Backup Exec Private Cloud Services

Planning and Deployment Guide

Chapter 1

Introducing Backup Exec Private Cloud Services

This chapter includes the following topics:

About Backup Exec Private Cloud Services

Backup Exec Private Cloud Services is intended for managed service providers (MSP) who are interested in offering managed backup services to their customers. Backup Exec Private Cloud Services lets partners host backup storage within their datacenters as a "private cloud" configuration.

Managed service providers can provide backup services over the Internet to the partner private cloud as an alternative to managing offsite copies of tapes. Backups are encrypted and deduplicated, making transportation over a WAN secure and efficient. Local backups are still available on-premise for fast restore capability. Additionally, Backup Exec Private Cloud Services lets users perform backups directly to the cloud. Users can restore full or granular data directly from the cloud.

Backup Exec Private Cloud Services is also intended for Backup Exec customers with widely distributed networks. Customers can send duplicate copies of backups from remote offices to disk storage and tape storage within a central datacenter private cloud location.

The following table further explains some Backup Exec terms that are important to understanding Backup Exec Private Cloud Services.

Table 1-1 Backup Exec terms

Term	Definition
Deduplication disk storage	A deduplication disk storage device provides integrated deduplication on the Backup Exec server.

	Note: You can use Symantec NetBackup 5000/5020 series deduplication storage appliances instead of an integrated Backup Exec deduplication storage device in the cloud. An appliance may provide a more scalable option.
Optimized duplication	A type of duplication that enables deduplicated data to be copied directly from one OpenStorage device to another OpenStorage device from the same vendor.
Granular Recovery Technology (GRT)	A backup option that lets you restore individual items from database backups. A separate backup of the individual items is not required for you to recover one item.

See "Security considerations for Backup Exec Private Cloud Services"

See "System requirements for Backup Exec Private Cloud Services"

See "Configuring Backup Exec Private Cloud Services"

See "About the Backup Exec Private Cloud Services configurations"

Security considerations for Backup Exec Private Cloud Services

Backup Exec Private Cloud Services uses Backup Exec's current job and resource credential model to provide a secure experience. Additionally, Symantec recommends that you use a secure network connection between the customer location and the datacenter using a VPN solution. Various IPsec, SSL layer, and other VPN solutions are available.

You should use VLAN or routing restrictions to keep customer networks isolated from each other when you use any configuration that supports multiple customers.

See "About Backup Exec Private Cloud Services"

System requirements for Backup Exec Private Cloud Services

The following table lists the minimum system requirements and recommendations for running Backup Exec Private Cloud Services:

Table 1-2 System requirements for Backup Exec Private Cloud Services

Requirement	Description
Backup Exec servers	You can configure Backup Exec Private Cloud Services in one of three different ways. See "About the Backup Exec Private Cloud

	Services configurations" Any Backup Exec server in the cloud must include the Backup Exec Deduplication Option. The only requirement for local servers is that they comply with the requirements for Backup Exec 2012. You can find a list of compatible operating systems, platforms, and applications at the following URL: http://entsupport.symantec.com/umi/V-269-1
Deduplication Option license	You must install the Symantec Backup Exec Deduplication Option on both the private cloud server and any local Backup Exec servers. You do not have to create a deduplication disk storage device on the local Backup Exec server. However you must install the Deduplication Option on the local Backup Exec server to be able to access the shared deduplication disk storage device on the server in the cloud. All configurations require a deduplication disk storage device on the cloud Backup Exec server.
Central Admin Server Option license	You must install the Symantec Backup Exec Enterprise Server Option with Central Admin Server Option on the local or the cloud computers if you use the offsite copy configurations.
An active Internet connection	You must have an active Internet connection to transfer data to your private cloud deduplication disk storage device.
Virtual private network (VPN)	Symantec recommends that you use a secure network connection between the customer location and the datacenter using a VPN solution. Various IPsec and SSL layer VPN solutions are available.

Chapter 2

Configuring Backup Exec Private Cloud Services

This chapter includes the following topics:

Configuring Backup Exec Private Cloud Services

To configure Backup Exec Private Cloud Services, you must complete the following steps.

Table 2-1 How to configure Backup Exec Private Cloud Services

Step	Description
Step 1	You must configure the VPN between the private cloud Backup Exec server instance and any computers running on the local network.
Step 2	Consider which of the Backup Exec Private Cloud Services configurations best suits your needs and select one. You can choose to use a dedicated offsite copy to cloud configuration or direct backup configuration for each customer. See "About the Backup Exec Private Cloud Services configurations" You must configure Backup Exec Private Cloud Services. See "Setting up the offsite copy to cloud configurations"

	See "Setting up the direct backup configuration"
Step 3	After you configure the VPN and Backup Exec, you can begin working with Backup Exec Private Cloud Services. See "About working with Backup Exec Private Cloud Services for the offsite copy configurations" See "About working with Backup Exec Private Cloud Services and the direct backup configuration"
Step 4	If you use a VPN gateway with port restrictions, you may need to open port exceptions on both the on-premise and cloud VPN gateways. Port exceptions allow the Backup Exec Backup Exec server that is located in the cloud to communicate with the on-premise Backup Exec servers and agents. You should also change the CAS Backup Exec SQL port from a dynamically assigned port to a static port. The following Backup Exec support articles list all the port numbers that Backup Exec requires and which ones must be opened: http://www.symantec.com/business/support/index?page=content&id=HOWTO22990#id-SF700155293 http://www.symantec.com/business/support/index?page=content&id=HOWTO22989 http://www.symantec.com/business/support/index?page=content&id=HOWTO23022 The following Backup Exec support article details how to configure the SQL static port: http://www.symantec.com/business/support/index?page=content&id=HOWTO22985

About the Backup Exec Private Cloud Services configurations

You can configure Backup Exec Private Cloud Services in one of three ways.

Table 2-2 Specific configurations for Backup Exec Private Cloud Services

Configuration type	Details
Offsite copy to cloud managed Backup Exec server	The offsite copy to cloud managed Backup Exec server configuration uses a managed Backup Exec

	server, central administration server, and domain controller. The configuration provides offsite copy capabilities to a managed Backup Exec server that is located in the private cloud. This configuration requires one managed Backup Exec server per customer. See "About the offsite copy to cloud managed Backup Exec server configuration"
Offsite copy to cloud central administration server	The offsite copy to cloud central administration server configuration is similar to the first, except the locations of the central administration server and managed Backup Exec server are reversed. The configuration provides offsite copy capabilities to a central administration server that is located in the private cloud. This configuration requires one central administration server per customer. See "About the offsite copy to cloud central administration server configuration"
Direct backup	The direct backup configuration uses the Backup Exec Agent for Windows or the Backup Exec Agent for Linux instead of the managed Backup Exec server or central administration server. The configuration provides direct backup capabilities using a Backup Exec server that is located in the private cloud. This configuration requires one Backup Exec server per customer. See "About the direct backup configuration"

See "Setting up the offsite copy to cloud configurations"

See "Setting up the direct backup configuration"

About the offsite copy to cloud managed Backup Exec server configuration

The offsite copy to cloud managed Backup Exec server configuration involves three computers.

Table 2-4 Offsite copy to cloud managed Backup Exec server configuration

Computer	Role
Computer 1	The first computer (C1) is a Windows 64-bit server that has Backup Exec 2012 installed on it. C1 is configured as a managed Backup Exec server and it is located in the private cloud.
Computer 2	The second computer (C2) is a Windows 64-bit server that has Backup Exec 2012 installed on it. C2 is a central administration server that is located

	on the local area network. Note: You may use a 32-bit local Backup Exec server for C2 if you do not want to use a local deduplication disk storage device.
Computer 3	The third computer (C3) is a domain controller and DNS.

The network connection between the central administration server and the managed Backup Exec server is not always required to be active. The network connection is only necessary when you run any jobs that involve the managed Backup Exec server in the private cloud. The network connection does not need to be active for local jobs.

Figure 2-2 Offsite copy to cloud managed Backup Exec server

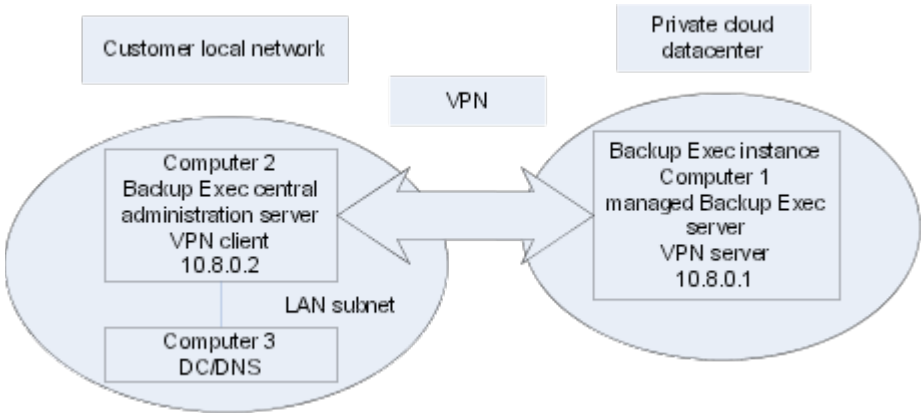

See "About the Backup Exec Private Cloud Services configurations"

About the offsite copy to cloud central administration server configuration

The offsite copy to cloud central administration server configuration involves three computers.

Table 2-5 Offsite copy to cloud central administration server configuration

Computer	Role
Computer 1	The first computer (C1) is a Windows 64-bit server that has Backup Exec 2012 installed on it. C1 is configured as a central administration server and it is located in the private cloud.
Computer 2	The second computer (C2) is a Windows 64-bit server that has Backup Exec 2012 installed on it. C2 is a managed Backup Exec server that is located on the local area network.

	Note: You can use a 32-bit local Backup Exec server for C2 if you do not want to use a local deduplication disk storage device.
Computer 3	The third computer (C3) is a domain controller and DNS.

This configuration lets you manage all your Backup Exec jobs within the private cloud's datacenter. It does, however, require that the network connection between the central administration server and the managed Backup Exec server is active at all times. The network connection must be active even when you run jobs locally.

Figure 2-3 Offsite copy to cloud central administration server

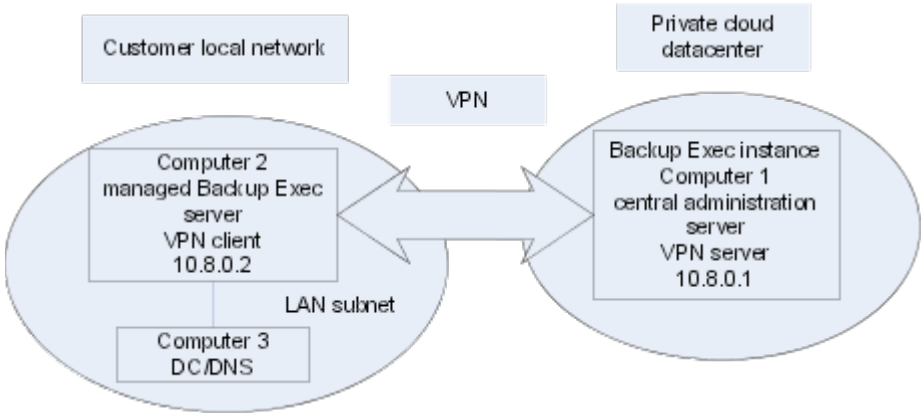

See "About the Backup Exec Private Cloud Services configurations"

About the direct backup configuration

The direct backup configuration involves a minimum of three computers.

Table 2-6 Direct backup configuration

Computer	Role
Computer 1	The first computer (C1) is the Windows 64-bit server Backup Exec 2012 server that is located in the private cloud datacenter.
Computer 2	The second computer (C2) is the Agent for Windows or Agent for Linux client that is located on the local area network. You can configure multiple agent client computers.
Computer 3	The third computer (C3) is a domain controller and DNS.

Figure 2-4 **Direct backup**

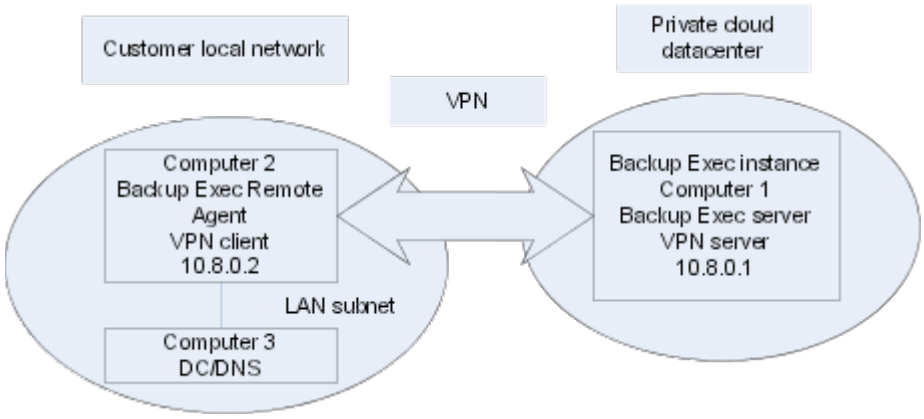

See "About the Backup Exec Private Cloud Services configurations"

Setting up the offsite copy to cloud configurations

After you have configured the VPN on the private cloud server, you should configure the Backup Exec server or servers.

See "Configuring Backup Exec Private Cloud Services"

You can select from one of two offsite copy to cloud configurations:

See "About the offsite copy to cloud managed Backup Exec server configuration"

See "About the offsite copy to cloud central administration server configuration"

Table 2-7 **How to configure the offsite copy to cloud configurations**

Step	Description
Step 1	Install the Backup Exec central administration server. See "Installing the Backup Exec central administration server"
Step 2	Install the managed Backup Exec server. See "Installing the managed Backup Exec server"
Step 3	Configure storage devices. See "Setting up storage devices for the offsite copy configurations"
Step 4	Seed the deduplication disk storage device with data. See "About seeding the deduplication disk storage device for the offsite copy configurations"

Installing the Backup Exec central administration server

You must install Backup Exec for Windows Servers on the computer that serves as the Backup Exec central administration server.

See "Setting up offsite copy to cloud configurations"

If you use the offsite copy to cloud managed Backup Exec server configuration, the central administration server is installed on a local office Backup Exec server (computer 2 or C2). Otherwise, the central administration server is installed as a cloud Backup Exec server (computer 1 or C1) for the offsite copy to cloud central administration server configuration.

You must add the central administration server to a domain. Install the Enterprise Server Option with Central Admin Server Option (CASO) on the central administration server.

Table 2-8 How to install the Backup Exec central administration server

Step	Description
Step 1	Add the Backup Exec server to your local domain by completing the following steps: • Using the Computer Properties dialog box in Windows, add the server to the domain. • Restart the computer when you are prompted to do so.
Step 2	After the server has restarted, log on with the domain account that you want to have administrator rights to your local Backup Exec instance.
Step 3	Use the proper license keys to install Backup Exec 2012. For more information on installing Backup Exec, see the *Symantec Backup Exec Administator's Guide*. Backup Exec partners can obtain licensing information from the Symantec PartnerNet Web site at the following link: https://partnernet.symantec.com/Partnercontent/Login.jsp
Step 4	Include the Enterprise Server Option with Central Admin Server Option (CASO) when you install Backup Exec. For more information on installing CASO, see the *Symantec Backup Exec Administrator's Guide*. Install the Deduplication Option when you use the offsite copy to cloud central administration server configurations. Using a local deduplication disk storage device on the central administration server is optional for the offsite copy to cloud managed Backup Exec server configuration.
Step 5	Use domain credentials for the default system logon account when you install Backup Exec.
Step 6	If you want to run incremental Exchange GRT duplicate backup jobs to the cloud, set the following registry value to **1** when the installation is complete. Changing the registry value disables the deduplication disk storage device's GRT-to-GRT duplicate copy capability on the Backup Exec server.

	dword HKEY LOCAL MACHINE\SOFTWARE\Symantec\Backup Exec for Windows\Backup Exec\Engine\Misc\DisablePDI2PDISetCopy This computer is now the central administration server that controls the managed Backup Exec server across the WAN. For more information about the limitations of offsite copy Granular Recovery Technology (GRT), refer to the following topic: See "Limitations of Granular Recovery Technology with offsite copy"

Installing the managed Backup Exec server

You must install the managed Backup Exec server. If you use the offsite copy to cloud managed Backup Exec server configuration, the managed Backup Exec server is installed as the cloud Backup Exec server (computer 1 - C1). Otherwise, the managed Backup Exec server is installed on a local office Backup Exec server (computer 2 - C2).

See "Setting up the offsite copy to cloud configurations"

To install the managed Backup Exec server

1. Do the following:

   Add the Backup Exec server to your local domain by completing the following steps:

   - Use the Computer Properties dialog box in Windows to add the server to the domain.
   - Restart the computer when you are prompted to do so.

2. After the server has restarted, log on with the domain account that has administrator rights to your local Backup Exec server.
3. Install Backup Exec 2012 on the server and select the **Managed Backup Exec server** installation option.
4. At the prompt, specify the same system logon account credentials that you used to install the central administration server.
5. If you want to use the offsite copy to cloud managed Backup Exec server configuration, select **Deduplication Option**. Using a local deduplication disk storage device on the managed Backup Exec server is optional for the offsite copy to cloud central administration server configuration.
6. When Backup Exec prompts you for the central administration server, enter the information for your local Backup Exec central administration server.
7. Select the **Centrally managed Backup Exec server** option.

8. If you want to run incremental Exchange GRT duplicate backup jobs to the cloud, set the following registry value to **1** when the installation is complete. dword HKEY LOCAL MACHINE\SOFTWARE\Symantec\Backup Exec for Windows\Backup Exec\Engine\Misc\DisablePDI2PDISetCopy Changing the registry value disables the deduplication disk storage device's GRT-to-GRT duplicate copy capability on the Backup Exec server.
9. Open Backup Exec on the central administration server.
10. Select the Storage tab, and then double-click the Backup Exec server that is located in the private cloud datacenter.
11. In the left pane, click Settings.
12. In the Private cloud server field, select **Enabled**.

Setting up storage devices for the offsite copy configurations

Before you can run backup jobs to the private cloud, you must configure storage devices.

See "Setting up the offsite copy to cloud configurations"

Table 2-9 How to set up storage devices for the offsite copy configurations

Step	Description
Step 1	Create new local disk storage devices on the local computer 2 (C2). You can create a deduplication disk storage device, if desired. For more information on creating storage devices, see the *Symantec Backup Exec Administrator's Guide*.
Step 2	Create a new deduplication disk storage device on your private cloud Backup Exec instance. For more information on creating a deduplication disk storage device, see the *Symantec Backup Exec Administrator's Guide*. Symantec recommends that you use a dedicated volume for the deduplication disk storage device if possible. Give the deduplication disk storage device a unique name to make it easy to differentiate from the local deduplication disk storage device, if you created one.
Step 3	If you want the at-rest data to be encrypted on your private cloud deduplication disk storage device, select **Yes, encrypt data during transmission to this deduplication disk storage device and while the data is stored on it** when you configure a new deduplication disk storage device. For an existing deduplication device, you can modify the Encryption field in the deduplication device's properties.

	Note: The VPN encrypts the data in transit between the local Backup Exec server and the cloud Backup Exec server.
Step 4	Share the new cloud deduplication disk storage device with your local Backup Exec computer. For more information on sharing deduplication disk storage devices, see the *Symantec Backup Exec Administrator's Guide*.
Step 5	Use the Backup Exec Services Manager to stop and restart all Backup Exec services on the local Backup Exec server. The process of sharing your cloud deduplication disk storage device with your local Backup Exec server is now complete. The private cloud deduplication disk storage device should appear and be accessible from both C1 and C2 now.

About seeding the deduplication disk storage device for the offsite copy configurations

To avoid long transfer times over the Internet, you can seed your deduplication disk storage device in the cloud with the data you need to get started. Seeding your deduplication disk storage device is the process of placing any initial configuration files or backup sets in the deduplication disk storage device to prepare it for use. Transfer times depend on the amount of data to be copied and backed up to the private cloud Backup Exec instance.

You can seed the initial data using one of two methods, depending on the type of data:

- You can seed the deduplication disk storage device with System State operating system backups. Seed the deduplication disk storage device by running duplicate backup jobs of System State data of other computers running in the private cloud. Back up System State data for the computers that run the same operating system as the local computers that you want to back up. See "Seeding operating system files for the offsite copy configurations"
- You can send a physical transfer drive that contains backup sets with the relevant data from the local Backup Exec server to the private cloud datacenter. See "About using a transfer drive to seed the deduplication disk storage device for the offsite copy configurations"

Seeding operating system files for the offsite copy configurations

To avoid long transfer times over the Internet, you can seed your deduplication disk storage device in the cloud with the data you need to get started. One way to seed the deduplication disk storage device is to use the System State backup data from other co-located computers.

See "About seeding the deduplication disk storage device for the offsite copy configurations"

Table 2-10 **How to seed operating system files for the offsite copy configurations**

Step	Description
Step 1	Install the Agent for Windows or the Agent for Linux on any computers that are co-located in the private cloud. For more information on installing Backup Exec agents, see the *Symantec Backup Exec Administrator's Guide*. The computers should be running the same operating system versions as the servers that are to be backed up on the local customer networks.
Step 2	Create and run backup jobs on the private cloud Backup Exec server. Back up the System State and system volumes of these co-located computers to the private cloud deduplication disk storage device.

About using a transfer drive to seed the deduplication disk storage device for the offsite copy configurations

To avoid long transfer times over the Internet, you can seed your deduplication disk storage device in the cloud with the data you need to get started. One way to seed the deduplication disk storage device is to use a physical transfer drive.

See "About seeding the deduplication disk storage device for the offsite copy configurations"

Symantec provides a calculator tool that lets you compare the time that is needed to use a transfer drive with the time that is needed to copy data over the Internet. You can find the calculator at the following link:

http://entsupport.symantec.com/umi/V-269-34

To seed your private cloud Backup Exec instance using a transfer drive, complete the following procedure:

See "Seeding the deduplication disk storage device using a transfer drive for the offsite copy configurations"

Seeding the deduplication disk storage device using a transfer drive for the offsite copy configurations

You can use a physical transfer drive to seed your private cloud Backup Exec deduplication disk storage device. Seeding your deduplication disk storage device with the files it takes to get started can save you the time of performing a large backup over the Internet.

See "About using a transfer drive to seed the deduplication disk storage device for the offsite copy configurations"

To seed your deduplication disk storage device using a transfer drive for the offsite copy configurations

1. Create disk storage on a portable drive on the local Backup Exec server, which is computer 2 (C2).

2. Copy a backup set to the disk storage and encrypt the data with software encryption using one of the following methods:

If you did not create the "DisablePDI2PDISetCopy" registery key during installation, then you can duplicate backup sets	Complete the following steps: • Select to duplicate the latest full backup sets of the data that you want to use to seed your private cloud deduplication disk storage device. • Select the disk storage that you created as the storage destination on the Duplicate Job dialog. • Configure software encryption on the Duplicate Job dialog. You must create or select an encryption key for software encryption.
If you created the "DisablePDI2PDISetCopy" registry key during installation, then you should create a full backup job	Complete the following steps: • Create a full backup job that uses the disk storage for any applications that are capable of Symantec's Granular Recovery Technology (GRT). • Disable GRT for any specific GRT-capable applications that you want to back up. Refer to the following topic for more information about the limitations of offsite copy to GRT. See "Limitations of Granular Recovery Technology with offsite copy" • Enable software encryption on the Storage panel. You must create or select an encryption key for software encryption.

3. Run the job you created in the previous step.
4. Ship the portable disk to the private cloud datacenter.
5. Attach the portable disk to the private cloud Backup Exec server.
6. Create disk storage on the attached portable drive using the disk storage that you originally created on the drive.
7. Create and run a Backup Exec inventory operation on the portable disk storage device.
8. Create and run a Backup Exec catalog operation on the portable disk storage device.
9. Duplicate the backup sets on the disk storage device and use the cloud deduplication disk storage device as the destination storage device.
10. When the duplicate operation is complete, you can use Backup Exec to retire and delete the files in the disk storage. Use a disk utility to wipe the portable drive clean. When you have successfully seeded your private cloud deduplication disk storage device, you have completed the configuration process. You can proceed to the following topic to begin working in Backup Exec: See "About working with Backup Exec Private Cloud Services for the offsite copy configurations"

Setting up the direct backup configuration

You should configure the Backup Exec server or servers.

See "Configuring Backup Exec Private Cloud Services"

The direct backup configuration involves a minimum of three computers.

See "About the direct backup configuration"

Table 2-11 How to configure the direct backup configuration

Step	Description
Step 1	Configure the private cloud deduplication disk storage device. See "Configuring the private cloud deduplication disk storage device for the direct backup configuration"
Step 2	Seed the private cloud deduplication disk storage device with data. See "About seeding the deduplication disk storage device for the direct backup configuration"

Configuring the private cloud deduplication disk storage device for the direct backup configuration

You must create the Backup Exec disk storage device and the deduplication disk storage device on the private cloud instance.

See "Setting up the direct backup configuration"

Table 2-12 How to configure the private cloud Backup Exec instance deduplication disk storage device

Step	Description
Step 1	Log onto C1 using the domain account that has administrator rights to your local server.
Step 2	Install Backup Exec 2012 on C1 and specify a system logon.
Step 3	On C1, in Backup Exec, create a new deduplication disk storage device. If you want the at-rest data to be encrypted on your private cloud deduplication disk storage device, select **Yes, encrypt data during transmission to this deduplication disk storage device and while the data is stored on it** when you configure a new deduplication disk storage device. For an existing deduplication device, you can modify the Encryption field in the deduplication device's

	properties. Note: The VPN encrypts the data in transit between the local Backup Exec server and the cloud Backup Exec server. For more information on creating a deduplication disk storage device, see the *Symantec Backup Exec Administrator's Guide*.
Step 4	Enable the private cloud server setting: • Open Backup Exec on the Backup Exec server. • Click the Backup Exec button, select Configuration and Settings, and then click **Local server properties**. • In the left pane, click **Settings**. • In the Private cloud server field, select **Enabled**.

About seeding the deduplication disk storage device for the direct backup configuration

To avoid long transfer times over the Internet, you can seed your deduplication disk storage device in the cloud with the data that you need to get started. Seeding your deduplication disk storage device is the process of placing any initial configuration files or backup sets in the deduplication disk storage device to prepare it for use. Transfer times depend on the amount of data to be copied and backed up to the private cloud Backup Exec instance.

You can seed the initial data using one of two methods, depending on the type of data that you want to seed:

- You can seed the deduplication disk storage device with System State operating system backups. Seed the deduplication disk storage device by running backup jobs of System State data of other computers running in the private cloud. Back up System State data for the computers that run the same operating system as the local computers that you want to back up. See "Seeding operating system files for the direct backup configuration"
- You can send a physical transfer drive that contains backup sets with the relevant data from the local Backup Exec server to the private cloud datacenter. See "Seeding the deduplication disk storage device using a transfer drive for the direct backup configuration"

Seeding operating system files for the direct backup configuration

To avoid long transfer times over the Internet, you can seed your deduplication disk storage device in the cloud with the data that you need to get started. One way to seed the deduplication disk storage device is to use the System State backup data from other co-located computers.

See "About seeding the deduplication disk storage device for the direct backup configuration"

Table 2-13 **How to seed operating system files for the direct backup configuration**

Step	Description
Step 1	Install the Agent for Windows and the Agent for Linux on any computers that you intend to back up on the local customer networks. For more information on installing Backup Exec agents, see the *Symantec Backup Exec Administrator's Guide*. The computers that you use to seed the data should run the same operating system versions as the computers that are to be backed up.
Step 2	Create and run backup jobs on the private cloud Backup Exec server. Back up the System State and system volumes of these co-located computers to the private cloud deduplication disk storage device.

Seeding the deduplication disk storage device using a transfer drive for the direct backup configuration

You can use a physical transfer drive to seed your private cloud Backup Exec deduplication disk storage device. Seeding your deduplication disk storage device with the files it takes to get started can save you the time of performing a large backup over the Internet.

See "About seeding the deduplication disk storage device for the direct backup configuration"

Table 2-14 **How to seed the deduplication disk storage device using a transfer drive for the direct backup configuration**

Step	Description
Step 1	Attach a portable drive to computer (C2).
Step 2	Copy the seed files from C2 to the portable drive.
Step 3	Encrypt the files on the disk using any 3rd party encryption tool.
Step 4	Ship the transfer drive to the private cloud datacenter.
Step 5	Connect the transfer drive to computer 1 (C1).
Step 6	Temporarily unencrypt the data on the transfer drive by using the same tool that was used to encrypt the data.
Step 7	Create and run a backup job that backs up the unencrypted files. Use the deduplication disk storage device in the cloud as the destination.
Step 8	When the backup job is complete, you can delete the copied source files. Use a disk utility to wipe the

	portable drive clean.

When you have successfully seeded your private cloud deduplication disk storage device, you have completed the configuration process.

You can proceed to the following topic to begin working with Backup Exec.

See "About working with Backup Exec Private Cloud Services and the direct backup configuration"

Chapter 3

Working with Backup Exec Private Cloud Services

This chapter includes the following topics:

About working with Backup Exec Private Cloud Services for the offsite copy configurations

Backup Exec Private Cloud Services lets you manage backup definitions using the Central Admin Server Option (CASO) and the Deduplication Option.

Symantec provides a helpful calculator tool that lets you estimate the time that is involved in copying data over the Internet. The cloud backup time calculator can be useful for planning your cloud backup strategy. You can use the calculator to determine if your system resources are adequate for backing up the customers' data within an allotted backup window. The time estimates can help you decide how much data you can reasonably support and how much time you should dedicate to cloud backups.

You can find the calculator at the following link:

http://entsupport.symantec.com/umi/V-269-34

See "Creating backup definitions for the offsite copy configurations"

See "About restoring data from the private cloud using the offsite copy configurations"

See "Restoring data from the private cloud using the offsite copy configurations"

See "Restoring data from the private cloud with a transfer drive using the offsite copy configurations"

Creating backup definitions for the offsite copy configurations

You can copy back up data to your private cloud Backup Exec instance by creating a backup definition with a duplicate stage. The backup definition resides on the central administration server. The definition contains backup jobs that back up data to the local deduplication disk storage device. The definition also contains a duplicate stage that then copies those backup sets to the private cloud deduplication disk storage device.

Optionally, you can add an additional duplicate stage to the backup definition to replicate the copied backup set from the cloud deduplication storage device. You can duplicate the backup set to a tape device that is also located in the cloud or to another deduplication storage device on a managed Backup Exec server. The managed Backup Exec server can be located in the private cloud or in a different physical location.

Note: For more information on creating backup definitions, see the *Symantec Backup Exec Administrator's Guide*.

To create backup definitions for the offsite copy configurations

1. On the central administration server, open Backup Exec.
2.

On the

   - To back up a single server, right-click the server name.
   - To back up multiple servers, Shift + click or Ctrl + click the server names, and then right-click one of the selected servers.

3. On the Backup menu, select the backup option that you want to use.
4. In the Name field, type a unique name for the backup definition.

Note: If you back up data from multiple servers, Backup Exec appends the server name to the text that you enter in the Name field. Backup Exec uses the server name and the text that you entered to create unique names for each backup definition.

5. Do any of the following:

To test or edit the credentials that Backup Exec uses to access backup selections	In the Selections box, click **Test/Edit Credentials**.
To change the backup selections	In the Selections box, click **Edit**.
To add a stage to the backup definition	Complete the following steps: • In the Backup box, click **Add Stage**. • Click **Duplicate** to add the duplicate stage. • In the Duplicate box, click **Edit**. • On the Storage pane, select the private cloud deduplication disk storage device as the storage for the duplicate operation.

- Complete any other settings as necessary. Symantec recommends that you verify the duplicate operation as a separate job. If you select to verify the operation at the end of the job, the job performance is degraded. You can configure the verify operation on the Verify pane.

Note:	You can add additional duplicate stages to the backup definition. You may want to send additional copies to a co-located tape device or to a deduplication storage device on a remote managed Backup Exec server, for example.

To modify the job settings

Complete the following steps:

- In the Backup box, click **Edit**.
- On the Storage pane, select the local deduplication disk storage device as the storage for the backup job.
- Complete any other settings as necessary.

6. When you are finished configuring the backup definition, click **OK** on the Backup Properties dialog box.

See "About working with Backup Exec Private Cloud Services for the offsite copy configurations"

About restoring data from the private cloud using the offsite copy configurations

After you have backed up data to the private cloud Backup Exec instance, you can restore it at any time. Restoring data from a private cloud Backup Exec deduplication disk storage device is very similar to restoring data normally in Backup Exec.

See "Restoring data from the private cloud using the offsite copy configurations"

It may be more efficient to restore a large amount of data from a Backup Exec private cloud instance using a physical transfer drive. You can use the transfer drive to transfer the data to the local Backup Exec server. Then use the local Backup Exec server to run the restore job.

See "Restoring data from the private cloud with a transfer drive using the offsite copy configurations"

Restoring data from the private cloud using the offsite copy configurations

You can restore data from the private cloud Backup Exec instance to the local Backup Exec client computers.

See "About restoring data from the private cloud using the offsite copy configurations"

To restore data from the private cloud using the offsite copy configurations

1. Ensure that the server that you restore contains the network route command that allows it to communicate with computer 1 (C1) as described in the following procedure: See "Configuring local network routing"
2. Open Backup Exec on the central administration server.
3. On the Backup and Restore tab, click Restore.
4. Select the data that you want to restore and any other necessary job options, and then submit the job.

See "About working with Backup Exec Private Cloud Services for the offsite copy configurations"

Restoring data from the private cloud with a transfer drive using the offsite copy configurations

You can copy data from the private cloud Backup Exec instance to the local Backup Exec server using a transfer drive. Using a transfer drive can be useful if you want to restore a large amount of data at one time. A large restore job can affect your system resources, depending on the amount of available bandwidth and time to complete the job.

See "About restoring data from the private cloud using the offsite copy configurations"

To restore data from the private cloud using a transfer drive and the offsite copy configurations

1. Create disk storage on a portable drive on computer 1 (C1), the private cloud Backup Exec instance.
2. Duplicate the backup sets that you want to restore from the cloud-based deduplication disk storage device. Select the disk storage that you created as the destination storage device. Make sure that you select to encrypt the data using software encryption. You must create or select an encryption key for software encryption. For more information on encrypting data, see the *Symantec Backup Exec Administrator's Guide*.
3. After the job is complete, ship the transfer drive to the local office.
4. After the portable drive arrives, connect the drive to the local Backup Exec server.
5. Create disk storage on computer 2 (C2) using the portable drive as the path.
6. Create and run Backup Exec inventory and catalog operations on the disk storage.
7. Restore the data from the new disk storage to the appropriate destination.
8. Erase the data from the transfer drive.

See "About working with Backup Exec Private Cloud Services for the offsite copy configurations"

Restoring data from a managed Backup Exec server in the event of a central administration server failure

If a hardware failure or other disaster affects your central administration server, it makes it impossible for your managed Backup Exec server to run backup or restore jobs. You can recover the central administration server by configuring a replacement computer and reinstalling the Backup Exec central administration server. You can also, however, convert a managed Backup Exec server to a standalone Backup Exec server to restore the central administration server.

To convert a managed Backup Exec server to a standalone Backup Exec server to restore the central administration server

1. The following steps must be followed before a CAS disaster occurs and in preparation for such an event:

   a, On the managed Backup Exec server, note the names and directory paths of any local disk storage.

Note: Double-click the disk storage on the Storage tab. Then click **Properties** in the left pane to view the storage properties.

   b. If the managed Backup Exec server has its own deduplication disk storage device, note the device's name, path, logon account, and password properties.

Note: Double-click the deduplication disk storage device on the Storage tab. Then click **Properties** in the left pane to view the storage properties.

2. Open the Programs and Features (or Add or Remove Programs) dialog or the Uninstall a program dialog from the Windows Control Panel.
3. Select the **Change** option for Symantec Backup Exec.
4. In the left pane, select **Additional Options**, if it is not already selected.
5. Click **Next** until you reach the Configure Managed Backup Exec Server panel.
6. Select the **Locally Managed Backup Exec server** option.
7. Click **Next**.
8. Do one of the following when you receive the message "Unable to contact {central administration server}. Ensure that the central administration server is running." When the installation is complete, the computer is no longer a centrally managed Backup Exec server.

If the central administration server is unavailable and you want this managed Backup Exec server to be managed locally	Click **OK** to continue.
If you want to retry this operation when the central administration server is running	Click **Cancel** to end the procedure.

9. Click **Next**.
10. Restart the computer if you are prompted to do so.
11. Open Backup Exec and select the Storage tab. If Backup Exec fails to connect to the Backup Exec server, restart the Backup Exec services and then try again.
12. Recreate any local disk storage by importing the original disk storage using the same names and paths that you noted in step 1.
13. Recreate any deduplication disk storage devices by importing the original deduplication disk storage devices using the same information that you noted in step 2.

Note: It might take much longer to recreate an existing storage device than it would to create a new storage device. The amount of time depends on how many backup sets the storage device contained and whether this managed Backup Exec server has access to its

domain controller and DNS.

14. Create and run Backup Exec inventory and catalog operations on each storage device that you recreated. You can now use the standalone Backup Exec server to restore any backup sets that were stored on the Backup Exec server's storage devices.
15. If you use the standalone Backup Exec server to recover the central administration server, you may need to delete the existing central administration server resource on the standalone Backup Exec server. Then push install the Agent for Windows to the central administration server before you restore it. Once the central administration server has been recovered, you may convert the locally managed Backup Exec server back into a centrally managed Backup Exec server using the Backup Exec change installation dialog again. Select the centrally managed Backup Exec server option to reconfigure the computer as a managed Backup Exec server.

About working with Backup Exec Private Cloud Services and the direct backup configuration

Backup Exec Private Cloud Services lets you manage backup definitions with client-side deduplication for the direct backup configuration.

Symantec provides a helpful calculator tool that lets you estimate the time that is involved in copying data over the Internet. The cloud backup time calculator can be useful for planning your cloud backup strategy. You can use the calculator to determine if your system resources are adequate for backing up the customers' data within an allotted backup window. The time estimates can help you decide how much data you can reasonably support and how much time you should dedicate to cloud backups.

You can find the calculator at the following link:

http://entsupport.symantec.com/umi/V-269-34

See "Enabling client-side deduplication for the direct backup configuration"

See "Creating backup definitions for the direct backup configuration"

See "Restoring data from the private cloud with a transfer drive using the direct backup configuration"

Enabling client-side deduplication for the direct backup configuration

Before you can create and run direct backup jobs to the private cloud Backup Exec instance, you must enable client-side deduplication.

To enable client-side deduplication for the direct backup configuration

1. On the Storage tab, double-click the storage for which you want to edit properties.
2. In the left pane, click **Properties**.
3. In the Client-side deduplication field, select **Enabled**.
4. Click **Apply**.

5. Restart the Backup Exec services.

Note: You must stop and restart Backup Exec services on C1.

After you have enabled client-side deduplication, you can create and run direct backup jobs.

For more information on creating backup jobs that use client-side deduplication, see the *Symantec Backup Exec Administrator's Guide*.

See "Creating backup definitions for the direct backup configuration"

See "About working with Backup Exec Private Cloud Services and the direct backup configuration"

Creating backup definitions for the direct backup configuration

After you have configured the VPN and enabled any additional computers for remote agent sharing and client-side deduplication, you can create and run direct backup jobs.

See "Enabling client-side deduplication for the direct backup configuration"

Note: For more information on creating backup definitions, see the *Symantec Backup Exec Administrator's Guide*.

Use the following procedure to back up data directly to the private cloud Backup Exec instance.

To create backup jobs for the direct backup configuration

1. On computer 1 (C1), open Backup Exec.
2.

On the

- To back up a single server, right-click the server name.
- To back up multiple servers, Shift + click or Ctrl + click the server names, and then right-click one of the selected servers.

3. On the Backup menu, select the backup option that you want to use.
4. In the Name field, type a unique name for the backup definition.

Note: If you back up data from multiple servers, Backup Exec appends the server name to the text that you enter in the Name field. Backup Exec uses the server name and the text that you entered to create unique names for each backup definition.

5. Do any of the following:

To test or edit the credentials that Backup Exec uses to access backup selections	In the Selections box, click **Test/Edit Credentials**.
To change the backup selections	In the Selections box, click **Edit**.

To add a stage to the backup definition	In the Backup box, click **Add Stage**.
To modify the job settings	Complete the following steps: • In the Backup box, click **Edit**. • Ensure that the option **Enable the remote computer to directly access the storage device and to perform client-side deduplication, if it is supported** is selected. • Complete any other settings as necessary.

6. When you are finished configuring the backup definition, click **OK** on the Backup Properties dialog box.

See "About working with Backup Exec Private Cloud Services and the direct backup configuration"

Restoring data from the private cloud with a transfer drive using the direct backup configuration

You can create a normal restore job to restore data from the private cloud Backup Exec instance to the local client. However, if you want to restore a large amount of data at one time, it may make sense to use a physical transfer drive. The time it takes to transfer a large amount of data depends on the amount of available bandwidth and the time to complete the job.

To restore data from the private cloud with a transfer drive using the direct backup configuration

1. Create and run a restore job on computer 1 (C1) to restore the files to a folder on a portable disk drive.
2. After the job has completed, encrypt the files on the disk using any 3rd party encryption tool.
3. Ship the portable drive to the local office.
4. When the portable drive arrives, unencrypt the files using the same tool that you used to encrypt them.
5. Transfer the unencrypted files to their proper destination on computer 2 (C2).
6. Erase or wipe the files from the transfer drive completely to ensure that the data is permanently removed.

See "About working with Backup Exec Private Cloud Services and the direct backup configuration"

About cloud disaster recovery service

The Backup Exec 2012 Simplified Disaster Recovery (SDR) feature and the conversion to virtual machine feature let service providers or customers provide cloud disaster recovery services. Backup data that is stored in the cloud can be used to create temporary replacement virtual or physical servers in the private cloud in the event of a disaster.

Specific network configurations and failure conditions can affect the specific steps that are required for failover and failback. This section provides only basic guidelines for using the SDR and conversion to virtual machine features within a Backup Exec private cloud environment to provide disaster recovery services.

There are two main disaster recovery scenarios that can occur. The first scenario is the server failover and failback in which one or more on-premise servers fail, but the on-site network remains intact. The second scenario is the site failover and failback in which an entire site has failed.

See "Recovering a server or a site from failover"

See "Recovering a server or a site from failback"

Recovering a server or a site from failover

To prepare for a server failover scenario, you should configure and run regularly scheduled Simplified Disaster Recovery (SDR) enabled backup definitions for any business critical servers. The backup definitions must include duplicate stages that copy the backup data to the private cloud deduplication disk storage device. When the server failover occurs, you use the private cloud Backup Exec server to recover replacement virtual or physical servers.

See "About cloud disaster recovery service"

To recover a replacement physical server, use the Simplified Disaster Recovery Disk to perform a bare metal restore. Use the most recent SDR-enabled backup on the private cloud deduplication disk storage device. You can transport the replacement server to the on-premise site to replace the failed server. A site failover requires that an entire group of business critical servers be replaced with virtual machines in a hypervisor environment that is located in the cloud.

For more information on Simplified Disaster Recovery, see the *Symantec Backup Exec Administrator's Guide*.

Note: Specific network configurations and failure conditions can affect the specific steps that are required for failback. The following procedure provides only basic guidelines for using a Backup Exec private cloud environment to provide disaster recovery services.

To recover a server or a site from failover

1. Create a Hyper-V or VMWare ESX hypervisor environment in the cloud location.
2. Create a fenced virtual network for the replacement virtual machine or virtual machines that will run on the hypervisor. The replacement servers should retain their original, on-premise IP address for an entire site failover scenario.

Note: When you recover a site, the replacement servers should retain their original, on-premise IP address. You should restore the replacement computers in a logical order. For example, you should restore any domain controller and DNS servers first.

3. Do one of the following:

To failover from a physical computer	Complete the following steps: • Create and run a conversion to a virtual machine. Convert the SDR point-in-time system volume and the System State data into virtual machines for all replacement computers. The virtual machines should be targeted to the hypervisor. Do not select any application resources at this time.

	• Configure any fixed IP addresses for the replacement virtual machines, if necessary. • Establish network connectivity between the replacement virtual machine or virtual machines and the private cloud Backup Exec server or servers. • Create and run restore jobs from the same SDR-enabled point-in-time backups for each of the replaced servers. Select all of the computer's resources that are available for that point in time. Redirect the restore data to the replacement server or servers.
To failover from a virtual machine	Create and run a redirected restore job from each of the replacement servers' most recent SDR point-in-time backups. The same type of hypervisor should be used for both the on-premise and the cloud servers.

4. To recover only a single server, establish VPN connectivity between the replacement virtual server and the on-premise network and configure any on-premise DNS entries for the replacement virtual machines' IP addresses.
5. Expose any new external addresses from the cloud network and change any external DNS records if the failed server or servers were exposed through external IP addresses (an Exchange mail server, for example).
6. Configure and run regularly scheduled hypervisor host backup definitions for the replacement virtual machine or virtual machines. Use the private cloud deduplication disk storage device as the backup destination. If the on-premise Backup Exec server or servers have local deduplication disk storage, the backup definitions must include a duplicate stage that copies the backups to the on-premise deduplication disk storage device.

Recovering a server or a site from failback

You can recover a server or a site in the event of failback. A site failback scenario requires that an entire group of business critical servers be restored to on-premise physical servers or virtual machines.

See "About cloud disaster recovery service"

You may want to recover on-premise servers gradually instead of recovering them all at one time. You can recover some servers initially and leave others to be recovered over a period of days or weeks. This strategy likely requires VPN connectivity and IP address changes for the remaining replacement cloud servers that connect to the on-premise network.

For more information on Simplified Disaster Recovery, see the *Symantec Backup Exec Administrator's Guide*.

Note: Specific network configurations and failure conditions can affect the specific steps that are required for failback. The following procedure provides only basic guidelines for using a Backup Exec private cloud environment to provide disaster recovery services.

To recover a server or a site from failback

1. Run a Simplified Disaster Recovery (SDR) enabled backup and include any duplicate stages.
2. Turn off the replacement virtual machine or virtual machines.
3.

If the SDR-enabled backup definition did not include a duplicate stage that sent backup sets to the on-premise deduplication disk storage, complete the following steps:

- Add a portable disk storage device to Backup Exec on the private cloud Backup Exec server or servers.
- Duplicate the backup sets from the final point-in-time backup of all of the replacement computer or computers' data. Use the portable disk storage device as the destination.
- Ship the portable disk storage device to the on-premise location.
- Add the portable disk storage device to Backup Exec on the on-premise Backup Exec server or servers.
- Inventory and catalog the disk storage device on the on-premise Backup Exec server or servers.

4. Do either of the following:

To fail back to an on-premise physical server or servers	Complete the following steps: • Use the Simplified Disaster Recovery Disk to perform a bare metal restore. Select the most recent SDR-enabled backups on the on-premise Backup Exec server or servers. • Configure a fixed IP address for the recovered computers, if necessary. • Configure any on-premise DNS entries for the recovered computer or recovered computers' IP addresses, if necessary.
To fail back to an on-premise virtual server or servers	Complete the following steps: • Create and run a redirected restore job from the replacement server or servers' most recent point-in-time backups. The same type of hypervisor should be used for both the on-premise and the cloud servers. • Configure a fixed IP address for the recovered virtual machines, if necessary. • Configure any on-premiseDNS entries for the recovered virtual machine or virtual machines' IP addresses, if necessary.

5. If the failed server or servers was exposed through an external IP address (an Exchange mail server, for example), restore the original address or addresses in the external DNS records.

6. Delete the replacement cloud server or servers' backup definitions.
7. Resume running the original backup definition or definitions for any restored on-premise computers.

Backup Exec deduplication disk storage device requirements

The Backup Exec deduplication disk storage device requirements apply to all private cloud configurations. If you reach the share limit on a particular cloud Backup Exec server, you must add additional cloud Backup Exec servers.

For more information on deduplication disk storage device requirements, see the *Symantec Backup Exec Administrator's Guide*.

Limitations of WAN latency

If your network has high levels of network latency, it can adversely affect the performance of your initial direct cloud backup job. Latency can also affect some duplicate backup jobs that transfer data between the local office and the private cloud Backup Exec server. You may experience performance issues even if you seeded the deduplication disk storage device with a transfer drive, although you always improve performance by seeding devices. During the initial backup job, Backup Exec identifies and caches information about data segments, which provides more efficient performance for subsequent jobs.

Note: High latency values can be considered any average round-trip latency of over 30 milliseconds. The higher the latency, the more Backup Exec's performance is affected.

This limitation does not apply to duplicate backup jobs, when both the source device and target device are deduplication disk storage devices.

The following are limitations to using Backup Exec Private Cloud Services in high latency environments:

- Duplicate backup jobs that use a source device other than a deduplication disk storage device and a private cloud deduplication disk storage device as the destination may experience performance issues. Avoid these performance issues by using a deduplication disk storage device as the local source storage device.
- You may find that using the direct back up to cloud configuration is not suitable for backing up large amounts of data.
- If you delete and recreate backup definitions for the same resources, Backup Exec must cache the data fingerprints all over again. So you may experience the same performance issues as with the initial direct cloud backup job.

Limitations of Granular Recovery Technology with offsite copy

The following are limitations to using Backup Exec's Granular Recovery Technology (GRT) option with the offsite copy configuration:

- Backing up local Exchange incremental GRT-enabled backup sets to a private cloud deduplication disk storage device creates backup data in an MTF tape format. You can restore granular data from these backup sets, but it requires staging the backup set on the cloud Backup Exec server during the restore job. This limitation does not exist for direct backup of GRT-enabled backup sets to the cloud deduplication disk storage device.

- Copying duplicate GRT-enabled sets from local tape devices directly to a cloud deduplication disk storage device is not recommended and can result in excessive job run times.
- Backing up GRT-enabled sets directly to the cloud Backup Exec server may cause reduced performance times in high latency environments. You may experience reduced performance even after the initial backup. If performance continues to be a problem, you may want to disable GRT for direct backups.

1.