

Backup Exec Best Practices

This section includes the following topics:

- [Best practices for Backup Exec software encryption](#)
- [Best practices for Backup Exec installation](#)
- [Best practices for Backup Exec tape management](#)
- [Best practices for Backup Exec disk-based storage](#)
- [Best practices for Backup Exec data lifecycle management \(DLM\)](#)
- [Best practices for Backup Exec catalogs](#)
- [Best practices for Backup Exec backups](#)
- [Best practices for Backup Exec backing up critical system components](#)
- [Best practices for Backup Exec Agent for Microsoft Exchange Server](#)
- [Best practices for Backup Exec Agent for Microsoft SQL Server](#)
- [Best practices for Backup Exec Agent for Linux](#)
- [Best practices for Backup Exec Agent for Microsoft SharePoint](#)
- [Best practices for Backup Exec Central Admin Server Option](#)
- [Best practices for Backup Exec Agent for Oracle on Windows and Linux Servers](#)
- [Best practices for Backup Exec NDMP Option](#)
- [Best practices for Backup Exec reports](#)
- [Best practices for Backup Exec and Veritas Update](#)
- [Best practices for Backup Exec Simplified Disaster Recovery](#)
- [Best practices for Backup Exec Agent for Enterprise Vault and the Backup Exec Migrator](#)
- [Best practices for Backup Exec Granular Recovery Technology](#)
- [Best practices for Backup Exec Remote Media Agent for Linux](#)
- [Best practices for Backup Exec Agent for Microsoft Hyper-V](#)
- [Best practices for Backup Exec Agent for VMware](#)
- [Best practices for using Backup Exec with server clusters](#)
- [Best practices for Backup Exec Deduplication Option](#)

- Best practices for using Backup Exec Deduplication Option with the Central Admin Server Option
- Best practices for using hot-pluggable devices such as USB devices in a drive rotation strategy
- Best practices for Backup Exec database encryption keys
- Best Practices for Using the Veritas Backup Exec Cloud Connector

Best practices for Backup Exec software encryption

Best practices include tips and recommendations to help you use software encryption effectively with Backup Exec. For more information about software encryption, see the *Backup Exec Administrator's Guide*.

The following best practices can ensure smooth operations when you administer software encryption:

- Create strong pass phrases by doing the following:
 - Use more than the minimum number of characters that are required.
 - Use a combination of upper- and lower-case letters, numbers, and special characters.
 - Avoid literary quotations in pass phrases.
 - Keep your pass phrases secure.
- Run Backup Exec Services in FIPS mode and use 256-bit AES encryption to be FIPS-compliant.

To enable FIPS compliance, select the Use FIPS 140-2 compliant software encryption option in the Network and Security settings.

- Use software compression with encryption if the data must be compressed.
- Avoid using hardware compression with software encryption.
- If you do not use software encryption when you back up data to disk-based storage, use File System encryption to prevent unauthorized access.
- Use the same encryption key for all of the templates in a synthetic backup policy.

Best practices for Backup Exec installation

Best practices include tips and recommendations to help you install or upgrade Veritas Backup Exec more effectively. For more information about installing Backup Exec, see the *Backup Exec Administrator's Guide*.

Best practices for preparing to install or upgrade Backup Exec

- Visit the Veritas Support Web site to check for updates to the documentation.
- Review the Readme document and Documentation Addendum for updates to the Backup Exec Administrator's Guide.
- Review the "Features or requirements no longer supported by Backup Exec" section of the readme.
- Use only standard ANSI characters for the computer name of the computer on which you want to install Backup Exec. You may receive errors if you install Backup Exec on a computer with a name that uses non-standard characters.
- Ensure that the most recent version of Microsoft.NET framework is installed on the computer on which you want to install Backup Exec. Installing the .NET framework expedites Backup Exec's installation process.
- Run Microsoft Windows update.
- Run Symantec LiveUpdate before you upgrade Backup Exec to install the latest feature pack or service pack for the version of Backup Exec that is currently installed.
- Restart the server to resolve any pending restarts that are required for system updates.
- Document your current configuration and settings before you upgrade Backup Exec. You can verify that your configuration after the upgrade is complete.
- Upgrade the central administration server first if you perform a rolling upgrade in an existing Central Admin Server Option (CASO) environment. Upgrade the managed Backup Exec servers as soon as possible afterwards.
- Back up your server before you install or upgrade any software, including Backup Exec. Copy Backup Exec's Catalogs and Data folders for additional security.
- Pause or stop all jobs before upgrading Backup Exec. If you run Backup Exec jobs during an upgrade, the jobs fail.
- Pause all communications with managed Backup Exec servers in CASO environments.
- Delete any unnecessary job history and alert history before an upgrade to help expedite the process.

- Disable any monitoring utilities that may restart services.
- Run the Environment Check Utility. The Environment Check Utility automatically runs during local installations. Veritas recommends that you also run the utility for remote installations. You should address all warnings and errors before you install Backup Exec.
- Perform database maintenance on your Backup Exec Database immediately before an upgrade.
- Close all instances of the Backup Exec Administration Console before the upgrade.

Best practices during the installation process and the upgrade process

- Use an uninterrupted power supply (UPS) for your Backup Exec server during the installation. A UPS helps ensure that you do not have a failed installation due to a power outage.
- Install Backup Exec to a drive that is neither compressed nor encrypted if you use SQL Express for the Backup Exec Database.
- Run the installation wizard from the local server, from a DVD image on the local server, or by push-installing from the local server.
- Ensure that the path to the Backup Exec installation media does not exceed 60 characters. If the Backup Exec installation media is on a network share or in a custom folder that exceeds 60 characters, the installation may fail. To resolve this issue, shorten the path name or move the installation media to the root of the drive on the local system, and then run the installation again.
- Ensure that you review and acknowledge the upgrade notice about the disk reclamation process called data lifecycle management (DLM) during the upgrade.
- Ensure that you review and acknowledge the Migration Report after the migration completes to continue the upgrade process.
- Wait until after the installation to make configuration changes. Do not make configuration changes during the installation.

Best practices after the installation process or the upgrade process

- Run Symantec LiveUpdate to check for any Backup Exec updates.
- Run Microsoft Windows Update. Backup Exec uses many Microsoft technologies that may have been updated since Backup Exec's release.
- Monitor your disk space regularly to prevent disk space problems. Backup Exec's space requirements vary depending on usage and installed options. The requirements in the

Administrator's Guide do not include space estimates for the Simplified Disaster Recovery files, catalogs, or job logs, for example.

- Consult any of the following resources on the Help and Documentation menu if you have questions or difficulties:
 - Use the Administrator's Guide for comprehensive information about Backup Exec.
 - Use the Backup Exec Help for searchable, topic-based documentation.

Best practices for Backup Exec tape management

Best practices include tips and recommendations to help you use Veritas Backup Exec to manage tapes effectively. You should also review the best practices for the agents, options, or features that you use for more information. If you use backup jobs for which Granular Recovery Technology (GRT) is enabled, there are additional best practices for media management.

See [Best practices for Backup Exec Granular Recovery Technology](#).

For more information about tape management, see the *Backup Exec Administrator's Guide*.

The following best practices are for effective tape management:

- Be aware of the following consequences of the infinite setting for the overwrite protection period for all tape media:
 - Backup data may consume tape capacity quickly.
 - Tapes do not become recyclable automatically. You must specify when to overwrite each tape by setting the append and overwrite protection periods in the media set that you associate the tape with.
- Create new media sets with the append and overwrite protection periods that accommodate your needs. When the overwrite protection periods expire, the tapes are recyclable and Backup Exec has access to overwritable tapes.
- You should overwrite tapes periodically to keep the media family at a manageable size so that Backup Exec can rebuild the catalog if necessary. You can use a tape rotation strategy so that tapes are periodically overwritten, or select the option *Overwrite media* when you run a full backup.

“ ”

Note: A media family is a set of tapes that are related because one or more backup jobs span from one tape to another. For example, you back up Resource A to Tape 1. When Tape 1 runs out of space, Backup Exec uses Tape 2 to complete the backup. Tape 1 and Tape 2 are now in the same media family because the backup job started on Tape 1 and finished on Tape 2. If you append more backups to Tape 2 until it runs out of space, then Backup Exec uses Tape 3. Tape 3 is now part of the same media family as Tape 1 and Tape 2. The media family continues to grow as Backup Exec appends backup jobs. A new media family starts when a backup job overwrites the first tape it uses.

“ ”

Best practices for Backup Exec disk-based storage

Best practices include tips and recommendations to help you use Veritas Backup Exec to manage disk-based storage effectively. You should also review the best practices for the agents, options, or features that you use.

For more information about disk-based storage, see the *Backup Exec Administrator's Guide*.

The following best practices are for effective disk-based storage management:

- If you want to keep backup data longer than the period that you specify when you create the backup job, you should duplicate the backup sets. A duplicate backup job can copy the backup data from the original storage device to tape or to disk cartridge, which you can then send for long-term or off-site storage.
- When you use the Configure Storage wizard to create disk storage, Backup Exec provides a list of disks on which you can create disk storage. The disks do not appear in alphabetical order. Instead, the disk that appears first in the list has the most amount of disk space. You can select any disk that you want, but the disk that Backup Exec recommends for use appears at the top of the list. The disk that you use as the system drive always appears last in the list. Veritas recommends that you do not configure disk storage on the system drive.
- Veritas recommends that you use a dedicated hard disk or iSCSI attached device as disk storage for backup data.
- You should not configure disk storage and deduplication disk storage on the same disk.
- Before you create the disk storage on a network share, you must give read and write permissions to the Backup Exec service account. The Backup Exec service account is on the Backup Exec server that you want to access the network share.

- Do not delete or edit the contents of the BEControl folder, which Backup Exec creates on the root of the volume. Do not copy the BEControl folder to other volumes or drive letters.
- Do not delete or edit the changer.cfg or folder.cfg files. These files store information about the backup files.

“ ”

Note: If Windows data deduplication is enabled on the disk storage volume, Backup Exec excludes the backup data in the folder \BEData from deduplication, unless the \BEData folder already exists. Backup Exec must exclude backup data from deduplication for you to use Simplified Disaster Recovery (SDR) to perform a local recovery of the Backup Exec server. If Windows data deduplication is enabled on the disk storage volume, local disaster recovery using SDR fails. The Windows Preinstallation Environment (Windows PE) that SDR uses cannot read the files that Windows data deduplication processes.

“ ”

Best practices for Backup Exec data lifecycle management (DLM)

Best practices include tips and recommendations to allow the most efficient use for the Veritas Backup Exec data lifecycle management (DLM) feature. You should also review the best practices for the agents, options, or other features that you use.

For more information about DLM, see the *Backup Exec Administrator's Guide*.

The following best practices are for efficient DLM:

- To prevent DLM from deleting an expired backup set, you can manually retain the backup set or you can change the expiration date of the backup set. Backup Exec automatically retains all dependent backup sets as well. When you no longer want to retain a backup set, you must release it so that DLM can delete it.
- To prevent the inadvertent loss of backup sets when you enable the option 'Allow Backup Exec to delete all expired backup sets', do the following:
 - Ensure that when you create jobs, the backup sets are kept longer than the amount of time between full backups. Otherwise, the backup sets from the last full backup job may expire before the next full backup runs.

- Ensure that you rerun failed or missed jobs before the backup sets from the previous full backup expire.
- When setting up the schedules for backups in a backup definition, avoid adding many incremental backups between full backups. The DLM process must search through each backup set to check dependencies; therefore, the more incrementals there are, the longer the DLM process takes.
- To monitor the backup sets that DLM deletes, you can view the Backup Set Retention category in the audit log. You can also run the audit log report to view the backup sets that DLM deletes.

“ ”

Warning: DLM deletes all expired backup sets that are created by a one-time backup job. DLM does not keep the last backup set after the retention date expires if the backup set is from a one-time backup.

“ ”

Best practices for Backup Exec catalogs

Best practices include tips and recommendations to help you use Veritas Backup Exec to manage catalogs effectively. You should also review the best practices for the agents, options, or features that you use for more information.

For more information about catalogs, see the *Backup Exec Administrator's Guide*.

The following best practices are for effective catalog management:

- Do not delete or edit the files in BE_INSTALL\Catalogs\directory. The catalog files store the metadata for the backup sets.
- Back up the catalog files in BE_INSTALL\Catalogs\directory. Consider backing up the catalog files after you back up or perform maintenance on the Backup Exec Database.
- If you want to change the catalog location, Veritas recommends that you use Backup Exec Utility. If you change the location using the Catalog path field in the Backup Exec settings, you must also manually copy the existing catalogs to the new location and then restart the Backup Exec services. Refer to the following knowledge base article for more information:

https://www.veritas.com/support/en_US/article.TECH210578

The following recommendations are to help you choose the best catalog location to use in a Central Admin Server Option (CASO) environment:

- The distributed catalog location is the default location, and is efficient and suitable for a CASO environment that uses a wide area network (WAN). An example of this environment is if the central administration server is located in a cloud and the managed Backup Exec servers are on a local network.
- The centralized catalog location is best suited for use in a CASO environment that uses SAN-connected shared storage with a stable, high-bandwidth network between the central administration server and the managed Backup Exec servers. Use a centralized catalog location if the managed Backup Exec servers have minimal CPUs and disk space, and if the central administration server has significantly more CPUs and memory.
- The replicated catalog location is best suited for a CASO environment that has a stable, high-bandwidth network between the central administration server and the managed Backup Exec servers. A replicated catalog location can provide better performance through catalog redundancy since queries for catalog information are performed locally.

Best practices for Backup Exec backups

Best practices include tips and recommendations to help you use Veritas Backup Exec backup jobs effectively. For more information, see the *Backup Exec Administrator's Guide*.

The following best practices help ensure effective backup jobs:

- Test to make sure that you have the appropriate credentials to access the content that you want to back up before you run a backup job. If a credentials test fails, you can enter new credentials for the content so that Backup Exec can access it.
- You should run a backup job to your storage device before you run a test run job. Backup Exec does not recognize the capacity of a storage device until an actual backup job sends data to the device. If you create a test run job before any other jobs, Backup Exec cannot check that the device has sufficient capacity to perform the backup job. After at least one backup job has send data to a device, Backup Exec can determine the capacity.
- You should always run a full backup job before and after upgrading Backup Exec, the operating system, or any applications.
- Be sure to run full backup jobs periodically, in addition to any incremental backup jobs that you run. When you restore data that was backed up using incremental backups, Backup Exec restores the data from the initial full backup plus any data that was backed up in subsequent incremental backups. However, if one of the incremental backups is corrupt or missing, it can

cause the restore to fail. Running full backups periodically can help ensure that all of the data is accessible when you need to restore it.

- You should avoid using hardware compression with software encryption. Hardware compression is performed after encryption. Data becomes randomized during the encryption process. Compression does not work effectively on data that is randomized.
- You should not use software compression or encryption for GRT-enabled backup jobs. The compression and encryption process are resource-intensive. Enabling either software compression or encryption can result in degraded performance for GRT-enabled backup jobs.
- You should run a verify operation after all backup jobs. Running a verify operation can help you to determine if you will be able to restore the backup sets created by a backup job. If a verify operation fails, you can rerun the backup job to ensure that your data is protected. Otherwise, you may not realize that the media is inaccessible until you try to restore from it. By default, Backup Exec automatically verifies backed up data at the end of a backup job. However, you can also schedule the verify operation to take place at a later time or manually verify backup sets at any time.
- If you need to keep data longer than four weeks, you should duplicate it. You can duplicate the backup data from the original storage device to tape, for example, which you can then send for long-term or off-site storage.

Best practices for Backup Exec backing up critical system components

Best practices include tips and recommendations to help you use Veritas Backup Exec to protect critical system components effectively.

For more information about backing up critical system components, see the *Backup Exec Administrator's Guide*.

To create backup sets that are capable of disaster recovery, you must select to back up each of your critical system components in full. In Backup Exec, all of your critical system components are selected by default when you create a backup.

A full backup of your critical system components is necessary for any of the following restore scenarios:

- Simplified Disaster Recovery
- Conversion to virtual machines
- Complete online restore of a Microsoft Windows computer

You can select critical system components explicitly, by manually selecting them in the backup selections pane. Or you can select them implicitly, by selecting the parent server node and letting Backup Exec dynamically include the critical component as a child. When you add new backup sources to an existing backup source, Backup Exec automatically selects the new sources and includes them in the backup.

You can only create backup sets that are capable of disaster recovery on computers that run Windows. Creating disaster-recovery capable backup sets is not supported for clustered virtual nodes or Oracle RAC.

The following system components are considered critical:

- System volume (including BIOS, EFI, UEFI, and utility partitions)
- Boot volume (executing operating system)
- Services application volumes (boot, system, and automatic startup)
- System State devices and volumes (including Active Directory, system files, etc.)
- Windows Recovery Partition (WinRE) on Windows 8/Server 2012/8.1/Server 2012 R2

You must make a full selection of the critical system components. Backing up only the System State does not ensure the complete recovery of a server. If any directory or file on the critical system component is excluded from the backup selections, the resulting backup creates backup sets that are not capable of disaster recovery.

When all of the critical system components are included in your backup job selections, the Simplified Disaster Recovery indicator on the selections pane reads ON. If you deselect one or more critical system components, the indicator changes to OFF.

The following best practices will help you to back up critical system components more effectively:

- You can deselect any non-critical volumes or applications to create a backup job that protects only the operating system and any critical system components. Run the job periodically and before and after any system upgrades. Then you can create and run separate backup jobs to back up data volumes and applications on a more frequent basis.
- You should isolate user data on disk storage devices that are designated as data volumes. User data can include application data or user share data, for example. Placing user data on separate volumes lets you include application data or user share data, for example. Placing user data on separate volumes lets you include or exclude backup selections from those volumes without affecting the critical system components. Then you can perform full backups on your critical system components to create disaster recovery compatible backup sets. This

configuration lets you perform less frequent critical system component backups and more frequent application data and user data backup jobs.

- Create only one critical system component backup for each server or backup source.
- Upgrade previous Backup Exec server and Agent for Windows installations before performing backups with a new version of Backup Exec.

If your Backup Exec disk storage is located on a critical system volume, you will automatically back up the .bkf files that contain any backed up data whenever you run a critical system component backup job. Backing up the .bkf files can result in large, redundant backups that require a great deal of storage space. As an additional best practice, you may want to enable a registry key to let Backup Exec automatically exclude any .bkf files when you back up a volume.

To automatically exclude .bkf files from backups, locate the following key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Symantec\Backup Exec For Windows\Backup  
Exec\Engine\Backup
```

Set the value `Disable Backup of Disk Storage Devices` to 1 to automatically skip .bkf files. Set the value to 0 to re-enable the default behavior in which .bkf files are backed up.

Best practices for Backup Exec Agent for Microsoft Exchange Server

Best practices include tips and recommendations to help you use the Exchange Agent effectively. For more information about the Exchange Agent, see the *Backup Exec Administrator's Guide*.

Best practices for preparing the Exchange Server for backup

- Circular logging must be disabled if you want to do the following:
 - Run incremental and differential backups.
 - Recover data up to the point of failure.
- Put transaction log files on a separate physical disk from the database. If the disk that contains the database is damaged, the transaction logs are available as a recovery resource.
- Set the retention period for deleted items and mailboxes to a length of time that is appropriate for the available disk space. The longer the retention period, the more disk space is required. However, some retention period can prevent you from having to restore a mailbox or database. If possible, configure the Exchange server so that items are not deleted until a full backup is performed.

- Make Write Cache unavailable on the SCSI controller. Data corruption can occur if the computer fails before the operation is written to disk.
- Monitor the Application, Security, and System logs for any relevant events that may affect Exchange Server functionality.
- Allow sufficient disk space for maintenance and recovery procedures. Refer to your Microsoft documentation for details.
- Document the Exchange server configuration in detail.
- Avoid making the Exchange server a domain controller. You can restore Exchange more easily if you don't have to restore the Active Directory first. Also, it may add rights to the Exchange Trusted Subsystem account.
- Install the Exchange Server into a domain that has at least two domain controllers. With two domain controllers in a domain, databases on a failed domain controller can be updated with replication.
- You must have local administrator rights on each node of a database availability group (DAG) and on the Microsoft Exchange mailbox server to back up and restore Microsoft Exchange database files.
- For Exchange 2010/2013, use a Database Availability Group (DAG) with at least one passive database copy for each database to protect against data loss. If you can make more than one passive copy, the second passive copy should use a log replay delay of 24 hours.
- Divide user mailboxes between two or more exchange databases according to the *Microsoft Exchange Capacity Planning Guide*.
- Keep moderate-sized Exchange databases; if databases are large, backup times may increase.

Best practices for backing up Exchange Information Store data

- When you run full backups, enable the option for Granular Recovery Technology (GRT). The GRT option lets you restore individual mail messages and folders from a database backup without the need for a separate mailbox backup.

“ ”

Note: For information on best practices for Backup Exec Granular Recovery Technology (GRT) with an Exchange Information Store backup, refer to the Granular Recovery Technology Best Practices. Exchange 2007 and later does not support individual mailbox backup.



- Veritas recommends that you do not send an incremental GRT-enabled Exchange backup to a deduplication disk storage device. The transaction logs contain primarily unique data that does not deduplicate well. For best results, create a backup definition that runs a full backup of Exchange to a deduplication disk storage device, and then runs an incremental backup to a disk storage device.
- Change your default staging location if you run GRT-enabled backup jobs. The default location is used for recovery as well as staging GRT-enabled restore jobs. You should change the location to a volume that is not your system volume for faster performance.
- You should have less than 75,000 transaction log files for GRT-enabled backups. If you have more than 75,000 transaction log files, it may increase the amount of time that it takes to complete the backup job.
- Ensure that the scheduled maintenance for the Information Store does not run at the same time as the database backup. If you run these operations at the same time, it can cause issues with the Exchange Server databases.
- Run Exchange backup jobs separately from other backup jobs.
- Back up the Active Directory on a regular basis.
- Run a regular backup for System State and Shadow Copy Components, if applicable. These selections back up the Internet Information Service (IIS) metabase and the Windows registry.
- Run a backup after you make any changes to system settings or application settings.
- When you run offline backups, back up all of the files that make up the storage group, including any .Edb and .Stm files, and all transaction log files.
- For Exchange 2010/2013 DAGs that have three or more copies of the database, the consistency check can be disabled.

Best practices for recovering data for all versions of Exchange Information Store

- Be aware of the effect of the Restore all transaction logs; do not delete existing transaction logs option. After an operation runs with this option enabled, transactions in existing transaction logs are applied when you start or mount the Information Store database. If those transactions include any deletions that occurred after the backup ran, those deletions are also applied. As a result, the very data that you intend to recover may be deleted. In this situation, enable the Purge existing data and restore only the databases and transaction logs from the backup sets option. This option discards the Exchange data that was generated after the backup.

Alternatively, you can use a second recovery server. You can also use the Recovery Storage Group feature in Exchange 2007 or Exchange 2010 or later recovery database to perform the restore.

- If you must use the Microsoft Eseutil utility to repair the database, ensure that the recovery server has sufficient disk space. You may need as much as 125% of the actual size of the Information Store database. You can also specify another disk or volume as a temporary location on which to run the Eseutil utility. Refer to your Microsoft documentation for details.

Best practices for restoring data for Exchange Server 2007 or later

- Ensure that you specify a valid temporary location on the Exchange server for log and patch files. The temporary location must have enough space to accommodate the transaction logs that you want to recover.
- Read the Restore.env file if issues occur when you mount a database after a restore operation. Information in this file can help you troubleshoot issues. To read the file, run the Eseutil utility with the /cm switch. Refer to your Microsoft documentation for details.
- Select the Commit after restore completes option when you configure a restore job so that the database can be mounted. Run the Eseutil utility with the /cc switch to perform a manual hard recovery. Refer to your Microsoft documentation for details.
- Ensure the following if you restore to an Exchange server other than the source server:
 - Ensure that the recovery server is in the same Exchange forest as the original server.
 - Ensure that the Exchange server uses the same version of Exchange with the same service pack level or higher as the original server.
 - If you redirect the database to another database name, you must use a database name that is different from the name of the source database for Exchange 2010 or later. Also, an empty database must already exist on the target server with the option to allow overwrites enabled.

Best practices to plan for disaster recovery of an Exchange Server

- Perform tests periodically to ensure that disaster recovery and data recovery scenarios produce the expected results.
- Become familiar with the Microsoft documentation for Exchange database management, disaster plans, and recovery.
- Document the Exchange Server configuration in detail. Document any subsequent changes. Note all hotfixes and service packs that are applied.

Best practices for Backup Exec Agent for Microsoft SQL Server

Best practices include tips and recommendations to help you use Veritas Backup Exec Agent for Microsoft SQL Server (SQL Agent) effectively. For more information about the SQL Agent, see the *Backup Exec Administrator's Guide*.

General recommendations for using the SQL Agent

- Back up the entire Microsoft SQL Server.

Include the following in the backup job:

- Full SQL database backups
- Windows System State
- System drive backups of the hard drive or drives where Microsoft SQL resides
- System drive backups of the hard drive or drives where the Microsoft SQL databases reside
- Exclude all database files from an anti-virus scan.
- Use database, differential, and log backups to maximize your backup window. Combine these backup methods with backup strategies that address the following issues:
 - How much data loss can you accept if a failure happens between the time of the last backup and the time the loss occurred?
 - How many transactions are processed each day?
 - What are your users' expectations when a recovery is required? For example, do they expect a full recovery to the point at the time when the data loss occurred?
- Use only the SQL Agent to perform SQL full, differential, and log backups. If you use a third-party application, Backup Exec makes a new full backup with the SQL Agent.
- Run transaction log backups if the database is configured for the full recovery model to prevent unlimited log file growth. Backup Exec generates a success with exception job warning based on the current size of the log file.
- With the simple recovery model, copies of the transactions are not stored in the log file, which prevents transaction log backups from being run.
- Use snapshot technology with the backup jobs that use deduplication devices.
- If you restore a large database that use deduplication devices, ensure that you use the SQL Management Studio to reduce the amount of memory that the SQL instance uses.
- Veritas recommends that you run regular database consistency checks to verify the integrity of the database. Run a consistency check either before or after a SQL backup and after a SQL

restore. If you back up a database, transaction log, or file group that contains errors, these errors still exist when the backup is restored. In some cases, these errors can prevent a successful restore. Backup Exec lets you check the logical and physical consistency of the data before and after a backup. SQL reports any consistency check failures in the Backup Exec job log.

You should specify the following SQL backup and restore options for the consistency check:

- Consistency check before backup - Physical check only
- Continue with backup if consistency check fails
- Consistency check after restore - Physical check only
- To ensure recovery after a disaster, run periodic test restore jobs and ensure that they are included in your disaster preparation plan.
- To ensure recovery from a deduplication device after a disaster, restore the backup set that was created after the baseline deduplication backup set.

Best practices for security and database access with the SQL Agent

- Ensure that the Windows user account that you use to back up the SQL instances has System Administrator privileges.
- Ensure that Backup Exec has rights to the following registry keys:
 - HKEY_LOCAL_MACHINE\Software\Microsoft\Microsoft SQL Server
 - HKEY_LOCAL_MACHINE\Software\Microsoft\MSSQL

Best practices for Backup Exec Agent for Linux

Best practices include tips and recommendations to help you use Backup Exec Agent for Linux (Linux Agent) effectively. For more information about the Linux Agent, see the *Backup Exec Administrator's Guide*.

The following best practices help you use the Linux Agent effectively:

- Use the Linux Agent to back up user data and configuration files. However, because of the lack of open file handling support, Veritas does not recommend the Linux Agent as a complete system for disaster recovery.
- Select the Using modified time option for the backup method that you choose.
- Disable file locking for backup jobs.

- Exclude pipes and socks from backup jobs.
- Add the same directories manually to the Backup Exec Exclude list if you have added directories to the autofs automounter utility. When you add these directories to the Exclude list, you can inadvertently avoid backing up data from sources such as CD-ROM drives.
- Assign administrator permissions to the user or add the user in the beoper group, who backs up the computer.
- For files and directories on a host that you do not want to back up, enter file and directory exclusion information in the `ra1us.cfg` file. After you add the information, the file and the directories are ignored during backup.
- Run a small restore job periodically to test the validity of your storage media.
- If you have gvfs mounted directories on your Linux server, unmount the directories before trying to browse selections for Linux backups or running actual backups. Linux agent does not support backing up or restoring to gvfs filesystems. If a gvfs mounted directory is present, backup browse actual backup does not work and gives an error.
- Use a port number that is not in use by another application or service when you add a Linux agent in Backup Exec.

To change the port number, see the following URL:

https://www.veritas.com/support/en_US/article.000028883

Best practices for Backup Exec Agent for Microsoft SharePoint

Best practices include tips and recommendations to help you effectively use Backup Exec Agent for Microsoft SharePoint (SharePoint Agent). For more information about the SharePoint Agent, see the *Backup Exec Administrator's Guide*.

General recommendations for using the Agent for Microsoft SharePoint

- Use an account that already exists in the SharePoint Farm Administrators Group for backup and restore jobs.
- Make SharePoint backup selections from the appropriate SharePoint farm resource container in the list of servers on the Backup and Restore tab instead of selecting individual SharePoint servers. The SharePoint farm resource container represents the entire farm topology in your environment. When you select the SharePoint resources from the resource container, you ensure that you select all available SharePoint resources for backup.

- Use the Backup Exec SQL Agent, in addition to the SharePoint farm resources, to back up the master, model, and msdb databases to fully protect each SQL instance for disaster recovery purposes.
- Perform a backup of the SharePoint server after you install service packs or hot fixes. SharePoint restore jobs may not complete successfully if the databases were backed up with different patch levels.
- Perform full system backups of the SharePoint servers, including System State, system volume, and any other volumes that contain SharePoint-related files and folders. When you make full system backups, you fully protect the entire SharePoint environment.
- To restore individual items, ensure that Granular Recovery Technology (GRT) is enabled before you run backups. GRT is enabled by default for the SharePoint Agent. You can enable or disable GRT for SharePoint globally in the Backup Job Defaults dialog box or on the Backup Options dialog box for individual backup jobs.
- To improve the performance of restore jobs, duplicate GRT-enabled backup sets on tape to disk storage first and then run the restore jobs from the disk-based backup sets. A staging location is not required to restore GRT-enabled data from a disk storage device.
- Ensure that you use the same version of the Backup Exec Agent for Windows on all of the SharePoint resources.
- Do not restore the SharePoint Configuration DB resource unless you are in a disaster recovery situation. To recover from a disaster, restore the database after you restore all of the other SharePoint databases, components, and applications.

Best practices for Backup Exec Central Admin Server Option

Best practices include tips and recommendations to help you use the Backup Exec Central Admin Server Option (CASO) effectively. For more information about CASO, see the *Backup Exec Administrator's Guide*.

Best practices for upgrading CASO

The following best practices help you upgrade the Central Admin Server Option effectively:

- Ensure that no jobs are running and put scheduled jobs on hold before you perform an upgrade.
- Ensure that the most recent Backup Exec service packs are installed on all Backup Exec servers before you perform an upgrade.

- Stop all Backup Exec services on the central administration server and all managed Backup Exec servers before you perform an upgrade of the central administration server. Upgrade one managed Backup Exec server at a time and stop the services on each managed Backup Exec server before you upgrade it.
- Review the guidelines for upgrading CASO on the Veritas Knowledge Base at the following URL:

https://www.veritas.com/support/en_US/article.000029459

Best practices for running test jobs in a CASO environment

The following best practice applies to test jobs:

- Perform test run jobs only at a local Backup Exec server. You cannot dispatch test jobs from a central administration server to a remote managed Backup Exec server.

Best practices for optimizing network bandwidth in a CASO environment

The following best practices help you to effectively optimize your network bandwidth:

- Ensure that distributed catalogs are used. CASO uses distributed catalogs by default. Catalog files are stored on the managed Backup Exec server. However, the central administration server includes some catalog metadata to enable centralized restore.
- Use the Export-BEBackupDefinition command in the Backup Exec Command Line Interface to recreate an existing backup definition on a different server. This command is similar to the Copy Job option that was available in previous versions of Backup Exec.
- Use the Copy Settings to Other Servers option to set options, such as default job options, schedules, error-handling rules, and alert configurations from one server to another server. This option can be accessed from the Backup Exec button > Configuration and Settings.
- Use the Create Simplified Disaster Recovery Disk wizard on the central administration server to create recovery media for any managed Backup Exec server or central administration server.
- Access the Settings options from the central administration server's Storage tab to reduce the frequency of job status updates that are sent from the managed Backup Exec server to the central administration server. To access the Send status updates to the central administration server every option, you must select Custom in the Connection settings field, and then select Yes in the Send active job status updates to the central administration server field.
- Access the Settings options from the central administration server's Storage tab to increase the amount of time that Backup Exec waits before it changes the Backup Exec server's status if the

Backup Exec server becomes unresponsive. To access the Communication stalled option, you must select Custom in the Connection settings field.

- Enable communications between the managed Backup Exec server and the central administration server before you delete a managed Backup Exec server from a CASO configuration. By enabling communications, Backup Exec can remove all of the necessary components for the deleted server from the database.
- Contact Veritas Technical Support for assistance before you reconfigure CASO. Reconfiguration errors can require a recovery of the managed Backup Exec servers, the central administration server, or the entire CASO environment.

Best practices for Backup Exec Agent for Oracle on Windows and Linux Servers

Best practices include tips and recommendations to help you use Backup Exec Agent for Oracle on Windows and Linux Servers (Oracle Agent) effectively. For more information about the Oracle Agent, see the *Backup Exec Administrator's Guide*.

The following best practices help you use the Oracle Agent effectively:

- Enable the Oracle archive log mode and the Oracle automatic archival of log files.
- Know the DBID and other important configuration details of the database.
- Know the names of the `init<SID>.ora` and the `spfiles` for the instances on the Oracle server.
- Do not store the RMAN Repository on the same server that holds the database that you want to back up.
- Back up your current control file when you run a Database Administrator (DBA) initiated job. If you have a backup of the current control file, then you do not have to search media to find a control file that is available for recovery.
- Test recovery scenarios often to get comfortable with the restore procedures. Oracle recovery can be complex and is often time-sensitive due to the nature of the data involved. We recommend that you coordinate test plans and configuration activities with your Oracle DBA to be sure that restore procedures are confirmed.
- Use RMAN scripts to do the following:
 - Delete all archive log copies in a multiplexed archive log configuration.

“ ”

Note: You can use Backup Exec to delete all non-multiplexed, single location archive logs.

“ ”

- Run RMAN optimization.

The following best practices should be considered when you back up Oracle databases:

- Take a full backup whenever you make structural changes to a database.
- Do not delete archived log files unless you have two confirmed backups of each log.
- Create Oracle-specific media sets and backup jobs for the following reasons:
 - RMAN can manage media retention and can communicate to the Backup Exec server that backup sets are expired. RMAN can successfully manage the media's retention period as long as unrelated backup sets are not present. Unrelated backup sets may have retention periods that are longer than the RMAN retention period.
 - Media sets for Oracle backups should have a retention period that is greater than the CONTROL_FILE_RECORD_KEEP_TIME setting. By default, the CONTROL_FILE_RECORD_KEEP_TIME is 7 days. If the media sets for the Oracle backups have a greater retention period, backup sets are not overwritten and RMAN is not updated.
 - When you configure multiple job streams in Oracle, additional resources such as file systems can cause more devices than expected to allocate drives.
- Load balance Oracle jobs between managed Backup Exec servers in a CASO environment. However, this scenario means that archived log file backups may reside on multiple managed Backup Exec servers, which makes restores complicated.
- Consider port re-assignments when you use RALUS in a Linux environment. Applications such as Webmin that uses port 10000 can interfere with RALUS operations.
- Enable the Oracle block change tracking for faster incremental backups.
- Enable Backup Exec compression when you configure general options for backup jobs.
- Ensure that you enter the fully qualified domain name of the Oracle server when you add it to list of servers on the Backup and Restore tab.
- Ensure that you add the fully qualified domain name of the Oracle server and the logon account name to the Backup Exec server's list of Oracle servers and authentication credentials.

The following best practices must be considered if you use the Oracle 12c database:

- Take a full backup of a container database (CDB) whenever there are any structural changes, such as addition of a new pluggable database (PDB).
- Include the root of the CDB in the backups to ensure that metadata of the CDB is always backed up.
- Run the database in the archive log mode to ensure that the database can be recovered to point in time.
- If the CDB is in no-archive log mode, then before backing up the PDBs, shut down the CDB. To avoid shutting down the CDB, you can either run the database in archive log mode or run a DBA-initiated backup of PDBs.
- Oracle recommends users to not only restore the root because it might cause metadata inconsistencies. Instead, you should recover the whole CDB.
- If the PDB point-in-time (PIT) restore fails, then consecutive jobs might also fail with the following error message:

ORA-19852: Error creating services for auxiliary instance.

This error occurs because the previous failed PIT restore attempts were not cleaned up properly. To solve this issue, you must clean the failed database PIT restore attempts. Perform the following steps to clean up the failed auxiliary service creation attempts:

- Use the Database (DB) PIT recovery package to clean up the auxiliary instance in case of failed PIT jobs:

```
SQL> exec dbms_backup_restore.manageauxinstance ('DBPITR',1);
```

- Also, run the last set of commands in the RMAN script to clear the RMAN configuration.

```
CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' CLEAR;
```

```
CONFIGURE AUXILIARY CHANNEL DEVICE TYPE 'SBT_TAPE' CLEAR;
```

Best practices for Backup Exec NDMP Option

Best practices include tips and recommendations to help you use the Backup Exec NDMP Option effectively. For more information about the NDMP Option, see the *Backup Exec Administrator's Guide*.

The following best practices help you use the NDMP Option effectively:

- Match the port number that Backup Exec uses to the port number that the device's NDMP service uses when you add an NDMP device.
- Add the NDMP port number that you want to use to the Backup Exec server file `%systemdrive%\Windows\System32\drivers\etc\services`.
- Do not include NDMP resources in a job that is software-encrypted. Software encryption is not supported for NDMP resources. Therefore, if you include NDMP resources in a job that is software-encrypted, the status of the job is "Failed."
- Verify that the NDMP protocol is enabled on the NDMP device and note the NDMP logon information.
- You cannot back up data from an NDMP server to a simulated tape library or to a tape device that is attached to a Backup Exec Remote Media Agent for Linux.

Best practices for Backup Exec reports

Best practices include tips and recommendations to help you use the Veritas Backup Exec reports effectively. For more information about the Backup Exec reports feature, see the *Backup Exec Administrator's Guide*.

Backup Exec includes standard reports that show detailed information about your system. You can also create the reports that contain information to meet the specific requirements of your organization.

Report formats

Reports can be viewed and printed in the following formats:

- PDF
- HTML
- XML
- Microsoft Excel (XLS)
- Comma Separated Value (CSV)

The following best practices can ensure smooth operations when you format or print Backup Exec reports:

- To properly format Backup Exec reports, you must configure a default printer in Windows even if you don't intend to print reports.
- To print reports in HTML, PDF, or Excel format, use the printer's landscape orientation.

- To create complex custom reports, you need a working knowledge of SQL and detailed information on Backup Exec's report database schema.

Best practices for Backup Exec and Veritas Update

Best practices include tips and recommendations to help you use Backup Exec and Veritas Update effectively. For more information about Veritas Update, see the *Backup Exec Administrator's Guide*.

The following best practices can help ensure the effective operation of Veritas Update:

- Run Veritas Update before and after each Backup Exec installation or upgrade.
- Schedule Veritas Update to run when backup jobs are idle to optimize system performance.
- On the Backup and Restore tab, right-click the remote computer and then select Update to update the Agent for Windows on remote computers with the same patches that were installed on the Backup Exec server. Veritas Update does not update the Agent for Windows on remote computers. The Downloads section of the Symantec Technical Support website contains information about what each Veritas Update service pack contains.

Best practices for Backup Exec Simplified Disaster Recovery

Best practices include tips and recommendations to help you use Backup Exec Simplified Disaster Recovery (SDR) effectively. For more information about SDR, see the *Backup Exec Administrator's Guide*.

The following best practices help you use SDR effectively:

- Review all of the requirements for SDR in the *Backup Exec Administrator's Guide*.
- When you run backup jobs with critical system components selected for Simplified Disaster Recovery preparation, Veritas recommends that you do the following:
 - If you install Backup Exec into an existing SQL instance, use the Backup Exec Agent for Microsoft SQL to periodically back up the SQL system database.
 - Avoid excluding files from the backup using the Selection Details tab.
- When you make considerable hardware changes to the computer such as if you change the Host Bus Adaptor or network interface card of an SDR protected system, follow these steps:
 - Run an SDR backup of the system to back up the new hardware drivers.
 - Create a new SDR disk or customize the existing SDR disk. This ensures embedding the new drivers into the SDR disk that might be required during SDR restore.

- Use SDR to perform a test recovery of a non-production computer before a disaster occurs. You can use a virtual environment such as VMware or Hyper-V for test recovery purposes. A test recovery lets you familiarize yourself with the SDR recovery process before it is needed.
- For disaster recovery purposes, consider the following:
 - The installed hard disks should be the same size or larger than the original.
 - The latest RAID, SCSI, or NIC (if remote) drivers are required.
- SDR supports recovering the computers that use the Unified Extensible Firmware Interface (UEFI) standard. However, backups of UEFI-based computers cannot be restored to standard BIOS-based computers.

“ ”

Note: UEFI computers use GPT-style disks. MBR-style disk data cannot be restored to GPT-style disk, and vice versa.

“ ”

For the computers that support both UEFI and BIOS firmware types, you must start the computer using UEFI firmware if you backed up the computer in that mode.

- If you have OEM partitions such as Dell Utility partitions on the system, they are considered part of a computer's critical system components and are backed up and restored as such.
- Ensure that Backup Exec Database maintenance runs after you configure new storage devices on your Backup Exec server but before you run an SDR-enabled backup job. This ensures that the Backup Exec Database contains the latest storage device configuration details, which can be restored as part of an SDR recovery. You can set the schedule to run Backup Exec Database maintenance by enabling the option Enable Backup Exec database maintenance. Click the Backup Exec button, select Configuration and Settings, and then select Backup Exec Settings. In the left pane, click Database Maintenance.
- Specify an alternate location where Backup Exec can store the disaster recovery information files. These files contain specific information for each computer that you back up with SDR. The alternate location should be on another computer or on a different physical drive than the default location.

Best practices for Backup Exec Agent for Enterprise Vault and the Backup Exec Migrator

Best practices include tips and recommendations to help you use Backup Exec Agent for Enterprise Vault and the Backup Exec Migrator effectively. For more information about the Enterprise Vault Agent and the Backup Exec Migrator, see the *Backup Exec Administrator's Guide*.

The following best practices help you use the Enterprise Vault Agent and the Backup Exec Migrator effectively:

- Use the Enterprise Vault service account or an account with rights to access the restore selections as the default logon account. Otherwise, you may have to enter credentials for each Enterprise Vault resource that you restore.
- Avoid backing up the Enterprise Vault folders and databases using SQL and NTFS agents. SQL and NTFS agents can interfere with Enterprise Vault Agent backup jobs and may affect the Enterprise Vault Agent's ability to restore the data.

“ ”

Note: You cannot select Enterprise Vault entries such as NTFS folders and SQL databases for backup. However, you can select Microsoft's Common Internet File System (CIFS) folders on network-attached storage devices where these entities may reside.

“ ”

- Back up the Enterprise Vault Directory database after you make any configuration changes in Enterprise Vault.
- Make sure Enterprise Vault 8.x components are not in Backup mode before you back up the Enterprise Vault 8.x Directory database.
- Do not allow the backup window and archive window to overlap.
- Do not allow the backup window and the Backup Exec migration window to overlap.
- Restore the Directory database in a separate job when you restore an Enterprise Vault installation. After you restore the Directory database, you can restore other Enterprise Vault components and partitions.
- Restore all Full, Differential, and Incremental backup sets of the vault store database in a single restore job.

- Run the Enterprise Vault recovery tools after you restore Enterprise Vault. The recovery tools synchronize Enterprise Vault with the newly restored databases.
- Ensure that the appropriate Enterprise Vault product scripts (SQL scripts) run after a redirected restore job completes, but before the redirected restore job is scheduled again.
- If you install both the Veritas Backup Exec NDMP Option and the Enterprise Vault Agent, pick only one product to protect an Enterprise Vault partition that resides on NDMP filers.
- Create separate backup jobs for NDMP backups of Enterprise Vault data. You may see a performance improvement when you use the Veritas Backup Exec NDMP Option.
- Do not change the recovery model of any database that Enterprise Vault creates. Enterprise Vault configures each database in full recovery mode when it creates them.

Consider the following best practices when you use the Backup Exec Migrator:

- Veritas recommends that you regularly back up the Backup Exec catalogs.

If the catalogs become corrupt, you can restore them from backups. After you restore the catalogs, you must re-catalog the storage media on which Backup Exec Migrator data is stored. Re-cataloging the storage media ensures that the latest catalog entries are available.

- For best performance, configure the Backup Exec Migrator to migrate data to a backup-to-disk folder and then to a tape device by using a duplicate job.
- In the Enterprise Vault Migration options tab, set the time period for Remove collection files from primary storage to something greater than zero days.

If you set the time period to zero days, Enterprise Vault immediately deletes the migrated data from the partition.

If you set the time period to zero days, Veritas recommends the following:

- Increase the number of concurrent jobs that are allowed for the backup-to-disk folder you use for migration purposes.

Increase the number of concurrent jobs based on the following formula:

$$\langle \text{number of recommended concurrent jobs} \rangle = \langle \text{number of installed tape drives plus two} \rangle$$

For example, if you have two installed tape drives, you should configure the backup-to-disk folder to allow four concurrent jobs.

Concurrent jobs let the Backup Exec Migrator to continue migrating data to disk storage while tape drives process duplicate jobs in a staged migration environment.



Note: You can increase the number of concurrent jobs that run by increasing the total concurrency level of the backup-to-disk devices.



- Veritas recommends that you first collect all of the archived files in one collection and migration operation, and then migrate these files in the next collection and migration operation. This process ensures that the Backup Exec Migrator creates a single job for each migration operation, which improves the migration performance.

Best practices for Backup Exec Granular Recovery Technology

Best practices include tips and recommendations to help you use Veritas Backup Exec and Granular Recovery Technology (GRT) effectively. For more information about Granular Recovery Technology, see the *Backup Exec Administrator's Guide*.

The following best practices can help ensure the effective operation of Granular Recovery Technology:

- Ensure that GRT is enabled before you run backups if you intend to be able to restore individual items.

GRT is enabled by default for the following resources. It can be enabled or disabled in the Backup Options dialog when you create a backup job:

- Microsoft Active Directory
- Microsoft Exchange Server
- Microsoft SharePoint
- VMware
- Hyper-V
- Back up your current or most recent GRT-enabled backup jobs to disk. It is more convenient to work with GRT-enabled jobs on the volumes that do not have file size limitations. You can create duplicate backup jobs and send copies of your backups to tape for archival purposes.
- Use disk storage on a volume that does not have file size limitations as the destination for any backups that are enabled for GRT. An NTFS drive is an example of a volume without file size

limitations. Some examples of volumes that have file size limitations include FAT and FAT32 volumes.

- Review the requirements for staging locations in the Administrator's Guide.

You must use a staging location for GRT-enabled jobs in the following scenarios:

- You back up to or restore from a volume with file size limitations.
- You back up to tape.

Tape backups require a staging location that is at least as large as the data that you back up. Backup Exec extracts the granular data to the staging location while it is being cataloged. When you restore granular data from a tape backup, you must specify a staging location to store all of the backup sets that are required for the restore job as well as a separate staging location of at least 1 GB for the GRT processing.

- You back up Active Directory data or Exchange data to a disk.
- Use a volume that is not your system volume for a staging location. The volume on which the staging location resides should have at least as much available space as the size of your largest GRT-enabled backup job. You can change the default staging locations in the Backup Exec Settings.
- Do not allocate the maximum size for backup files. If you enable the Preallocate disk space incrementally up to the maximum file size option in the storage details, Backup Exec creates a file that is as large as the maximum file size that you specified. Since GRT information is stored in IMG media, the file does not hold backup data. The extra space that the file occupies can often lead to failed jobs because of low disk space.
- Run a full GRT-enabled backup job periodically if your backup strategy uses frequent incremental GRT-enabled jobs. Each incremental GRT-enabled job requires a small amount of internal storage. If this storage amount increases too much, it can negatively affect your system resources.
- Duplicate GRT-enabled backup sets to disk storage first and then run the restore jobs from the disk-based backup sets, if you must run multiple restores from the same backup set on tape. GRT restores from backup sets on tape must be staged to disk first. The staged data is not retained after the restore completes. Duplicating the backup sets to disk storage eliminates the need to stage the data multiple times and improves the performance of the restore jobs.
- Monitor your processor, disk, and memory usage if you experience any performance issues. Recovery and staging of GRT data may require more than the minimum system requirements, depending on the volume of data in the backup sets.

- Do not use software compression or encryption for GRT-enabled backup jobs. The compression and encryption processes are resource-intensive. Enabling either software compression or encryption can result in degraded performance for GRT-enabled backup jobs.
- Backup Exec does not store the granular backup sets on disk in encrypted form when you enable encryption for the GRT-enabled backup jobs that are sent to disk, deduplication, and disk cartridge devices. Only the backup sets for the backup sources that do not support GRT are stored in encrypted form. All the backup sets for the backup jobs that are sent to cloud, OpenStorage, and tape devices are stored in encrypted form.

Best practices for Backup Exec Remote Media Agent for Linux

Best practices include tips and recommendations to help you use Veritas Backup Exec Remote Media Agent for Linux (RMAL) effectively. For more information about the RMAL, see the *Backup Exec Administrator's Guide*.

The following best practices help you use the RMAL effectively:

- Ensure that the storage device that is attached to the Linux server is supported. Also, confirm that the operating system can access the device before you start the RMAL.

You can find a list of compatible devices at the following URL:

<http://www.veritas.com/docs/000017788>

- Ensure that a minimum of 500 MB of available storage space exists on the Linux server when you use the Tape Library Simulator Utility. The available space includes hard disk space, flash drives, and USB drives. If there is not enough space, the jobs fail with an end-of-media error. You must either create available disk space or you must direct the jobs to another volume, and then start the jobs again.
- Add an RMAL from the central administration server if you use the Backup Exec Central Admin Server Option (CASO). You cannot add an RMAL from a managed Backup Exec server.
- If you use the Shared Storage option when you add a RMAL, use the Linux server's hostname and not its IP address. If you use the IP address, the Backup Exec Database cannot distinguish which device path to use for jobs.
- Only Microsoft Tape Format media are supported if you use an RMAL to restore data from the tapes that other applications create. Backup Exec does not support GRT-enabled backup and restore jobs targeted to an RMAL device.

Best practices for Backup Exec Agent for Microsoft Hyper-V

Best practices include tips and recommendations to help you effectively use the Backup Exec Agent for Microsoft Hyper-V. For more information about the Agent for Microsoft Hyper-V, see the *Backup Exec Administrator's Guide*.

General recommendations for using the Agent for Hyper-V

- Perform a periodic full backup.
- Back up to a disk-based storage device instead of to a tape device.
- Ensure that each virtual machine has a unique name. If you have two virtual machines with the same display name, backup jobs may fail.
- Ensure that the correct version of the Hyper-V Integration Components is installed. The version of the Hyper-V Integration Components should match the version of Hyper-V that you use. Note that Hyper-V Integration Components cannot be disabled.

Best practices for using Granular Recovery Technology (GRT) with the Agent for Hyper-V

- To use GRT, install the Agent for Windows on the virtual machine on which data is to be restored.
- Do not use GRT for system file restore for disaster recovery of a virtual machine. Instead, restore the full virtual machine.
- If you want to use GRT for some virtual machines, but not for others, set up jobs as follows:
 - For virtual machines that do not require GRT, create a backup job with all GRT options disabled.
 - For virtual machines that require GRT, create one of the following jobs:
 - For virtual machines that require file/folder GRT, but do not have SQL, Exchange, Active Directory, or SharePoint installed, create a backup job and select the option Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual files and folders from virtual machines.
 - For virtual machines that require GRT for SQL, Exchange, Active Directory, or SharePoint, create a backup job and select the appropriate application GRT checkboxes. For example, to use GRT for Exchange, select the option Enable GRT for Microsoft Exchange databases and mailbox items on virtual machines. File/folder GRT can also be enabled in this job if you also require file/folder GRT. The Agent for Windows must be installed on the virtual machines that you include in this backup job.

You must provide the appropriate credentials to access the virtual machines and the applications that reside on them.

- For virtual machines that do not support GRT, such as Linux and Mac, create a separate backup job and deselect the four options to Enable GRT for \<Microsoft application name\> on virtual machines. If a virtual machine does not support GRT, but the options to enable application GRT are selected, the job will complete with exceptions.

“ ”

Note: Do not include the same virtual machine in multiple backup jobs. If Agent-based backups are also being performed, schedule them so that they do not overlap with the GRT-enabled backups of the virtual machines.

“ ”

- If you have Microsoft SQL installed on a virtual machine, you can select the option Run a SQL Log backup after backing up the virtual machine to back up the SQL logs for the databases that use logging. After the logs are backed up, the data from the logs is committed to the database and the log is emptied in order to receive more data. If you do not select the option to perform SQL log backups, the SQL logs continue to grow until the disk is full or until you perform a manual backup job to back up the logs. The option to run a SQL log backup is located on the Virtual Machines dialog box.
- To use GRT for individual items from Microsoft Active Directory, Exchange, SharePoint, or SQL, do the following:
 - Install the Agent for Windows on the guest virtual machine.
 - Ensure that you have a valid license for each application that you want to protect.
 - Ensure that the credentials for the guest virtual machine are valid for the application that you want to protect.
 - Use the same requirements for backing up Microsoft Exchange on a physical computer when you back up Exchange on a virtual machine.

Best practices when using the Instant GRT option for GRT-enabled jobs

- In a CASO environment, ensure that the logon accounts used for backups are added to the list of logon accounts on the central administration server and the managed Backup Exec servers.

- The storage that hosts the backup sets must be online when you browse for individual items that you want to restore because Backup Exec mounts the backup sets dynamically. For incremental and differential backup sets, all such related backup sets should also be accessible during restore.
- If a CASO environment, if a Backup Exec server tries to browse the backup sets of another Backup Exec server and if a firewall is configured between them, you must open ports on the servers.

Veritas recommends to browse backup sets either from the managed Backup Exec server on which the backup jobs were run or from the central administrative server.

For the list of ports, see the "Backup Exec ports" and "Backup Exec listening ports" topics in the *Backup Exec 15 Administrator's guide*.

Best practices for Backup Exec Agent for VMware

Best practices include tips and recommendations to help you effectively use the Backup Exec Agent for VMware. For more information about the Agent for VMware, see the *Backup Exec Administrator's Guide*.

General recommendations for using the Agent for VMware

- Review the Backup Exec Software Compatibility List (SCL) to confirm that the applications that you want to back up are supported in this version of Backup Exec.

https://www.veritas.com/support/en_US/article.000017788

- Perform a weekly full backup.
- Back up to a disk-based storage device instead of to a tape storage device.
- If you have VMware Tools for vSphere 6.0 or earlier installed, ensure that only one VSS Provider is installed on guest virtual machines. The VMware VSS Provider and the Veritas VSS Provider cannot reside on the same guest virtual machine.

Both of the VSS Providers can be installed on virtual machines that have VMware Tools for vSphere 6.5 or later installed.

Best practices for selecting VMware data to back up

- You cannot select individual drives, folders, and files from the virtual machines that appear when you expand VMware vCenter and ESX Servers.

- Note that if you select the vCenter server or the ESX server as a backup resource, all virtual machines are backed up. For a recurring job, any virtual machines that you add after the job is created are included in the backup.
- Note that if you select to back up the vCenter server or the ESX server, the backup job does not include the following:
 - Configuration files for the vCenter server or the ESX server
 - Physical Raw Disk Mapping (RDM) devices
 - Independent disks

Best practices for using Granular Recovery Technology (GRT) with the Agent for VMware

- Install the Backup Exec Agent for Windows on any virtual machines on which you want to use Backup Exec's Granular Recovery Technology.
- If you want to use GRT for some virtual machines, but not for others, set up jobs as follows:
 - For virtual machines that do not require GRT, create a backup job with all GRT options disabled.
 - For virtual machines that require GRT, create one of the following jobs:
 - For virtual machines that require file/folder GRT, but do not have SQL, Exchange, Active Directory, or SharePoint installed, create a backup job and select the option Use Backup Exec Granular Recovery Technology (GRT) to enable the restore of individual files and folders from virtual machines.
 - For virtual machines that require GRT for SQL, Exchange, Active Directory, or SharePoint, create a backup job and select the appropriate application GRT checkboxes. For example, to use GRT for Exchange, select the option Enable GRT for Microsoft Exchange databases and mailbox items on virtual machines. File/folder GRT can also be enabled in this job if you also require file/folder GRT. The Agent for Windows must be installed on the virtual machines that you include in this backup job. You must provide the appropriate credentials to access the virtual machines and the applications that reside on them.
 - For virtual machines that do not support GRT, such as Linux and Mac, create a separate backup job and deselect the four options to Enable GRT for \<Microsoft application name\> on virtual machines. If a virtual machine does not support GRT, but the options to enable application GRT are selected, the job will complete with exceptions.

“ ”

Note: Do not include the same virtual machine in multiple backup jobs. If Agent-based backups are also being performed, schedule them so that they do not overlap with the GRT-enabled backups of the virtual machines.

“ ”

- To use GRT for individual items from Microsoft Active Directory, Exchange, SharePoint, or SQL, do the following:
 - Install VMware Tools on the guest virtual machine.
 - Install the Agent for Windows on the guest virtual machine. You must install the VMware Tools before you install the Agent for Windows.
 - Ensure that you have a valid license for each application that you want to protect.
 - Ensure that the credentials for the guest virtual machine are valid for the application you want to protect.
 - Use the same requirements for backing up Microsoft Exchange on a physical computer when you back up Exchange on a virtual machine.

Best practices when using the Instant GRT option for GRT-enabled jobs

- In a CASO environment, ensure that the logon accounts used for backups are added to the list of logon accounts on the central administration server and the managed Backup Exec servers.
- The storage that hosts the backup sets must be online when you browse for individual items that you want to restore because Backup Exec mounts the backup sets dynamically. For incremental and differential backup sets, all such related backup sets should also be accessible during restore.
- If a CASO environment, if a Backup Exec server tries to browse the backup sets of another Backup Exec server and if a firewall is configured between them, you must open ports on the servers.

Veritas recommends to browse backup sets either from the managed Backup Exec server on which the backup jobs were run or from the central administrative server.

For the list of ports, see the "Backup Exec ports" and "Backup Exec listening ports" topics in the *Backup Exec 15 Administrator's guide*.

Best practices for using Backup Exec with server clusters

Best practices include tips and recommendations to help you use Backup Exec effectively in a clustered environment. For more information about Backup Exec and clustered environments, see the *Backup Exec Administrator's Guide*.

The following best practices apply when you use Backup Exec in clustered environments:

- To prevent possible loss of data in failover situations, Veritas recommends that you use the backup method Incremental - Using modified time when you configure backup jobs. Do not use the backup method, Incremental - Using archive bit (reset archive bit).

When you apply Backup Exec patches to clustered Backup Exec servers or to a central administration server, Veritas recommends the following to prevent possible data loss:

- Use the Microsoft Cluster Administrator to set each of the clustered Backup Exec services so that they do not automatically restart after a failure occurs. By doing that, you prevent the cluster from changing nodes in the event of a failure.

After you finish applying the patches, you should reset each of the clustered Backup Exec services so that they restart after a failure occurs.

Best practices for Backup Exec Deduplication Option

Best practices include tips and recommendations to help you effectively use the Backup Exec Deduplication Option. For more information about the Deduplication Option, see the *Backup Exec Administrator's Guide*.

The following best practices apply to deduplication disk storage devices:

- Use a dedicated volume for a deduplication disk storage device.
- Ensure that the volume you use for a deduplication disk storage device has enough space to hold data from multiple servers. Veritas recommends that you use a dynamic volume that can grow as needed.
- Perform regular backups of the deduplication disk storage device to removable media. Create this job separately from the backup job that you use to protect the rest of your system. In the backup selection list, the Backup Exec Deduplication Storage node is located under the User Data node under Shadow Copy Components.

- Perform a duplicate backup job so that you have redundant backups of the data in your deduplication disk storage device. This method is preferred over performing full backups using the deduplication writer because this method is easier to control.
- Use a dedicated logon account when creating a deduplication disk storage device.
- Use the spausser.exe utility if you have to change the password for the Backup Exec logon account that you specified when you created a deduplication storage folder. For more information, refer to the following tech note:

https://www.veritas.com/support/en_US/article.000005534

- Exclude the deduplication disk storage device from your antivirus scans. If an antivirus scanner deletes or quarantines the files from the deduplication disk storage device, access to the deduplication disk storage device may be disabled.
- Do not use Backup Exec's Simplified Disaster Recovery (SDR) to recover local data from the local deduplication disk storage device. SDR recovery from deduplication disk storage devices is supported for remote resources only.
- On the drive properties where the deduplication disk storage device is located, ensure that the following option is not selected: Allow files on this drive to have contents indexed in addition to file properties.

To access the drive properties, in Windows Explorer, right-click the drive on which the deduplication disk storage device is located. Click Properties, and then click the General tab.

- Do not delete any files from the deduplication disk storage device. To reclaim space from the deduplication disk storage device, refer to the following tech note:

https://www.veritas.com/support/en_US/article.000017049

- Do not directly add any files to the deduplication disk storage device. Backup Exec does not recognize files that are placed in the deduplication disk storage device by other applications.
- Review the tech note titled "Getting the most out of the Deduplication Option and Deduplication Storage Folders" for more details about using deduplication disk storage devices. The tech note is available from the following link:

https://www.veritas.com/support/en_US/article.000009541

The following best practices apply to using deduplication with SQL:

- Set the option Use snapshot technology as a backup job property for the SQL backup job. This option allows data to be deduplicated in the most effective way.

- Use Microsoft SQL Server Management Studio to restrict SQL's use of physical memory. This option can be set for each SQL instance by selecting Properties > Memory. Restricting SQL to use 85% to 90% of total physical memory can prevent a situation in which backup rates may be reduced to 10% of the expected rate due to memory restrictions.

The following best practices help you to use the Deduplication Option effectively:

- Use client-side deduplication for all deduplication jobs, except in the following situations:
 - The remote computer has limited or no resources available to dedicate to the deduplication process during backup jobs.
 - Your environment contains multi-server configurations of Microsoft SharePoint servers, Microsoft Exchange 2010 DAG servers, and Veritas Enterprise Vault servers.
 - You want to back up virtual machines with the Agent for VMware.
- Ensure that the appropriate third-party vendor plug-in is installed for an OpenStorage Technology (OST) device and exists in the Backup Exec directory. The plug-in enables Backup Exec to detect the OST device and display the device in the server list.
- Disable RAID caching on the disk where the deduplication disk storage device is located.
- Ensure that each backup job has a verify job that runs after the backup job finishes, or that is scheduled to run as a separate job.
- Apply the following Microsoft hot fixes to improve performance:

<http://support.microsoft.com/kb/979612>

<http://support.microsoft.com/kb/982383>

- Computer disk speeds have the following effects on deduplication performance:
 - Computers with disk speeds greater than 200 MB per second have optimal read and write performance for deduplication.
 - Computers with disk speeds between 150-200 MB per second have sufficient read and write speed for deduplication.
 - Computers with disk speeds between 100-150 MB per second have some operations with degraded performance.
 - Computers with disk speeds less than 100 MB per second experience poor performance. You should improve disk reads and writes before you install and run deduplication.

- Connection storage area network (Fibre Channel or iSCSI), direct-attached storage (DAS), iSCSI, or internal disks are supported. Removable disks including USB, eSATA, and FireWire devices are not supported.
- The Backup Exec server should have redundant connectivity to the deduplication disk storage.
- The storage network must be a dedicated, low-latency network with a maximum of 1-millisecond latency per round trip.
- The storage network must have enough bandwidth to meet your throughput objectives. Veritas supports the following storage network bandwidths:
 - iSCSI SANs with a bandwidth of at least 10 Gb per second.
 - Fibre Channel SANs with a bandwidth of at least 4 Gb per second.
- Veritas requires a minimum bandwidth of 130 MB per second for read and write performance. Bandwidth that is less than 130 MB per second may be used in smaller, less resource-intensive environments. However, as usage increases, deduplication requires more bandwidth to ensure adequate throughput for deduplication processes and backups. Otherwise, performance and stability are negatively affected.

Best practices for using Backup Exec Deduplication Option with the Central Admin Server Option

Best practices include tips and recommendations to help you effectively use the Backup Exec Deduplication Option with the Central Admin Server Option (CASO). For more information about the Deduplication Option or CASO, see the *Backup Exec Administrator's Guide*.

The following best practices apply to using the Deduplication Option with the Central Admin Server Option:

- Create a deduplication disk storage device on managed Backup Exec servers in a CASO environment.
- Do not run jobs that use a deduplication disk storage device on a central administration server. Instead, assign backup jobs to managed Backup Exec servers and select a Backup Exec server pool as a setting for the destination device.
- Share the deduplication disk storage devices that are on multiple managed Backup Exec servers with the Backup Exec server that is the replication target. This provides the best use of optimized duplication. For example, the main office site that has a managed Backup Exec server deployment shares the deduplication disk storage device with the central administration server at a remote site for disaster recovery purposes.

- Use the following recommendations for optimal catalog placement:
 - For local LAN deployments, use replicated catalogs.
 - For site-to-site optimized duplication, use distributed catalogs.
- Do not use Backup Exec server pools as device targets in standalone backups. This prevents duplicate data from being hosted on multiple managed Backup Exec servers.
- Subdivide the remote computers that are enabled for direct access sharing among the servers that host the deduplication disk storage folders to best optimize several managed Backup Exec servers at a single site. For example, a large site with two managed Backup Exec servers and 10 remote computers should split the remote computers evenly between the managed Backup Exec servers.
- Limit the sharing of remote computers that are enabled for direct access sharing with other Backup Exec servers that use deduplication disk storage. This prevents duplicate data from being hosted on multiple backup servers.
- To perform optimized duplication between a central administration server and a managed Backup Exec server, the following WAN requirements must be met:
 - Less than one percent packet loss during transmissions
 - Less than 250 milliseconds network latency (round trip)
- Computer disk speeds have the following effects on deduplication performance:
 - Computers with disk speeds greater than 200 MB per second have optimal read and write performance for deduplication.
 - Computers with disk speeds between 150-200 MB per second have sufficient read and write speed for deduplication.
 - Computers with disk speeds between 100-150 MB per second have some operations with degraded performance.
 - Computers with disk speeds less than 100 MB per second experience poor performance. You should improve disk reads and writes before you install and run deduplication.
- Connection storage area network (Fibre Channel or iSCSI), direct-attached storage (DAS), iSCSI, or internal disks are supported. Removable disks including USB, eSATA, and FireWire devices are not supported.
- The Backup Exec server should have redundant connectivity to the deduplication disk storage.

- The storage network must be a dedicated, low-latency network with a maximum of 1-millisecond latency per round trip.
- The storage network must have enough bandwidth to meet your throughput objectives. Veritas supports the following storage network bandwidths:
 - iSCSI SANs with a bandwidth of at least 10 Gb per second.
 - Fibre Channel SANs with a bandwidth of at least 4 Gb per second.
- Veritas requires a minimum bandwidth of 130 MB per second for read and write performance. Bandwidth that is less than 130 MB per second may be used in smaller, less resource-intensive environments. However, as usage increases, deduplication requires more bandwidth to ensure adequate throughput for deduplication processes and backups. Otherwise, performance and stability are negatively affected.

Best practices for using hot-pluggable devices such as USB devices in a drive rotation strategy

Best practices include tips and recommendations to help you effectively use hot-pluggable devices in Backup Exec. For more information about using storage devices in Backup Exec, see the *Backup Exec Administrator's Guide*.

The following best practices apply to using hot-pluggable devices in Backup Exec:

- Do not run incremental Granular Recovery Technology (GRT)-enabled backups if you back up to hot-pluggable devices that you rotate.
- If you send differential backups to hot-pluggable devices that you rotate, then when the differential backups run, you must keep the hot-pluggable device that contains the last full backup attached to the server.
- If your environment contains devices to which vendors have assigned the same disk signature, you should check if a new hot-pluggable device has the same disk signature as another. To check for identical disk signatures, configure the devices one at a time. After you configure a device, view the Storage tab in the Backup Exec administration console and check if any devices that were previously configured are now online. If so, then most likely that device has the same disk signature as another device. Use the UNIQUEID DISK command in DISKPART to manually set each hot-pluggable device to a different disk signature.

Alternatively, you can attach all of the hot-pluggable devices and bring them online at the same time. You can attach the devices to the server if there are enough free ports, or attach the devices to a hub.

- Configure only one hot-pluggable device per disk storage. Do not configure all of your hot-pluggable devices on one disk storage.
- Create a storage pool just for the hot-pluggable devices that you want to use in a rotation strategy. Specify the storage pool as the storage destination when you run backups.
- If you change the default Windows setting for the removal policy in the properties of the disk in Device Manager, you should use the Safely Remove Hardware notification tray icon to disconnect the hot-pluggable device safely.
 - For Windows Server 2008 R2/2012/2012 R2, the default setting is called Quick Removal.
 - For Windows Server 2003/2008, the default setting is called Optimize for quick removal.

“ ”

Note: Be aware that when you unplug a hot-pluggable device, Backup Exec sends an alert that the device is offline; this is normal and expected.

“ ”

Best practices for Backup Exec database encryption keys

Best practices for Backup Exec database encryption keys

Best practices include tips and recommendations to help you use Backup Exec to manage the database encryption feature effectively. For more information about database encryption, see the *Backup Exec Administrator's Guide*.

The following best practices can help ensure the effective operation of database encryption:

- Export the database encryption key immediately after you install Backup Exec to ensure that you have a copy of it in the event of a server failure.

To export the database encryption key, complete the following steps:

- Click the Backup Exec button, select Configuration and Settings, and then click Backup Exec Settings.
- In the left pane, select Database Maintenance and Security.
- In the Path field, type the location to which you want to export the encryption key.

- Click Export.

The key is exported to the location that you specified. The key is named with a unique hash value. Backup Exec uses the name to identify the key later. Do not change the key's file name or file contents. If you want to export the key to additional locations, repeat the previous steps.

- Click OK.

Make sure that you export the database encryption key to a location that meets the following criteria:

- The destination is either on a physical volume that is assigned to a drive letter or a network share that is specified by a UNC path (network shares that are mapped to drive letters are not supported).
- The destination has enough disk space.
- The destination is accessible from the Backup Exec server.
- Backup Exec has permission to write to the destination.
- Save the database encryption key to a secure location. Veritas recommends that you save the key to an off-site location for increased security. It is your responsibility to ensure that the database encryption key is backed up.
- Exercise caution when you configure access rights for the Data folder in the Backup Exec install directory. The Data folder contains the Backup Exec Database, SSL certificates, and database encryption keys as well as other critical data. The Data folder is protected from unauthorized access using Windows Access Control Lists (ACL). You should ensure that only trusted users can access the Data folder.
- Refresh the database encryption keys periodically. Refreshing the database encryption keys helps to protect the server from any attacks that might try to decipher the keys.

For more information about refreshing the database encryption keys, refer to the *Backup Exec Administrator's Guide*.

Best Practices for Using the Veritas Backup Exec Cloud Connector

The Backup Exec Cloud Connector feature provides seamless and secure integration with 3rd-party cloud storage services, which enables direct-to-cloud backups and disk-to-cloud backups. To review the list of supported regions, see the Backup Exec hardware compatibility list at the following URL: <http://www.veritas.com/docs/000017788>

Before you begin using the Backup Exec Cloud Connector, you should be familiar with the following terms:

Common Cloud Terminology

- **Access key ID:** an alphanumeric code that allows access to the cloud storage. You must create an access key with the cloud storage service provider that you choose to use, such as Amazon S3 or Google, before you configure your cloud storage in Backup Exec. When you configure your cloud storage in Backup Exec, you must enter this key as the user name for the logon credentials for the cloud storage. If you copy and paste the key from the cloud storage provider, be sure that you don't copy white space.
- **Secret key:** an alphanumeric code that allows access to the cloud storage. You must create a secret key with the cloud storage service provider that you choose to use, such as Amazon S3 or Google, before you configure your cloud storage in Backup Exec. When you configure your cloud storage in Backup Exec, you must enter this key as the password for the logon credentials for the cloud storage. If you copy and paste the key from the cloud storage provider, be sure that you don't copy white space.
- **Bucket:** a logical unit of storage that stores objects, such as data and metadata. You must create a bucket in the cloud storage that you choose to use, such as Amazon S3 or Google, before you configure your cloud storage in Backup Exec. When you configure a backup-to-cloud job in Backup Exec, if you have created more than one bucket, then you must select the bucket that you want to use to store the data.

Pros and Cons of Common Cloud-based Backup Scenarios

Review the following table to determine the best backup scenario for your organization:

Table: Pros and Cons of Common Cloud-based Backup Scenarios

| BACKUP SCENARIO | PROS | CONS |
|------------------------------|--|--|
| Backup directly to the cloud | This is the simplest operation. It does not require any additional space on the local backup storage device. | Since cloud backups and restores may be slow depending on available bandwidth, they may not fit in your backup window. |

| BACKUP SCENARIO | PROS | CONS |
|--|--|--|
| | | A menu option to back up directly to the cloud is available. |
| | | For releases earlier than Backup Exec 16 FP1, to back up directly to the cloud, you must create a backup-to-disk job, and then edit the storage properties to select the cloud storage device. |
| Back up to disk, and then duplicate to the cloud | This operation provides a quick restore from a local copy. | This option requires additional disk space on the local backup storage device. |
| | This is usually faster than backing up directly to the cloud, so your backup window may not be impacted. | |
| | A menu option is available for this operation. | |
| Back up to a deduplication storage device, and then duplicate to the cloud | This operation provides a quick restore from a local copy. | The backup will no longer be deduplicated when it is copied to the cloud. |
| | In addition, it reduces the amount of disk space required for a local copy. | |

Using the Verify Option with Cloud-based Backups

Backup Exec can perform a verify operation to make sure that the media can be read after a backup job has been completed. By default, Backup Exec automatically verifies backed up data at the end of a backup job. However, you can also schedule the verify operation to take place at a later time or disable the verify operation altogether. You can change Backup Exec's verify options as part of the default backup settings or for individual backup jobs.

For cloud-based storage devices, by default, the Do not verify data for this job option is now selected in the Backup Options. Cloud vendors charge for operations that read data from and write data to the cloud. To avoid charges for reading data during the verify operation of a backup or duplicate job, this option is selected by default.

The Backup Exec Cloud Connector implements the integrity check mechanisms for S3 (Amazon, Google, and private cloud vendors supported by Backup Exec) and Azure compatible cloud storage. This is available from Backup Exec 16 and later.

Contact your cloud storage provider to determine the cost of reading data from and writing data to cloud storage. Then, you can weigh the cost of the verify operation versus the peace of mind of having verified data to determine if running the verify operation is worthwhile for your organization.

Using Amazon S3 versus Amazon Storage Gateway/VTL

Customers also have the option of leveraging a 3rd-party cloud gateway in addition to Backup Exec Cloud Connectors. One example of this is the Amazon Storage Gateway /VTL, which integrates with Backup Exec as an iSCSI disk storage or as a Virtual Tape Library (VTL).

The Amazon Storage Gateway is a separate service provided by Amazon and also has a separate monthly fee. Please refer to the Amazon website for more details about pricing.

The Amazon Storage Gateway acts as a local storage destination for Backup Exec and manages the data transfer to/from the cloud automatically and transparently.

The Amazon Storage Gateway VTL allows customers to create and store virtual tapes in a VTL powered by Amazon S3 and a Virtual Tape Shelf (VTS) powered by Amazon Glacier.

Amazon Storage Gateway /VTL does not charge for operations/requests and can be a viable alternative for customers who wish to avoid operation/request charges.

Refer to the AWS documentation for more details.

Using Granular Recovery Technology in Cloud-based Backups

Granular Recovery Technology (GRT) backups to Backup Exec Cloud Connector storage leverage Veritas OpenStorage Technology (OST). The data is then automatically staged locally as part of the restore job.

Cloud Vendor Changes

Cloud Storage providers charge for the use of their storage services based on multiple aspects, including (but not limited to):

- Time
- Storage amount
- Operations/requests
- Data transfers

Most vendors do not charge for data transfers into the cloud, but do charge for any transfers that exit their cloud; for example, for restores or even verify-operations that transfer information back to on-premises.

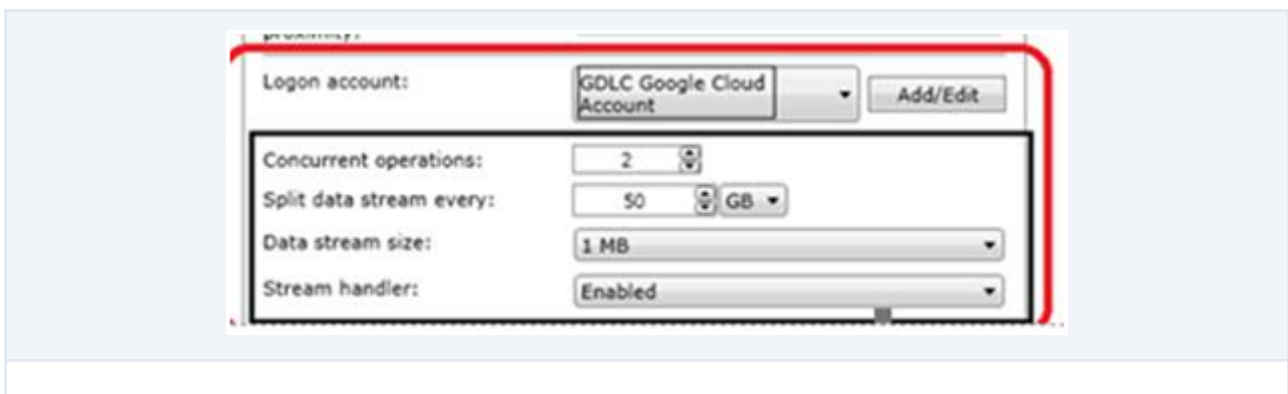
Additionally major vendors can provide multiple storage classes with different attributes, such as:

- Access frequency
- Access bandwidth
- Access delay
- Redundancy

Refer to the latest pricing information available on supported Cloud Storage vendor websites for more information and take advantage of the cost calculators that most vendors provide. Please keep in mind that the service pricing can be updated frequently.

Data Stream Size and Stream Handler Options for Cloud Connector

The Data stream size and Stream handler options currently have no effect for Backup Exec Cloud Connector. The Backup Exec Cloud Connector operates with data in 1MB chunks; this size is currently not configurable.



Note: Data Stream Size and Stream Handler Options are removed from Backup Cloud Connector device properties in the Backup Exec 16 FP1 user interface.



Tuning the Cloud Connectors for your environment

You can decrease and tune the Read/Write connections for Cloud-based backup jobs if there are bandwidth limitations between the Backup Exec server and the cloud backup destination. The default Backup Exec Read/Write connections settings assume a typical network bandwidth but in some cases the number of write or read connections may overwhelm the bandwidth available. With Backup Exec 16 FP1, the Read/Write connection settings are available in the user interface. If network bandwidth decreases, you can tune the Read/Write connections numbers to establish a consistent connection. You are not required to restart the Backup Exec services. The Read/Write tuning settings work for both public and private cloud Storage devices that are supported by Backup Exec. Below is the Read/Write settings table located in Backup Exec Configuration and Settings | Backup Exec Settings | Cloud Storage. The default settings shown can be decreased by cloud storage type to eliminate connection errors caused by insufficient bandwidth. Note that decreasing the Read/Write connections will cause a corresponding decrease in job throughput.

The screenshot shows the Backup Exec Settings window with the 'Cloud Storage' section selected in the left-hand navigation pane. The main area displays a table with the following data:

| Cloud server type | Connection range | Read connections | Write connections |
|-------------------|------------------|------------------|-------------------|
| Amazon China | 1-100 | 100 | 100 |
| Amazon | 1-100 | 100 | 100 |
| Azure China | 1-25 | 25 | 25 |
| Azure | 1-25 | 25 | 25 |
| Google | 1-100 | 100 | 100 |

Below the table, an information icon (i) is followed by the text: "The table displays the read and write connection values (range and selected value) for each type of cloud storage server that is supported by Backup Exec. Veritas recommends that you change the settings only during network or bandwidth issues for backup or restore jobs running on cloud devices. The default values are set as per the best practices suggested by Backup Exec. For more information on configuration of these values refer to the technote: <https://www.veritas.com/docs/000125021>"

Miscellaneous Tips

- To back up to disk and then duplicate to the cloud, you must configure two types of storage in Backup Exec; a local disk storage device to stage data from the local storage to the cloud, and the cloud storage.
- Create specific buckets/containers to use exclusively with Backup Exec.

- Use a different bucket/container for each cloud storage device. Do not use the same bucket for multiple cloud storage devices, even if these devices are configured on different Backup Exec servers.
- Ensure that bucket/container names contain only lowercase letters, numbers, and dashes or hyphens. Also, ensure that bucket names do not begin with a dash. Buckets are not available for use in Backup Exec if the bucket name does not comply with the bucket naming conventions.