*n*`.log`. Where installdir is the installation directory: For example: `C:\ProgramFiles\DataInsight`
	  ■ Collect the following log: `installdir\log\commd.0.0.log`
	  ■ Collect the following contents: *datadir*`\data\config.db.<n>`.
	  ■ Collect the following content: *datadir*`\data\collector\err\*.sqlite`.
	  ■ Collect Event Viewer logs.

Accumulation of backlog in DATADIR/outbox	■ Look for the date of oldest files (backlog): If the files are very old, investigate if the network is slow or communication is broken
	■ Collect the following information
	  ■ To see the schedule and last run status of CollectorJob and FileTransfer job processes on the collector, execute the following command: `configcli list_jobs`
	  ■ Collect Event Viewer logs.
	  ■ If communication issues occur, add the entries of the lindexer nodes in `etc/hosts` directory of collectors; and the Collector nodes' entries in `etc/hosts` folder of Windows File Servers.

Accumulation of backlog in DATADIR/changelog	■ Look for the date of the oldest file (backlog): If the files are very old, investigate if incremental scans are fail.
	■ Collect the following information:
	  ■ Last scan job
	  ■ Scan schedule

Table B-7 *(continued)*

Checks	Steps
Check for any errors in DATADIR/scanner (err)	■ Collect the following information: ■ Collect the following Scanner logs: `ext<x>_msu<y>.*log logs` Where x and y are the extract ID and the MSU ID, respectively. In case of SharePoint, sharepoint_scanner_<n>.log, where 'n' is the site collection ID. ■ Collect the following log: *installdir*\log\commd.0.0.log. ■ Collect Event Viewer logs

The following figure illustrates the smooth saw-tooth pattern in the Inbox and Outbox chart:

Figure B-6 Inbox and Outbox chart

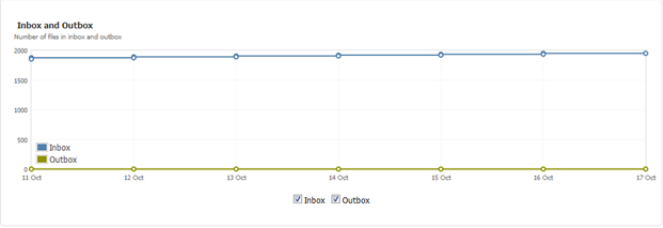

The following figure is a screen shot of the Event Pre-processor throughput:

Figure B-7 Event Pre-processor throughput

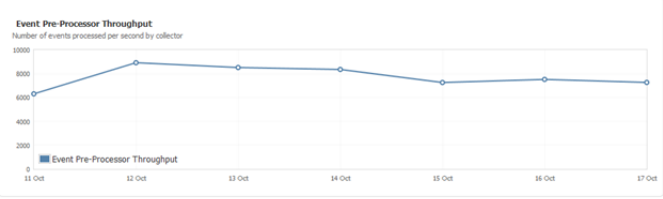

Data Insight Windows File Server checks

Table B-8 lists various Data Insight Windows File Server checks that you must perform and the related steps:

Table B-8 Windows File Server checks

Checks	Steps
File system scanner settings	See "Data Insight Collector checks" on page 342.

Accumulation of backlog and errors in DATADIR/collector (error and staging)	■ Look for the date of oldest files (backlog). If files are very old, it means pre-processing is not able to keep up with audit events collection.
	■ Verify pre-processing settings. You can set this value to more than 2 GB. You can increase the number of threads.
	■ Collect the following logs:
	  ■ `installdir\log\collector_n.log`
	  Where *installdir* is (?) and *n* is (?) (Query)
	  ■ `installdir\log\commd.0.0.log`
	  ■ `<datadir>\data\config.db.<n>`
	  ■ `<datadir>\data\collector\err\*.sqlite`
	  ■ Collect Event Viewer logs.
	■ To see schedule and last run status of the *CollectorJob* process, execute the command: `configcli list_jobs`.
	■ Look for kernel ring buffer error in `winnasd.log`.
	If the error is present, increase the **Maximum kernel ring buffer size** setting. (Windows File Server Agent setting)

Accumulation of backlog in DATADIR/outbox	■ Look for the date of oldest files (backlog): If the files are very old, investigate if the network is slow or communication is broken.
	■ Collect the following information:
	  ■ To see the schedule and last run status of *CollectorJob* and file *FileTransferJob* processes on the Collector, execute the following command: `configcli list_jobs`
	  ■ Collect Event Viewer logs.
	  ■ If communication issues occur, add entries of indexer nodes in the `etc/hosts` directory of the Collector nodes; and Collector nodes' entries in `etc/hosts` folders of Windows File Server.

Table B-8 Windows File Server checks *(continued)*

Checks	Steps

Accumulation of backlog in DATADIR/changelog	■ Look for the date of the oldest file (backlog): If the files are very old, investigate if incremental scans fail.
	■ Collect the following information:
	  ■ Last scan job
	  ■ Scan schedule

Data Insight SharePoint checks

Ensure that the SharePoint collector and SharePoint Agents are running the same version.

Understanding Data Insight best practices

The following topics describe the common practices and sizing guidelines:

- Common practices
- General sizing guidelines

Common practices

- Use 64-bit server architecture for better Data Insight performance.
- Use high performance disks for indexers, SAN disk for ease of expansion and backup
  - Disk concatenation or RAID 4 may generate hotspots in the array subsystem. Ensure proper distribution to avoid hotspots.
- Use separate disks for Install and Data directories.
- Assign faster and bigger Temp directory for indexers.
  - Temp directory (for example, in Windows, the directories that are indicated by TMP and TEMP system environment variables) must be on disk other than system disk (operating system disk). This best practice improves index vacuum (like defrag for segments) performance, and overall performance.
- Assign indexer memory of 1 GB per thread.
- Assign indexer memory sizing of 1 GB per 10 million audit events.
- Assign low latency high bandwidth (or dedicated) network links for Collector-Filer communication.

- Create exclude rules to ensure that the information that is reported is relevant and useful.
  - If a third-party application that generates a lot of events, resides on a volume, exclude that volume from auditing (like NetApp FPolicy) to restrict events.
- Schedule scans at off peak hours to minimize effect on users accessing the shares.
- Disable scans from running during peak usage.
- Manage job scheduling as per load.
- Assign appropriate roles and access for Data Insight users.
- Deploy DFS mappings to ensure that all names are recognizable and meaningful.
- Create containers to logically group related objects together for administration and reporting purposes.
- Set up event notifications to ensure that errors and warnings are reported.
- Define retention policies to ensure the database and log files are maintained over time.
- Define report retention in when configuring reports (“Reports to preserve” setting).
- Maintain adequate proximity – Management Server close to indexers, Collector close to filers.
- Use Chrome for faster UI.
- Refer to the *Veritas Data Insight Installation Guide* for the latest platform and version support information.

General sizing guidelines

The following sizing guidelines are typically recommended for Data Insight setups. The recommendations may change based on environment.

- Management Server
  - Up to 16 indexers
- Indexers
  - Up to 20,000 shares
  - Allocate 100 MB of disk space per million files
  - Allocate 20 MB of disk space per million events
  - Typical storage size range is 40 GB to 400 GB depending upon environment. Large deployments may even need more space.

- Collectors
  - Up to 10,000 shares or ten filers, whichever comes first
  - Up to 150 Windows NAS agents
  - Up to 20 SharePoint web front-end servers
  - Typical storage size range is 60-80 GB per collector. Additional space may also help in case of network outages where data is staged before it is transferred to Indexer.
- Indexer storage
  - 100 MB/1 million files
  - 20 MB/1 million events
- Scan times
  - 200 files per second per thread

Risk Dossier sizing guidelines

If you enable the computation of the Risk Dossier, there will be an overall increase in the space required for the index folder.

For the Management Server, the dossier database size can be calculated as:

Number of users * 2 MB (to store dossier date for 6 months)

Thus, for 50000 users, provision up to 100 GB space to store dossier data of 6 months.

Appendix C

Command File Reference

This appendix includes the following topics:

- fg.exe
- indexcli.exe
- reportcli.exe
- scancli.exe
- installcli.exe

fg.exe

`fg.exe` – A script that modifies the file group configuration for Data Insight.

SYNOPSIS

```
fg -C -N <name of file group>

fg -D -N <name of file group>

fg -L -d

fg -L -N <name of file group> -d

fg -R -N <name of file group> -t <name of extension>
```

Description

fg is a script used to modify the configuration for sorting files into file groups. By default, Data Insight sorts files into 18 file groups based on the file extensions.

Options

`-i <username>`
(Required) The fully-qualified user name of the user running the command, for example, user@domian. This user should have Server Administrator privileges in Data Insight.

`-A` Adds an extension to an existing file group.

`-C` Creates a new file group.

`-D` Deletes an existing file group.

`-L` Lists existing file groups.

`-R` Removes an extension from an existing file group.

`-N` Name of the file group to be created or deleted.

`-d` Shows file group details when listing existing file groups.

`-t <name of extension>`
The file extension to add or delete from the file group (For example, doc).

`-h` Prints the usage message.

EXAMPLES

EXAMPLE 1: The following command creates a new file group.

```
fg -i <username> -C -N <name of file group>
```

EXAMPLE 2: The following example adds a new extension to an existing file group.

```
fg -i <username> -A -N <name of file group>  -t <name of extension>
```

EXAMPLE 3: The following example deletes an extension from an existing file group.

```
fg -i <username> -R -N <name of file group>  -t <name of extension>
```

EXAMPLE 4: The following command deletes a file group.

```
fg -i <username> -D -N <name of file group>
```

EXAMPLE 5: The following command displays a detailed listing of all configured file groups.

```
fg -i <username> -L -d
```

EXAMPLE 6: The following command displays a detailed listing of a particular file group.

```
fg -i <username> -L -N <name of file group> -d
```

indexcli.exe

`indexcli.exe` – a utility that manages the index segments available on an Indexer worker node.

SYNOPSIS

```
indexcli.exe
             --display|--archive|--purge|--restore|--rearchive|--list-jobs
             |--stop-jobs [OPTIONS]

indexcli.exe -A <name of the index segments to be archived>

indexcli.exe -c

indexcli.exe -D <name of the index segments to be purged>

indexcli.exe -d

indexcli.exe -h

indexcli.exe -j

indexcli.exe -r

indexcli.exe -t

indexcli.exe -u
```

Archive options

```
indexcli.exe -A -a ¦ -f <FILERS> ¦ -m
<SHARES> ¦ -S <SITECOLLS> ¦ -w <WEBAPPS> ¦ -I
 <MONTHS>
```

`-a` Archives all index segments older than the specified interval.

`-f <name of filer(s)>`
Archives all index segments for the specified list of filers.

`-I <interval in months>`
Archives segments older than the specified interval. The segments which have been restored earlier are not archived.

`-m <name of share(s)>`
Archives all index segments for the specified list of shares.

`-S,--sitecoll <SITECOLLS>`

Archives segments for specified list of Microsoft SharePoint site collections.

`-w, --webapp <WEBAPPS><`

Archives segments for specified list of Microsoft SharePoint web applications.

Purge options

```
indexcli.exe -D -a ¦ -f <FILERS> ¦
-m <SHARES> ¦ -S <SITECOLLS> ¦ -w <WEBAPPS> ¦
-I <MONTHS>
```

`-a` Purges all index segments older than the specified interval.

`-f <name of filer(s).`

Purges all index segments for the specified list of filers.

`-I ,interval in months.`

Purges segments older than the specified interval. The segments which have been explicitly restored earlier, for which the lease is still valid, are not purged.

`-m <name of share(s)>`

Purges all index segments for the specified list of shares.

`-S,--sitecoll <SITECOLLS>`

Purges segments for specified list of Microsoft SharePoint site collections.

`-w, --webapp <WEBAPPS><`

Purges segments for specified list of Microsoft SharePoint web applications.

Display options

```
indexcli.exe -d -a ¦ -f <FILERS> ¦
-m <SHARES> ¦ -S <SITECOLLS> ¦ -w <WEBAPPS> ¦
-s <STATES>
```

`-a` Displays information for all shares.

`-f <name of filer(s)>`

Displays information for the specified list of filers.

`-m <name of share(s)>`

Displays information for the specified list of shares.

`-s <name of state>`

Displays index segments for the given state only. Multiple stars can be separated by comma. Possible states are, ARCHIVING, RE-ARCHIVING, ARCHIVED, RESTORING, RESTORED, RESTORE, FAILED, or DELETED.

`-S,--sitecoll <SITECOLLS>`

Displays information for a specified list of Microsoft SharePoint site collections.

`-w, --webapp <WEBAPPS><`

Displays information for a specified list of Microsoft SharePoint web applications.

Restore options

```
indexcli.exe -r -a ¦ -f <FILERS> ¦
-m <SHARES> ¦ -S <SITECOLLS> ¦ -w <WEBAPPS> -C -F <FROM> ¦ >
¦ -R <RANGE>
[-L <MONTHS> ¦ -l <MONTHS> ¦ -y]
```

`-a` Restores the index segments for all shares.

`-C` If the continue-on-error option is not specified, the restore command fails if the segment files required to restore data for the specified parameters are not available.

`-f <name of filer(s)>`

Restores all index segments for the specified list of filers.

`-F <month from which the segments need to be restored>`

Specify the month in the format, YYYY/MM. For example, `indexcli.exe -r -F 2010/01` restores segments from January 2010 till date.

`-L <interval in number of months>`

Resets lease on segments that are restored earlier using `-l` option. Specify the new lease interval in months. This command replaces the previous lease interval. Setting the value to 0 will make the lease permanent.

`-l <interval in number of months>`

Restores segments for a temporary lease in months. After the lease expires, restored segments are automatically re-archived. If this option is not specified, segments remain restored till you re-archive them with the `-u`, --rearchive, option.

`-m <name of share(s)>`

Restores all index segments for the specified list of shares.

`-S,--sitecoll <SITECOLLS>`

Restores segments for a a specified list of Microsoft SharePoint site collections.

`-w, --webapp <WEBAPPS><`

Restores segments for a specified list of Microsoft SharePoint web applications.

`-R <range in months>`

Restore all index segments for the specified month range. Specify the month in the format, YYYY/MM-YYYY/MM. For example, `indexcli.exe -r -R 2010/01-2010-03` restores segments from January 2010 to March 2010.

`-y` Instead of restoring segments, this option displays the list of files that must be available before restoring the specified segments.

Re-archive options

```
indexcli.exe -u -a ¦ -f <FILERS> ¦
-m <SHARES>  ¦ -S <SITECOLLS> ¦ -w <WEBAPPS>
-F <FROM> ¦ -R <RANGE>
```

`-a` Re-archives all previously restored index segments.

`-f <name of filer(s)>`

Re-archives previously restored index segments for the specified list of filers.

`-F <month FROM which the segments need to be restored>`

Specify the month in the format, YYYY/MM. For example, `indexcli.exe -u -F 2010/01` restores segments from January 2010 till date.

`-m <name of share(s)>`

Re-archives previously restored index segments for specified list of shares.

`-S,--sitecoll <SITECOLLS>`

Re-archives previously restored segments for a specified list of Microsoft SharePoint site collections.

`-w, --webapp <WEBAPPS><`

Re-archives previously restored segments for a specified list of Microsoft SharePoint web applications.

`-R <range in months>`

Restore all index segments for the specified month range. Specify the month in the format, YYYY/MM-YYYY/MM. For example, `indexcli.exe -u -R 2010/01-2010-03` restores segments from January 2010 to March 2010.

EXAMPLES

EXAMPLE 1: The following command archives index segments for specified list of filers.

```
indexcli.exe -A -f \\filer1,\\filer2,ID1,ID2
```

EXAMPLE 2: The following command archives index segments for specified list of shares.

```
indexcli.exe -A -m \\filer1\share1,\\filer2\shares2,ID3,ID4
```

EXAMPLE 3: The following command purges index segments for specified list of filers.

```
indexcli.exe -D -f \\filer1,\\filer2,ID1,ID2
```

EXAMPLE 4: The following command purges segments for specified list of shares.

```
indexcli.exe -D -m \\filer1\share1,\\filer2\shares2,ID3,ID4
```

EXAMPLE 5: The following command restores index segments for specified list of filers.

```
indexcli.exe -r -f <\\filer1,\\filer2,ID1,ID2>
```

EXAMPLE 6: The following command restores index segments for specified list of shares.

```
indexcli.exe -r -m \\filer1\share1,\\filer2\shares2,ID3,ID4
```

EXAMPLE 7: The following command re-archives previously restored index segments for specified list of filers.

```
indexcli.exe -u -f \\filer1,\\filer2,ID1,ID2
```

EXAMPLE 8: The following command re-archives previously restored index segments for specified list of shares.

```
indexcli.exe -u -m \\filer1\share1,\\filer2\shares2,ID3,ID4
```

EXAMPLE 9: The following command archives segments for specified list of Microsoft SharePoint site collections.

```
indexcli.exe -S,--sitecoll<http://sp_webapp:8000/sc1,ID2,ID3...>
```

EXAMPLE 10: The following command archives segments for specified list of Microsoft SharePoint web applications.

```
indexcli.exe - w,--webapp<http://sp_webapp:8000,ID2,ID3,...>
```

EXAMPLE 11: The following command purges segments for specified list of Microsoft SharePoint site collections.

```
indexcli.exe -S,--sitecoll<http://sp_webapp:8000/sc1,ID2,ID3...>
```

EXAMPLE 12: The following command purges segments for specified list of Microsoft SharePoint web applications.

```
indexcli.exe - w,--webapp<http://sp_webapp:8000,ID2,ID3,...>
```

EXAMPLE 13: The following command displays information for specified list of Microsoft SharePoint site collections.

```
indexcli.exe -S,--sitecoll<http://sp_webapp:8000/sc1,ID2,ID3...>
```

EXAMPLE 14: The following command displays information for specified list of Microsoft SharePoint web applications.

```
indexcli.exe - w,--webapp<http://sp_webapp:8000,ID2,ID3,...>
```

EXAMPLE 15: The following command restores segments for specified list of Microsoft SharePoint site collections.

```
indexcli.exe -S,--sitecoll<http://sp_webapp:8000/sc1,ID2,ID3...>
```

EXAMPLE 16: The following command restores segments for specified list of Microsoft SharePoint web applications.

```
indexcli.exe - w,--webapp<http://sp_webapp:8000,ID2,ID3,...>
```

EXAMPLE 17: The following command re-archives previously RESTORED segments for specified list of Microsoft SharePoint site collections.

```
indexcli.exe -S,--sitecoll<http://sp_webapp:8000/sc1,ID2,ID3...>
```

EXAMPLE 18: The following command re-archives previously RESTORED segments for specified list of Microsoft SharePoint web applications.

```
indexcli.exe - w,--webapp<http://sp_webapp:8000,ID2,ID3,...>
```

reportcli.exe

`reportcli.exe` – a utility to create reports using a properties file that contains the input parameters, execute and list configured reports, check the status of the reports, and cancel report runs.

SYNOPSIS

```
reportcli.exe --list-jobs¦──list-reports¦--list-outputs¦--create
--execute¦--cancel¦--help [OPTIONS]

reportcli.exe -c

reportcli.exe -e

reportcli.exe -h

reportcli.exe -j

reportcli.exe -l

reportcli.exe -o
```

Options

`reportcli.exe -n -r <name of report> -p <property file path> -u <user name of creator> [-rt <report type>] [--users <path of users' .csv file>] [-t <path of .csv file of paths>] [--custodian <path of custodian' .csv file>]`

Creates a report using the properties file in which the input parameters are specified. The following attributes apply:

`-r --report <name of report>`	Creates a report with the specified name.
`-p --properties <property file path>`	Property file containing the input parameters for the report. **Note:** By default, a sample properties file is installed in the INSTALL_DIR/DataInsight/reports folder.
`-u --creator <user name of creator>`	Creator of the report.
`-rt --type <report_type>`	Creates the specified report type. For example, Activity Details report.

`--users <path of users' .csv file>`	Path of the .csv file containing the names of users in the user@domain,<user group> format. Specify the path to input user information for the report.
`--paths <path of .csv file of paths>`	Path of the .csv file containing the fully qualified paths of the data for which you want to create the report.
`--custodian <path of custodian' .csv file>`	Path of the .csv file that contains information about custodians on configured paths. The names of the custodians are specified in the .csv file in the format user@domain. Specify the path to the custodian.csv file to include custodian information in the report.

`-j` Lists the report jobs that are currently running.

`-l` Lists all configured reports.

`-o -m <TOP_N> -r <Report Name>`

Lists all report outputs. The following attributes apply:

`-m --max <TOP_N>`	Limits output to specified number of records, and lists the latest output first. If the number of records is not specified, prints status for the last run.
`-r - -report <Report Name>`	Prints the status of jobs for the specified report. You can either specify the report ID or the report name.
`-rt - -type<Report Type>`	Prints the status of jobs for the specified report type.

`report.exe -e [-d <Output_Dir> -r <Report Name> -w <Max_Wait>`

Executes report. The following attributes apply:

`-d --output <Output_Dir>`	The generated report output, including the SQLite database is copied to the specified directory. If you specify this option, you do not have to pass the –-w option.

`-r --report<Report Name>`	Executes the specified report. You can either specify the report ID or the report name.
`--w --wait <Max_Wait>`	Returns the report output only after the report execution is complete or the specified wait time in minutes is exceeded. Specify –1 to wait forever.

```
reportcli.exe -e -r <name of report> [-rt <report type>]--p <property
file path> --creator <user name of creator> [--output <Output_Dir>]
[--users <path of users' .csv file>] [--paths <path of .csv file >]
[--wait <MAX_WAIT>] [--custodian <path of custodian' .csv file>]
```

Executes a report using the properties file in which the input parameters are specified. The following attributes apply:

`-r --report <name of report>`	Creates a report with the specified name.
`-p -properties <property file path>`	Property file containing the input parameters for the report.
`-u --creator <user name of creator>`	Creator of the report.
`--custodian <path of custodian.csv>`	Path of the .csv file containing the names of the custodians
`-d --output <OUTPUT_DIR>`	Copies the report database and output to specified location.
`-rt --type <report_type>`	Creates the specified report type. For example, Activity Details report.
`--users <path of users' .csv file>`	Path of the .csv file containing the names of users in the user@domain,<user¦group> format.
`--paths <path of .csv file >`	Path of the .csv file containing the fully qualified paths of the data for which you want to create the report.
`-w --wait<MAX_WAIT>`	If the wait time is specified, the command returns only after the report is executed OR the specified wait time in minutes is exceeded. Specify -1 to wait forever. Data Insight cancels the report execution if the wait time is exceeded.

`--custodian <path of custodian' .csv file>`	Path of the .csv file that contains information about custodians on configured paths. The names of the custodians are specified in the .csv file in the format user@domain. Specify the path to the custodian.csv file to include custodian information in the report.

`Reportcli.exe –sa op_excl_path [–ap <FILE_PATH>] –at <ATTRIBUTE NAME> –av <ATTRIBUTE VALUE>] -ty <ATTRIBUTE TYPE> --path –r <name of report> -rt <report type>`

Creates a report by excluding the paths specified in the `.csv` file for computing open shares data on the Data Insight dashboard. The following attributes apply:

`– ap <file path>`	Path of the .csv file containing the fully qualified paths that you want to exclude from open share computation for the dashboard.
`-rt --type <report_type>`	Creates the specified report type.
`-r --report <name of report>`	Creates a report with the specified name.

You can exclude any path on a file server or a SharePoint server from the dashboard data computation. However, Data Insight does not support the exclusion of DFS paths using this method.

`report.exe –c –i <JOB_ID>`

Cancels execution of the specified report job.

scancli.exe

`scancli.exe` – scancli.exe - a utility that scans shares and site collections.

SYNOPSIS

```
scancli.exe --start¦ --stop¦ --list-jobs¦ --help [OPTIONS]
```

`-s --start`

Scans the specified shares or site collections.

`-c --stop`

Cancels the scans for specified shares or site collections.

`-l --list-jobs`

Lists currently running jobs.

`-d --display`

Displays the scan status for specified shares or site collections. To view real time scan queue information, use the `-l --list-jobs`option.

`-h --help`

Displays help.

Scan options

```
scancli.exe -s -a ¦ -f <FILERS> ¦ -m <SHARES> ¦ -S <SITECOLLS> ¦w <WEBAPPS>
[-D] [-e <EXCLUDE] [-F ¦ -N ¦ -p] [-I <INCLUDE>] [-i <DAYS>] [-t]
```

`-a - - all`

Scans all shares and site collections.

`-D - -disabled`

By default, disabled devices or those for which scanning has been disabled are not included in the scan. Specify this option to include shares or site collections of disabled devices.

`-e - -exclude <EXCLUDE>`

Exclude shares or site collections matching specified patterns. Separate multiple patterns with a comma. You can specify one or more wildcards in the pattern, for example, vol*,*$.

`-f - -filer <FILERS>`

Scans shares of the specified filers. For example, `-f - -filer >\\filer1, filer2, ID1,..>`.

`-F - -failed`

Select the shares or site collections whose last scan failed completely. This does not include those shares or site collections that have never been scanned before or those which succeeded partially (*).

`-I - -Include <INCLUDE>`

Include the shares or site collections matching the specified patterns. Separate multiple patterns with a comma. You can specify one or more wildcards in the pattern. For example, `-I - -Include >vol*,*$ >`

`-i - -interval <DAYS>`

Select the shares or site collections that have not been scanned for specified number of days. This includes shares or site collections which have never been scanned before (*).

`-m - -share <SHARES>`

Scans specified list of shares. For example, `-m - -share >\\filer1\share1, share2, ID3...>`.

`-n - -never`

Select the shares or site collections that have never been scanned before (*).

`-p - -partial`

Select the shares or site collections whose last scan succeeded partially, that is, those shares or site collections for which the scan is complete but with failure to fetch information for some paths (*).

`-S - -sitecoll <SITECOLLS>`

Scans the specified list of Microsoft SharePoint site collections.

`- t - -top`

Adds shares or site collections to top of the scan queue.

`-w - -webapp <WEBAPPS>`

Scans site collections for specified list of Microsoft SharePoint web applications.

Note: (*) indicates that the option can only be used on the Management Server.

Stop scan options

```
scancli.exe -l -a ¦ -f <FILERS> ¦ -m <SHARES> ¦ -S <SITECOLLS> ¦w <WEBAPPS>
[-D] [-e <EXCLUDE] [-I <INCLUDE>]
```

`-a - - all`

Stops scans for all shares and site collections.

`-e - -exclude <EXCLUDE>`

Exclude shares or site collections matching specified patterns. Separate multiple patterns with a comma. You can specify one or more wildcards in the pattern, for example, vol*,*$.

`-f - -filer <FILERS>`

Stops scans for shares of the specified filers.

`-I - -Include <INCLUDE>`

Include shares or site collections matching the specified patterns. Separate multiple patterns with a comma. You can specify one or more wildcards in the pattern.

`-m - -share <SHARES>`

Stops scans for the specified list of shares.

`-S - -sitecoll <SITECOLLS>`

Stops scans for the specified list of Microsoft SharePoint site collections.

`-w - -webapp <WEBAPPS>`

Stops scans for site collections for specified list of Microsoft SharePoint web applications.

List job options

```
scancli.exe -l [-n --node <NODE>]
```

`-n --node <Node ID or Node name>`

Lists scan jobs on the specified node. Specify either node ID or node name. If not specified, localnode is assumed.

Display options

```
scancli.exe -d -a ¦ -f <FILERS> ¦ -m <SHARES> ¦ -S <SITECOLLS> ¦w <WEBAPPS>
[-D] [-e <EXCLUDE] [-F ¦ -N ¦ -p] [-I <INCLUDE>] [-i <DAYS>]
```

`-a - - all`

Displays scan status for all shares and site collections.

`-e - -exclude <EXCLUDE>`

Exclude shares or site collections matching specified patterns. Separate multiple patterns with a comma. You can specify one or more wildcards in the pattern, for example, vol*,*$.

`-f - -filer <FILERS>`

Displays scan status for the shares of the specified filers.

`-F - -failed`

Displays scan status for the shares or site collections whose last scan failed completely. The scan status does not include those that have never been scanned before or those which succeeded partially (*).

`-I - -Include <INCLUDE>`

Include shares or site collections matching the specified patterns. Separate multiple patterns with a comma. You can specify one or more wildcards in the pattern.

`-i - -interval <DAYS>`

Displays scan status for the shares or site collections that have not been scanned for specified number of days. The scan status includes the shares which have never been scanned before (*).

`-m - -share <SHARES>`

Displays scan status for specified list of shares.

`-n - -never`

Displays scan status for the shares or site collections that have never been scanned before (*).

`-p - -partial`

Displays scan status for the shares or site collections whose last scan succeeded partially, that is, those shares or site collections for which the scan is complete but with failure to fetch information for some paths (*).

`-S - -sitecoll <SITECOLLS>`

Displays scan status for the specified list of Microsoft SharePoint site collections.

`-w - -webapp <WEBAPPS>`

Displays scan status for the site collections for specified list of Microsoft SharePoint web applications.

Note: `-w - -webapp <WEBAPPS>` option can only be used on the Management Server.

Examples

EXAMPLE 1: The following command scans all shares of a filer, netapp1.

```
scancli - -start - - filer <netapp1>
```

EXAMPLE 2: The following command scans all shares and site collections for which a full scan failed 3 or more days ago.

```
scancli - -start - -all - -failed - -ineterval <3>
```

The following command scans all site collections of a web application that have not been scanned for the past 30 days or have never been scanned.

```
scancli - -start - -webapp https://sitecoll:8080 - -interval 30
```

installcli.exe

`installcli.exe` – A utility that is used to configure multiple Windows File Servers and Data Insight worker nodes simultaneously.

SYNOPSIS

```
installcli [-w winnas_csv [-q]] [-n node_csv [-q]] [-p operation_token]
[-l] [-h]
```

Options

`-w --winnas winnas_csv`

Installs Data Insight Windows File Server agents and configures the corresponding filer.

-w option uses a `.csv` file with the following details as input:

- The host name or IP address of the Windows File Server that you want Data Insight to monitor.
- The host name, IP address, or ID of the Collector node that is configured to scan the filer.
- The host name, IP address, or ID of the Indexer node that is configured for the filer.
- The credentials that Data Insight should use to install the agent on the Windows File Server. The credential should be in the format user@domain. installcli.exe also accepts Local System credentials as value _LOCAL_. The same credentials must be added to Data Insight as a saved credential previously.
- True or false value indicating if the filer is clustered.
- The IP addresses of the agents. Separate multiple IP addresses with a semi-colon. If you do not want to use an agent to monitor the filer, indicate this with a hyphen (-).
- The credentials required to scan the filer. The credential should be in the format user@domain. The same credentials should be added to Data Insight as a saved credential previously.
- In case of a Windows File Server agent upgrade, RP or Full value indicating the type of upgrade you want to perform. This parameter is optional. Optionally, the name of the installer. If not specified, an appropriate one will be picked up from installers folder on the collector.

- True or false value indicating whether event monitoring should be enabled. For example, *winnas.company.com,collector.company.com,indexer.company.com, Administrator@DOMAIN,FALSE,winnas.company.com, Administrator@DOMAIN,TRUE,TRUE,RP,Veritas_DataInsight_windows_winnas_5_1_3002_x64.exe*.
- True or false value indicating whether the scan should be enabled according to the specified schedule.

`-n --node node_csv`

Installs the Data Insight Collector and Indexer nodes. The `node_csv` file must be in the following format:

- The host name or IP address of the worker node that you want to install or upgrade.
- True or false value indicating whether the node is a Collector. Since Data Insight does not currently support push-install on Linux nodes, you must specify *true* as that value for this column. I
- True or false value indicating whether the node is a Indexer.
- The credentials that Data Insight should use to install the package on the worker node. The credential should be in the format user@domain. The same credentials must be added to Data Insight as a saved credential previously.
- The port used by the DataInsightComm service.
- The port used by DataInsightConfig service.
- The destination directory where you want Data Insight to be installed.
- The location where you want to store the product data.

The values for the Communication service port, query service port, the installation path, and the data directory are optional. You can enter ? to use default values.

`-p --poll operation_token`

Starts polling for the latest status of an operation.

`-l --list`

Lists status and progress information of all currently running and historic operations.

`-q --nowait`

Forks off an operation and does not wait for it to complete.

`-h --help`

Displays help.

Appendix D

Data Insight jobs

This appendix includes the following topics:

- Scheduled Data Insight jobs

Scheduled Data Insight jobs

Each Data Insight service performs several actions on a scheduled basis. These services are called jobs. The section explains the function of the important jobs that run in various services. The schedule for few jobs can be changed from the **Advanced Settings** tab of the Server details page.

Table D-1 Communication service jobs

Job	Description
ADScanJob	Initiates the `adcli` process on the Management Server to scan the directory servers. Ensure the following: ■ The directory servers are added to the Data Insight configuration. ■ The credentials specified when adding the directory server have permissions to scan the directory server.
CollectorJob	Initiates the collector process to pre-process raw audit events received from storage devices. The job applies exclude rules and heuristics to generate audit files to be sent to the Indexers. It also generates change-logs that are used for incremental scanning.
ChangeLogJob	The CollectorJob generates `changelog` files containing list of changed paths, one per device, in the `changelog` folder. There cab be multiple files with different timestamps for each device. The ChangeLogJob merges all `changelog` files for a device.

Table D-1 Communication service jobs *(continued)*

Job	Description
ScannerJob	Initiates the scanner process to scan the shares and site collections added to Data Insight. Creates the scan database for each share that it scanned in the `data\outbox` folder.
IScannerJob	Intiates the incremental scan process for shares or site-collections for paths that have changed on those devices since the last scan.
CreateWorkflowDBJob	Runs only on the Management Server. It creates the database containing the data for DLP Incident Management, Entitlement Review, and Ownership Confirmation workflows based on the input provided by users.
DlpSensitiveFilesJob	Retrieves policies and sensitive file information from Data Loss Prevention (DLP).
FileTransferJob	Transfers the files from the `data\outbox` folder from a node to the `inbox` folder of the appropriate node.
FileTransferJob_content	Runs every 10 seconds on the Windows File Server. Routes content file and CSQLite file to the assigned Classification Server.
FileTransferJob_Evt	Sends Data Insight events database from the worker node to the Management Server.
FileTransferJob_WF	Transfers workflow files from Management Server to the Portal service.
FileTransferJob_classify	Runs on all Data Insight nodes once every minute. It distributes the classification events between Data Insight nodes.
IndexWriterJob	Runs on the Indexer node; initiates the idxwriter process to update the Indexer database with scan (incremental and full), tags, and audit data. After this process runs, you can view newly added or deleted folders and recent access events on shares on the Management Console.
ActivityIndexJob	Runs on the Indexer node; It updates the activity index every time the index for a share or site collection is updated. The Activity index is used to speed up the computation of ownership of data.
IndexCheckJob	Verifies the integrity of the index databases on an Indexer node.
PingHeartBeatJob	Sends the heartbeat every minute from the worker node to the Data Insight Management Server.
PingMonitorJob	Runs on the Management Server. It monitors the heartbeat from the worker nodes; sends notifications in case it does not get a heartbeat from the worker node.

Table D-1 Communication service jobs *(continued)*

Job	Description
SystemMonitorJob	Runs on the worker nodes and on the Management Server. Monitors the CPU, memory, and disk space utilization at a scheduled interval. The process sends notifications to the user when the utilization exceeds a certain threshold value.
DiscoverSharesJob	Discovers shares or site collections on the devices for which you have selected the **Automatically discover and monitor shares on this filer** check box when configuring the device in Data Insight
ScanPauseResumeJob	Checks the changes to the pause and resume settings on the Data Insight servers, and accordingly pauses or resumes scans.
DataRetentionJob	Enforces the data retention policies, which include archiving old index segments and deleting old segments, indexes for deleted objects, old system events, and old alerts.
IndexVoldbJob	Runs on the Management Server and executes the command voldb.exe --index which consumes the device volume utilization information it receives from various Collector nodes.
SendNodeInfoJob	Sends the node information, such as the operating system, and the Data Insight version running on the node to the Management Server. You can view this information on the Data Insight **Server** > **Overview** page of the Management Console.
EmailAlertsJob	Runs on the Management Server and sends email notifications as configured in Data Insight.The email notifications pertain to events happening in the product, for example, a directory scan failure. You can view them on the **Settings** > **System Overview** page of the Management Console.
LocalUsersScanJob	Runs on the Collector node that monitors configured file servers and SharePoint servers. In case of a Windows File Server that uses agent to monitor access events, it runs on the node on which the agent is installed. It scans the local users and groups on the storage devices.
UpdateCustodiansJob	Runs on the Indexer node and updates the custodian information in the Data Insight configuration.
CompactJob	Compresses the `attic` folder and `err` folders in `<datadir>\collector`, `<datadir>\scanner`, and `<datadir>\indexer` folders. The process uses the Windows compression feature to set the "compression" attribute for the folders. The job also deletes stale data that's no longer being used.
Compact_Job_Report	Compresses the folders that store report output.
StatsJob	On the Indexer node, it records index size statistics to `lstats.db`. The information is used to display the filer statistics on the Data Insight Management Console.

Table D-1 Communication service jobs *(continued)*

Job	Description
MergeStatsJob	Rolls up (into hourly, daily and weekly periods) the published statistics. On the Collector nodes for Windows Filer Server, the job consolidates statistics from the filer nodes.
StatsJob_Index_Size	Publishes statistics related to the size of the index.
StatsJob_Latency	On the Collector node, it records the filer latency statistics for NetApp filers.
SyncScansJob	Gets current scan status from all Collector nodes. The scan status is displayed on the **Settings** > **Scanning Dashboard** > **In-progress Scans** tab of the Management Console.
SPEnableAuditJob	Enables auditing for site collections (within the web application), which have been added to Data Insight for monitoring. By default, the job runs every 10 minutes.
SPAuditJob	Collects the audit logs from the SQL Server database for a SharePoint web application and generates SharePoint audit databases in Data Insight.
SPScannerJob	Scans the site collections at the scheduled time and fetch data about the document and picture libraries within a site collection and within the sites in the site collection.
NFSUserMappingJob	Maps every UID in raw audit files for NFS and VxFS with an ID generated for use in Data Insight. Or generates an ID corresponding to each User and Group ID in raw audit files received from NFS/VxFS.
MsuAuditJob	Collects statistics information for all indexers on the Indexer.
MsuMigrationJob	Checks whether a filer migration is in process and carries it out.
ProcessEventsJob	Processes all the Data Insight events received from worker nodes and adds them to the yyyy-mm-dd_events.db file on the Management Server.
ProcessEventsJob_SE	Processes scan error files.
SpoolEventsJob	Spools events on worker nodes to be sent to Management Server.
WFStatusMergeJob	Merges the workflow and action status updates for remediation workflows (DLP Incident Remediation, Entitlement Reviews, Ownership Confirmation), Enterprise Vault archiving, and custom actions and update the master workflow database with the details so that users can monitor the progress of workflows and actions from the Management Console.
UpdateConfigJob	Reconfigures jobs based on the configuration changes made on the Management Server.

Table D-1 Communication service jobs *(continued)*

Job	Description
DeviceAuditJob	Fetches the audit records from the Hitachi NAS EVS that are configured with Data Insight. By default, this job runs in every 5 seconds.
HNasEnableAuditJob	Enables the Security Access Control Lists (SACLs) for the shares when a Hitachi NAS filer is added. By default, this job runs in every 10 minutes.
WorkflowActionExecutionJob	This service reads the request file created on the Management Server when a Records Classification workflow is submitted from the Portal. The request file contains the paths on which an Enterprise Vault action is submitted. When the action on the paths is complete, the job updates the request file with the status of the action. By default, this job runs in every 1 hour.
UserRiskJob	Runs on each Indexer. The job updates hashes used to compute the user risk score. By default, the job runs at 2:00 A.M. everyday.
UpdateWFCentralAuditDBJob	Runs only on the Management Server. It is used to update the workflow audit information in `<DATA_DIR>/workflow/workflow_audit.db`. By default, this job runs every 1 minute.
TagsConsumerJob	Parses the CSV file containing tags for paths. Imports the attributes into Data Insight and creates a Tags database for each filesystem object. By default, this job runs once every day.
KeyRotationJob	Run the job on demand to change the encryption keys. It is not an automatically scheduled job. It is recommended to run this job after the Data Insight servers including Windows File Agent server is upgraded to 5.2. If you want to run the KeyRotationJob without upgrading all the servers, restart all services on the servers that have not been upgraded after the KeyRotationJob is executed and the configuration database is replicated on these servers.
RiskDossierJob	Runs on each Indexer and computes the number of files accessible and number of sensitive files accessible to each user on each share. This job runs every day at 11.00 P.M. by default.
ClassifyInputJob	Runs every 10 seconds on the Management Server. The job processes the classification requests from the Data Insight console and from reports for the consumption of the book keeping database.

Table D-1 Communication service jobs *(continued)*

Job	Description
ClassifyBatchJob	Runs every minute on the Indexer. The job splits the classification batch input databases for the scanner's consumption, which are later pushed to the Collector.
ClassifyIndexJob	Runs once every minute on the Indexer node. Updates the index with classification tags and also updates the status of the book keeping database.
ClassifyMergeStatusJob	Runs once every minute on the Management Server. The job calls the files with the classification update status that are received from each indexer. These files are automatically created on the indexer whenever updates are available. It also updates the global book keeping database that is used to show high level classification status on the Console.

The following processes run in the Data Insight WatchDog service

Table D-2 WatchDog service jobs

Job	Description
SyncPerformanceStatsJob -	Runs only on the Management server. Fetches performance related statistics from all other servers.
SystemMonitorJob	Gathers statistics like disk usage, CPU, memory usage.
SystemMonitorJob_backlog	Gathers statistics for unprocessed backlog files.
UpdateConfigJob	Reconfigures its own jobs based on configuration updates from the Management Server.

The following processes run in the Data Insight Workflow service

Table D-3 Workflow service jobs

Job	Description
WFStepExecutorJob	Processes actions for Enterprise Vault archiving, requests for permission remediation, and custom actions configured in Data Insight.
WFStepExecutorJob_im	Processes workflows of type Entitlement Reviews, DLP Incident Remediation, and Ownership confirmation. It also sends email reminders containing links to the remediation portal to the custodians at a specified interval.
UpdateConfigJob	Updates its schedules based on the configuration changes made on the Management Server.
WFSpoolStatusJob	Reads the workflow data every minute, and if there are any new updates in last minute, it creates a status database with the new updates.
FileTransferJob_WF	Transfer workflow status databases from the Self-Service portal nodes to the Management Server.

The following processes run in the Data Insight Webserver service.

Table D-4 Webserver service jobs

Job	Description
CustodianSummaryReportJob	Periodically runs the custodian summary report, which is used to determine the custodians assigned in Data Insight for various resources. The output produced by this report is used in DLP Incident Remediation, Entitlement Review, and Ownership Confirmation workflows.
HealthAuditReportJob	Periodically creates a report summarizing health of the entire deployment, and stores it to `log/health_audit` folder on the Management Server. The report aids Veritas Support to troubleshoot issues on your setup.

Table D-4 Webserver service jobs *(continued)*

Job	Description
PolicyJob	Evaluates configured policies in the system and raises alerts.
PurgeReportsJob	Deletes older report outputs.
UpdateConfigJob	Updates configuration database on the worker nodes based on the configuration changes made on the Management Server.
UserIndexJob_merge	Consolidates user activity and permission map from all indexers.
UserIndexJob_split	Requests each Indexer for user activity and permission map.
UserRiskMergeJob	This job runs on the Management Server. Its default schedule is 6:00 A.M. every day. The job combines data from all MSUs into a single risk score value for each user. This job creates the `userrisk_dashboard.db` in the `DATA_DIR\conf` folder.

The following processes run in the Data Insight Classification service.

Table D-5

Job	Description
ClassifyFetchJob	Runs every minute on the server that is assigned the role of a Classification Server. It searches the `classification/inbox` folder for the input files and adds them to the priority queues. One input file can result in multiple snapshots with the name `<PRIORITY>_<CRID>_<BATCHID>_<NODEID>_<MSUID>_<TIMESTAMP>_snap<N>.csqlite`. The input file contains the location where the actual file has been kept in the classification/content folder. The job also keeps a list of files that could not be fetched. **Note:** Error logs are created in the `<Install directory>/log/fetch folder`.

Table D-5 *(continued)*

Job	Description
ClassifyFetchPauseJob	Runs once every minute on any node that acts as the Classification Server. Refreshes the pause or resume status of fetch jobs as per the duration configured for content fetching.
CancelClassifyRequestJob	Runs every 20 seconds in Communication Service and Classification Service. Fetches the list of classification requests that are cancelled and distributes this request between Data Insight nodes. Before classifying files, all the classification jobs consult this list to identify the requests that are marked for cancellation. If they observe any canceled request in the new request that is submitted for classification, then that request is deleted.
ClassifyJob	Runs once every minute on any node that acts as a Classification Server. Checks the `classification/inbox` folder for input files submitted for classification folder and adds them to three separate priority queues. It picks a file from the highest queue in FIFO order, and starts classifying content using Veritas Information Classifier. All files in that input file are submitted for classification. Once all paths in the file have been classified, result of the classification and any resulting errors are written to a database in the `classification/outbox` folder.
UpdateVICPolicyMapJob	Runs every ten seconds on the Management Server. It ensures that Data Insight configuration database is in sync with the Classification Policy Manager.
UpdateConfigJob	Reconfigures jobs based on the configuration changes made on the Management Server.

Table D-5 *(continued)*

Job	Description
CreateFeaturesJob	Runs once every week on Sunday at 00.01 A.M. on the Indexer. Checks if sufficient classified data is available for the supervised learning algorithm to create predictions (training sets). The job has a multi-threaded execution framework which executes actions in parallel. The default thread count is 2. You can set the value using the `matrix.classification.sl.features.threads` property at global or node level. **Note:** The node level property always takes precedence over the global level property.
PredictJob	Runs once every week on Sunday at 05.00 A.M. on the Indexer. Copies the prediction files from the temp output directory to a classification outbox.
SLCreateBatchesJob	Runs every 2 hours on the Indexer. The job creates batches of files for the consumption of Veritas Information Classifier. These files are classified with high priority.

See "Monitoring Data Insight jobs" on page 96.

Appendix E

Troubleshooting

This appendix includes the following topics:

- About general troubleshooting procedures
- About the Health Audit report
- Location of Data Insight logs
- Downloading Data Insight logs
- Migrating the data directory to a new location
- Enterprise Vault exceptions and their meanings
- Troubleshooting FPolicy issues on NetApp devices
- Troubleshooting EMC Celera or VNX configuration issues
- Troubleshooting EMC Isilon configuration issues
- Troubleshooting SharePoint configuration issues
- Troubleshooting Hitachi NAS configuration issues

About general troubleshooting procedures

This section provides an overview of the general troubleshooting procedures that you can use to help discover and troubleshoot common problems.

You can use the **Events** page on the Data Insight Management Console to get a quick overview of the node on which the error has occurred.

To troubleshoot a problem, it helps to consider the following:

- Check for prior occurrence.

Check existing troubleshooting information to see if the problem has occurred before and if there is a workaround available to troubleshoot the same. A good source for this type of information is the *Veritas Data Insight Release Notes*. The Release Notes contain a list of known issues for Data Insight and a list of possible workaround.

- Consider recent alterations.
  If a system is having problems immediately after some kind of maintenance, software upgrade, or other change, the problem might be linked to those changes.

Data Insight runs a daily health audit report that helps you identify potential problems in your environment.

About the Health Audit report

The Health Audit report gives you an overall status of the heath of Data Insight deployment. You can use the report to identify potential problems in the system and proactively take measures to fine tune the system settings. The report aids Veritas Support to troubleshoot issues on your setup.

The HealthAuditReportJob periodically creates the report at 5:00 A.M. daily and stores it on the `log/health_audit` folder.

Location of Data Insight logs

Veritas Data Insight log files are located in the Data Insight installation directory, `<INSTALLDIR>\log`. Typically, the installation directory is located at `C:\Program Files\DataInsight\log`. On Linux, the logs are located at `/INSTALL/DataInsight/log`

Table E-1 describes the logs that are relevant for troubleshooting.

Table E-1 Data Insight logs

`webserver0.0.log`	This file contains the log messages from the Web service process.
`commd0.0.log`	This file contains the log messages from the scheduler communication service.
`adcli.log`	This file contains the log messages from the Active Directory scanner process, `adcli.exe`.
`celerrad.log`	This file contains the log messages for DataInsightCelerra service.

Table E-1 Data Insight logs *(continued)*

`cli0.0.log`	This file contains the log messages for various command line utilities.
`collector.log.N`	This file contains the log messages for the audit pre-processor (collector.exe).
`dashboard.log`	This file contains the log messages for the Dashboard data generation report.
`dscli0.0.log`	This file contains the log messages for LDAP, NIS, NIS+ Directory scanner.
`scanner/extN_msuN.log`	This file contains the log messages for Full file system scans.
`scanner/extN_msuN.ilog`	This file contains the log messages for Incremental file system scans.
`sv3_<extac_id>_<msu_id>_<process_id>_<timestamp>.log`	This file contains the log messages for parallel scanner. **Note:** These logs are located in <*INSTALLDIR*>\log\scanner folder.
`fpolicyd.log`	This file contains the log messages for DataInsightFpolicy service.
`idxcheck.log`	This file contains log of index integrity check.
`indexer/index-N.log`	This file contains log messages for index updater process.
`localusers.log`	This file contains log messages for local users scanning process.
`mxpolicy.log`	This file contains log messages for Data Insight policy evaluation process.
`queryd.log`	This file contains log messages for DataInsightConfig service.
`sharepoint_audit.log`	This file contains log messages for SharePoint audit fetching utility.
`upgradecli0.0.log`	This file contains log messages of the upgrade utility.
`upgrade_to_X.log`	This file contains log messages of the upgrade process.
`vic.log`	This file contains log messages of Veritas Information Classifier.

Table E-1 Data Insight logs *(continued)*

`vic.0.0.log`	This file contains log messages of the interaction that happens between Veritas Information Classifier and Data Insight.
`watchdog0.0.log`	This file contains log messages of DataInsightWatchdog service.
`winnasd.log.0`	This file contains log messages of DataInsightWinnas service.
`winnas_util.log`	This file contains log messages of windows share discovery utility.
`workflowd0.0.log`	This file contains log messages from the DataInsightWorkflow service.
`birt.log`	This file contains log messages from the BIRT utility **Note:** This log file shows an exception “Could not load SQLite driver,” however this is by design and should be ignored as the SQLite driver is infact loaded.
`vic.log`	This file contains log messages of Veritas Information Classifier.
`vic.0.0.log`	This file contains log messages of the interaction that happens between Veritas Information Classifier and Data Insight.

Downloading Data Insight logs

To troubleshoot errors, you can download the Data Insight logs relevant to a file server, share, SharePoint web application, or SharePoint site collection from the **Settings** tab of the Management Console.

To download Data Insight logs

1. On the relevant listing page, click the **Select Action** drop-down, and select **Download Logs** for the data repository you want to troubleshoot.
2. On the **Download Logs** pop-up, select the check box for the information that you want to include in the logs.

   You can select one or all of the following information:

   - **Config database** - Select this option to include the configuration database in the download. Secret information, such as passwords are purged from the copied database.

- **Indexer database** - Select this option to include the index for the problematic shares or site collections in the download.
- **Error files** - Select this option to include scan or audit files that have not been indexed in the download.
- **User database** - Select this option to include the cached Active Directory information in the download.

Note: Contact Veritas Support to help you determine which of these options you should select when troubleshooting an issue.

Migrating the data directory to a new location

The data directory is the location where a Data Insight server stores the product data. You specify the location of the data directory during the Data Insight installation. You can find out the current location of the data directory by reading the `datadir.conf` file that is located at `C:\Program Files\Veritas\DataInsight\`.

To move the data directory from its current location

1. Stop all the Data Insight services that are running on the server.
2. Navigate to `C:\Program Files\Veritas\DataInsight\` and open the file `datadir.conf` in a text editor. Note the current location of the data directory. For example, `matrix.datadir=C:/DataInsight/data.`
3. Edit the value for the parameter *matrix.datadir* to indicate the new location of the data directory. For example, `matrix.datadir=E:/DataInsight/data.`
4. Copy the folder,`$DATADIR/data`, from the old location to the new location. For example, copy the folder from the original location `C:/DataInsight` to the new location `E:/DataInsight`.

   Note: If you choose to rename the data directory, do not use any space in the filename. Doing so will prevent the Data Insight services from starting.

   ---

5. Navigate to `C:\Program Files\Veritas\DataInsight\bin` using the command prompt. Execute the following command:

   ```
   configdb -c
   ```

6. Verify that the command output points to the new data directory location.

7 Execute the command `configdb -p -T node.`

Verify that the command output lists all the Data Insight servers that are in your deployment.

8 Start the Data Insight services on the server.

9 After all Data Insight services start successfully, delete the original data directory.

Enterprise Vault exceptions and their meanings

Sometimes, when you initiate an archive operation from the Data Insight Management Console, using the Enterprise Vault, you may encounter error messages. Here is a list of possible error messages and their description:

Table E-2 Enterprise Vault errors and descriptions

Exception	Description
`Exception in method Archive:System. ServiceModel.CommunicationException: The maximum message size quota for incoming messages (65536) has been exceeded. To increase the quota, use the MaxReceivedMessageSize property on the appropriate binding element. ---> System.ServiceModel. QuotaExceededException: The maximum message size quota for incoming messages (65536) has been exceeded. To increase the quota, use the MaxReceivedMessageSize property on the appropriate binding element.`	This error can arise when a large number of files are sent for archiving in a single batch. This causes an overflow of the product buffer.

Table E-2 Enterprise Vault errors and descriptions *(continued)*

Exception	Description
`Exception in method GetFileServerShares: System.ServiceModel.Endpoint NotFoundException: There was no endpoint listening at http://10.209.108.73/EnterpriseVault/FileSystemArchivingAPI that could accept the message. This is often caused by an incorrect address or SOAP action. See InnerException, if present, for more details. ---> System.Net.WebException: The remote server returned an error: (404) Not Found. 2013-05-07 03:37:51 INFO: #{27} [ArchiveStep.archiveFiles] [ARCHIVE OUT] --- End of inner exception stack trace --- 2013-05-07 03:37:51 INFO: #{27} [ArchiveStep.archiveFiles] [ARCHIVE OUT] at System.Service Model.Channels. HttpChannelUtilities.ProcessGet ResponseWebException(WebException webException, HttpWebRequest request, HttpAbortReason abortReason)`	This error occurs when the Enterprise Vault Admin service is down or the Enterprise Vault server is unreachable. For instance, when the disk on the Enterprise Vault server vault store is full, all the Enterprise Vault services are stopped. This could also occur if the name of Enterprise Vault server, protocol, or port are specified incorrectly. If the AdminService is not down, still this error can arise if the EVFileSvrArcMngr service is not running.

Table E-2 Enterprise Vault errors and descriptions *(continued)*

Exception	Description
`[Exception in method Archive:System.Service Model.FaultException` `1[www.symantec.com.Enterprise Vault.API.FileSystemArchiving. Data.ExecutionFailedFault]: This request operation sent to net.tcp://evserver1.tulip.matrixad .local:5114/TaskService/ 10E7B53932AB5754980C95B4591BC 72171012f00evserver1 did not receive a reply within the configured timeout (00:30:00). The time allotted to this operation may have been a portion of a longer timeout. This may be because the service is still processing the operation or because the service was unable to send a reply message. Please consider increasing the operation timeout (by casting the channel/proxy to IContextChannel and setting the OperationTimeout property) and ensure that the service is able to connect to the client. (Fault Detail is equal to www.symantec.com.EnterpriseVault. API.FileSystem Archiving.Data.Execution FailedFault).]`	This error occurs when the internal services of the Enterprise Vault times out when you try to archive a batch of files. To resolve this problem, you can consider changing the size of batch to a smaller value. Also you can consider reducing the number of files that are sent in a batch. Alternatively, you can change the timeout values in the configuration for the Enterprise Vault. To change the timeout value: **1** Open the `EvFileSvrArcMngr.exe.config` file. **2** Set the appropriate value for the key “<add key="OperationTimeoutInMin" value = "60"/>

Table E-2 Enterprise Vault errors and descriptions *(continued)*

Exception	Description
`Executing Function Archive: Exception in method Archive:System.ServiceModel.Fault Exception`1[www.symantec.co m.EnterpriseVault.API.FileSystem Archiving.Data.TimeoutFault]: The File System Ar chiving task service failed to start. Check that the File System Archiving task service is enabled in the configuration file, <Enterprise_Vault_installation_fol der>\EvFSAArchiving Task.exe.config. (Fault Detail is equal to www.symantec.com.E nterpriseVault.API.FileSystem Archiving.Data.TimeoutFault).`	This error occurs at random. To resolve this problem, re-run the failed archive operation. If you are an Data Insight administrator user, you can re-run the archive operation from the **Settings** > **Action Status** page of the Management Console.See "Viewing and managing the status of an operation" on page 323.
`2013-03-21 11:17:27 WARNING: Archive: Got exception while archiving - System.ServiceModel.Fault Exception`1[www.symantec.com.Enterprise Vault.API.FileSystem Archiving.Data.ServerTemporary UnavailableFault]: Unable to contact the Enterprise Vault Task Controller service. Check that the service is running. (Fault Detail is equal to www.symantec.com.Enterprise Vault.API.FileSystem Archiving.Data.ServerTemporary UnavailableFault).`	This error occurs when the Enterprise Vault task controller service cannot be reached for some reason.

Table E-2 Enterprise Vault errors and descriptions *(continued)*

Exception	Description
Exception in method Archive:System.ServiceModel.Fault Exception`1[www.symantec.com.Enterprise Vault.API.FileSystem Archiving.Data.ExecutionFailedFault]: Internal error occurred. Internal Error : <Error checking if file: \\?\UNC\vnxemc.tulip.matrixad.local\ Share4EV\Public Share Copy\Ganesh\di automaation \logs3\DI_Automation\Library\Copy (3) of KeyAction\Action1\SnapShots\ ssf35f0.html.z is a placeholder file.> (Fault Detail is equal to www.symantec.com.Enterprise Vault.API.FileSystem Archiving.Data.Execution FailedFault)	This error occurs when you attempt to re-archive an already archived file, but the placeholder service on the File System Archiving (FSA) agent is not running. To resolve this problem, run the placeholder services on the FSA agent and attempt to archive the file again.
[Exception in method Archive: System.Web.Services.P rotocols.SoapException: The socket connection was aborted. This could be caused by an error processing your message or a receive timeout being exceeded by the remote host, or an underlying network resource issue. Local socket timeout was ''''00:29:59.9810000''''., at evClient.Program. Archive(FileSystemArchivingService channel)]	This error occurs when you restart the Data Insight services, while Data Insight is still processing an archiving operation. To resolve this error, make sure that the File System Archiving (FSA) task and Enterprise Vault services are running.

Troubleshooting FPolicy issues on NetApp devices

Data Insight uses the FPolicy framework to collect information about access events from the NetApp filers. To receive access events from the NetApp filer, Data Insight uses a process called the DataInsightFPolicy service.

The following occurrences indicate that the DataInsightFPolicy service is unable to communicate with the NetApp filer:

- Frequent disconnections between the Collector node and the NetApp filer.
- No events are registered for a filer, even though the filer is properly configured.
- Warning or errors are displayed about connection problems.

Viewing FPolicy-related errors and warnings

To view the FPolicy related errors and warnings:

1. From the Management Console, navigate to **Settings** > **Filers**. The list of all the filers added to Data Insight is displayed.
2. Click the name of the filer to view the overview page for the filer.
3. Review the **Filer Health** section.
4. If FPolicy has disconnected from the filer, you will see an error message to that effect.
5. Alternatively, navigate to **Settings** > **Events**. From **Events** tab review the event-related errors and warnings which may indicate FPolicy connectivity problems.
6. To know more details about the connection problems you can also examine the `fpolicyd.log` file. You can locate the log file under the log folder of the Data Insight installation directory.

Resolving FPolicy connection issues

To resolve FPolicy connectivity issues perform the following checks:

- Check the network connectivity. See if you can connect to the Collector from the filer using the short name or the Fully Qualified Host Name of the Collector and vice versa.
- Ensure that host name or IP DNS lookup is working. The DNS lookup and reverse-lookup for host name of the Collector node from the filer should be working fine.
- Verify that the system-time on the Collector and the filer are within 5 minutes of the time on the Domain Controller.
- Check that the firewall settings or security software on the Collector are not blocking port numbers 139 and 445.
- Verify that the user whose credentials are used to configure the DataInsightFPolicy service belongs to the same domain as the Data Insight

Collector node and the NetApp filer. Ensure that this user is a part of the Backup Operators group on the filer.

- If the NetApp filer and Data Insight Collector node are in different domains, ensure that the DataInsightFPolicy service is running as Local System or there is bidirectional trust between the domains to which the filer and the Collector belong.

Once you performed all the checks mentioned in the earlier procedure, you might need to perform the following additional checks if the problem persists:

- Network Settings on the Collector: If you are using a Windows 2008 machine as a Collector, verify the local security policy for the named pipes that can be accessed anonymously.
  See "To verify the correct network security policy for FPolicy" on page 391.
- Setting entry in the hosts file on the filer: The hosts file entry on the filer and the format of the hosts file entry. The entry should be in the format shown:
  ```
  <IP_ADDRESS> <FQDN> <Short_name>
  ```
  For example,
  ```
  10.209.88.190 mycollector.tulip.matrixad.local mycollector
  ```
- SMB signing check on the filer and the Collector: Disable the SMB signing on the filer and the Collector.

To verify the correct network security policy for FPolicy

1. On the Collector node, navigate to the Control Panel.
2. Click**System and Security** > **Administrative tools** > **Local Security Policy** > **Local Policies** > **Security options**.
3. Check if the value for the policy called "Network access: Named Pipes that can be accessed anonymously" is NTAPFPRQ.
4. Restart the Collector node if you have made any changes.

Troubleshooting EMC Celera or VNX configuration issues

Data Insight registers with the EMC Celerra or the EMC VNX filer through the EMC Common Event Enabler (CEE) framework. The CEE framework can be installed on the same Windows server as the Data Insight Collector node or on a remote server in the same Active Directory domain. To fetch audit events, the computer running CEE must be able to maintain connectivity with the filer being monitored.

To monitor the activities of the CEE server, you can enable Verbose and Debug modes for the server.

To enable Verbose mode and Debug mode on the CEE computer

1. Open the Windows Registry Editor (Start > Run > regedit)
2. Navigate to:

   HKLM>**Software**>**EMC**>**CEE**>**Configuration**

   Set verbose to 3f.

   HKLM>**Software**>**EMC**>**CEE**>**Monitor**>**Configuration**

   Set debug to 3f.

   Once your debuggung session is over, set the values of the changed registry keys to 0.

Table E-3 Troubleshooting EMC Celera or VNX configuration issues

Issues/Indications	Checks	Steps to perform
Connectivity issues between the EMC filer and the host computer running EMC CEE service.	Ensure that the system time on the filer and the Windows host running the CAVA or the DataInsightCelerra service are in sync with the domain controller.	Perform the following steps: ■ Log in to the EMC Control Station using SSH or Telnet. ■ Execute the command for the corresponding data mover: `server_date` *`server_2`* `timesvc start ntp` *`<domain-controller>`* where *server_2* is the name of the data mover.
	Ensure that the CIFS servers can communicate with the domain controller.	To test the connectivity of the CIFS server with the domain controller: ■ Log in to the EMC Control Station using SSH or Telnet. ■ Execute the following command for the corresponding data mover: `server_cifs` *`<server_2>`* where *server_2* is the name of the CIFS server. ■ Verify that the output of the command contains the following line: FQDN=<Fully qualified name of the CIFS server> (Updated to DNS) Appearance of the line is the output confirms that connection is established.
		To check `cepp.conf` file entry: ■ Log in to the EMC Control Station using SSH or Telnet. ■ Execute the following command for the corresponding data mover: `server_file server_2 -get cepp.conf` ■ If you find any entry for msrprcuser, remove the entry. Ensure that you have configured the EMC CAVA service and DataInsightCelerra service to run using a domain user account with administrative privileges.

Table E-3 Troubleshooting EMC Celera or VNX configuration issues *(continued)*

Issues/Indications	Checks	Steps to perform
	Ensure that you have the following entry in the `cepp.conf` file: cifsserver=<CIFS server name> surveytime=90 pool name=matpool \ servers=<IP address of the host running CAVA service> \ postevents=* \ option=ignore \ reqtimeout=500 \ retrytimeout=50 Note that there is a space before the '\' character at the end of a line. Also note that there is no '\' character on the first and last line.	
Error when you attempt to start the CEPP service on the filer	To enable the EMC VNX or Celerra filer to send event information to Data Insight you must start the CEPP service on the filer. When you attempt to start the CEPP service you may encounter the error 13162905604.	To troubleshoot error while trying to start CEPP service on a VNX filer: ■ Log in to the EMC Control Station using SSH or Telnet. ■ View the `cepp.conf` file by executing the following command: `server_file server_2 -get cepp.conf` ■ Note the IP address mentioned in the servers section. ■ Ensure that the noted IP address is accessible and resolvable from the EMC filer Control Station.

Troubleshooting EMC Isilon configuration issues

Sometimes, Data Insight may fail to fetch audit events from an Isilon filer.

The Table E-4 describes some common issues and troubleshooting steps to resolve these issues.

Table E-4 Troubleshooting EMC Isilon configuration issues

Issues/Indications	Troubleshooting actions

Data Insight is unable to fetch audit events from the EMC Isilon filer.	■ Ensure that the Isilon hostname configured in Data Insight is the same as configured in the Isilon Console audit settings.
	■ Install Microsoft DebugView on the computer where you have installed EMC CEE.
	To print debug information modify the following registry settings and change their respective 'Debug' and 'Verbose' keys to value '3f':
	■ [HKEY_LOCAL_MACHINE\SOFTWARE\EMC\CEE\Configuration]
	■ [HKEY_LOCAL_MACHINE\SOFTWARE\EMC\CEE\Monitor\Configuration]
	■ On the Data Insight Management Server, set the Collector log level to FINE using the following command:
	`configcli backend_loglevel <node id of the Collector> collector FINE`
	You can use the following command to get the node ID:
	`configdb -p -T node`
	■ After you have set the Collector log level to FINE, you must restart DataInsightCelerra service on the Collector node. To change the log level back to INFO use the following command on the Management console:
	`configcli backend_loglevel <node id of the Collector> collector INFO`
	■ Restart DataInsightCelerra Service on the Collector.

Troubleshooting SharePoint configuration issues

Sometimes during the agent installation, your installer may return an error. Ensure that you have uninstalled the existing installation of the Data Insight SharePoint Agent before you upgrade the agent to a newer version.

The Table E-5 explains some useful troubleshooting steps:

Table E-5 Troubleshooting SharePoint configuration issues

Issues/Indications	Steps to perform

Issues when installing the Data Insight SharePoint agent on the SharePoint server	Perform the following steps:
	■ Log in to the SharePoint **Central Administration Console** and set the ULS log level to **Verbose**.
	■ Collect the Event Viewer logs and ULS logs from the following locations:
	  ■ For SharePoint 2007:
	  `C:\Program Files\Common Files\Microsoft Shared\Web Server Extensions\12\LOGS`
	  ■ For SharePoint 2010:
	  `C:\Program Files\Common Files\Microsoft Shared\Web Server Extensions\14\LOGS`
	  ■ For SharePoint 2013:
	  `C:\Program Files\Common Files\Microsoft Shared\Web Server Extensions\15\LOGS`
	  ■ For SharePoint 2016:
	  `C:\Program Files\Common Files\Microsoft Shared\Web Server Extensions\16\LOGS`
	■ Collect the Data Insight logs from the Collector node.
	■ Send the logs to Support for assistance.

The computer stops responding during the agent installation	Perform the following steps:
	■ Log in to the SharePoint **Central Administration Console**.
	■ Navigate to the **Solutions** page and check the state of the package called **sdispwebsvc.wsp**.
	  ■ If the state is **Deployed**, view the install log file to determine the cause of the problem.
	  ■ If the solution is stuck in **Deploying** state, go to the **Timer Job Definitions** page and delete the Timer Job definitions for the package. Check the state of the package again. If the state is not changed to **Deployed**, contact Support.
	  ■ If the state changes to **Deployed** after terminating the Timer job, check if the SharePoint feature, called **SDIEventHandler** is installed. To verify this browse to one of the monitored site collections URL. Navigate to **Site Actions** > **Site Settings** > **Site Collections Features**. Search for the keyword 'SDI' to locate the feature. Verify if the feature **SDIFeatureStapler** is in the **Active** state. Otherwise make it active.
	  ■ If the features are not visible, collect all the install logs and contact Support.
	  The install and uninstall logs are located in the Windows temp folder. The log file names start with the characters i4j**.

Table E-5 Troubleshooting SharePoint configuration issues *(continued)*

Issues/Indications	Steps to perform
Test Connection fails for a web application	In the Data Insight Management Console, when you click the **Test Connection** for a web application, Data Insight attempts to do the following: ■ Discover site collections ■ Check for scanning ■ Checks for auditing Ideally, all the three tasks should succeed. If any of the three tasks fail, look for the following log files for further troubleshooting: ■ sharepoint_util.log ■ sharepoint_audit.log ■ sharepoint_scanner_0.log
Errors during the addition of web applications	After you configure a web application, Data Insight discovers and adds site collections to the configuration automatically. Data Insight also enables the auditing flags on the site collections. To troubleshoot any errors during the addition of web applications, view the `sharepoint_util.log` file.

Table E-5 Troubleshooting SharePoint configuration issues *(continued)*

Issues/Indications	Steps to perform

Agent uninstallation issues	Perform the following steps:
	■ Navigate to the following locations and remove the files mentioned:
	■ Global Assembly Cache (GAC) under `C:\Windows\assembly` :
	■ `sdispwebsvc`
	■ `SDIEventHandler`
	■ SDIFeatureStapler and SDIEventHandler folders under `C:\Program Files\Common Files\Microsoft Shared\Web Server Extensions\12\Template\Features`
	■ `Elements.xml`
	■ `Feature.xml`
	■ `C:\Program Files\Common Files\Microsoft Shared\Web Server Extensions\12\ISAPI`
	Note that this location may differ slightly for different SharePoint versions.
	■ `sdispwebsvc.asmx`
	■ `sdispwebsvcdisco.aspx`
	■ `sdispwebsvcwsdl.aspx`
	■ Navigate to `C:\Program Files\Veritas\DataInsight\bin`, and execute the command:
	`SPManageFeature.exe -d`
	Run this utility as Farm Administrator.
	■ To remove features, run the following commands:
	`stsadm -o uninstallfeature -filename SDIFeatureStapler\Feature.xml`
	`stsadm -o uninstallfeature -filename SDIEventHandler\Feature.xml`
	■ To remove the DLL files from the Global Assembly Cache (GAC):
	■ Navigate to `C:\Windows\assembly`
	■ Right-click **sdispwebsvc** and choose **Uninstall**.
	■ Right-click **SDIEventHandler** and choose **Uninstall**.
	■ Repeat this step on all the SharePoint computers.
	■ Delete the following files:
	■ `Elements.xml`
	■ `Feature.xml`
	■ `sdispwebsvc.asmx`
	■ `sdispwebsvcdisco.aspx`
	■ `sdispwebsvcwsdl.aspx`
	■ Repeat this step on all the SharePoint computers.

Troubleshooting Hitachi NAS configuration issues

The Table E-6 explains some useful troubleshooting steps:

Table E-6 Troubleshooting Hitachi NAS configuration issues

Issues/Indications	Steps to perform
Error when Data Insight attempts to discover shares on the Hitach NAS (HNAS) filer.	Refer to the following log: `winnas_util.log`. You can find this log at the following location of your Collector node assigned to the Hitachi NAS file server: <*INSTALLDIR*>\log <INSTALLDIR> refers to the installation directory which is usually located in: `C:\Program Files\DataInsight\`
Data Insight Collector node is not able to fetch audit events from the HNAS filer.	Refer to the following log: `hansd_<filer_id>.log`

To view the audit logs that are generated on the Hitachi NAS filer	Perform the following steps:
	■ Using SSH or Telnet log in to Admin EVS.
	■ Switch to the EVS you are monitoring.
	■ Execute the following commands:
	`console-context --evs <EVS Name>`
	`audit-log-show <File System Name>`

To verify if Data Insight is able to fetch the audit events from the Hitachi NAS filer	Perform the following steps: ■ Navigate to the folder: `<datadir>\collector` <datadir> represents the data directory of Data Insight, which is located at: `C:\DataInsight\data` by default. ■ Look for the `hnas*.sqlite` files. The presence of such files indicates that Data Insight is able to successfully fetch audit events from Hitachi NAS filer.

Index

A

Active Directory domain scans
    scheduling 119
adding exclude rules 35
archiving
    adding new Enterprise Vault servers 320
    filer mapping 322
    managing the Enterprise Vault servers 320
    overview 318
archiving data
    overview 37

B

business unit mappings
    configuring 119

C

Classification Server
    registering 74
Clustered NetApp filers
    about configuration 147
configuring
    DFS target 246
    EMC filers 167
    SMB signing 133
    Windows File Server 193
    Workspace data owner policy 53
configuring product users
    reviewing current users and privileges 61
    Symantec Data Loss Prevention users 64
containers
    adding 122
    managing 123
    overview 122
current users and privileges
    reviewing 61

D

Data Insight
    adding exclude rules 35
Data Insight *(continued)*
    preparing to receive event notification 171
data retention
    configuring 38
deployment
    moving data directory 384
DFS utility
    overview 247
    running 248
directory domain scans
    overview 108
directory servers
    adding 109
    managing 114
directory service domain
    deleting 118

E

EMC Celerra filers
    configuration credentials 173
    preparing for CEPA 168
events
    configuring scanning 28
    email notifications configuring 291
    enabling Windows event logging 292

F

filers
    Add/Edit EMC Celerra filer dialog 220
    add/edit geneic device dialog 232
    Add/Edit Hitachi NAS file server 233
    Add/Edit NetApp filer dialog 211
    Add/Edit VxFS filer dialog 229
    Add/Edit Windows File Server dialog 225
    adding 210
    deleting 238
    editing configuration 236
    migrating 236
    viewing 208
Fpolicy
    overview 134

Fpolicy *(continued)*
preparing NetApp filer 136
preparing NetApp vfiler 138
preparing Veritas Data Insight 135

G

generic device
scanning credentials 205

H

Hitachi NAS filers
about configuration 188
configuration credentials 189

I

importing
custom attributes 120
DFS mappings 249
Indexers
Migration 103

L

licenses
managing 23

M

Management Console
configuring global settings 54
operation icons 14
Management Server
configuring SMTP settings 23
managing
undesired data 318

N

NetApp cluster
capabilities for adding non-administrator domain user 155
preparing non-administrator domain user 155
NetApp filers
capabilities for adding non-administrator domain user 141
configuration credentials 128, 132
handling events from home directories 145
preparing non-administrator domain user 141
prerequisites 127

O

operation progress status
tracking and managing 323
overview
administering Data Insight 14
configuring filers 208
DFS target 245
filtering accounts, IP addresses, and paths 33

P

patches and upgrades
viewing and installing recommendations 100
Permission remediation
About 314
exclude rules 317
managing 315
policies
managing 298
overview 296
preparing
EMC Celerra filer 168
NetApp filer for Fpolicy 136
NetApp vfiler for Fpolicy 138
Veritas Data Insight for Fpolicy 135
product user
adding 62
product users
deleting 64
editing 63
product users and roles
overview 60
purging data
overview 38

R

renaming
filer 276
SharePoint web application 276
storage device 276

S

saved credentials
managing 105
overview 105
scan errors
viewing 293
SharePoint servers
configuration credentials 251

shares
    Add New Share/Edit Share dialog 240
    adding 239
    deleting 245
    editing configuration 245
    managing 241
site collections
    managing 263
supported file servers 17
Symantec Data Loss Prevention
    configuring 44
Symantec Data Loss Prevention users
    configuring authorization 64
system events
    viewing 288

T

troubleshooting
    Celera/VNX 391
    HNAS 399
    Isilon 394
    NetApp 389
    SharePoint 395

V

Veritas Data Insight
    administering 14
    administration tasks 15
viewing
    configured filers 208
VxFS file server
    configuration credentials 200

W

Windows File Server agent
    installing
        using Upload Manager utility 102
Windows File Servers
    configuration credentials 194