Quest® NetVault® Backup Plug-in *for Exchange* 12.3

User's Guide

Contents

Introducing NetVault Backup Plug-in *for Exchange*

- NetVault Backup Plug-in for Exchange: at a glance
- Key benefits
- Feature summary
- Target audience
- Recommended additional reading

NetVault Backup Plug-in *for Exchange*: at a glance

Quest® NetVault® Backup Plug-in *for Exchange* (Plug-in *for Exchange*) increases confidence in the recoverability of Exchange. It also lets you create flexible backup policies that can account for multiple recovery scenarios. Support for online backups through Exchange Server's Volume Shadow Copy Service (VSS) lets you select your preferred backup method without requiring you to learn Exchange internals. Through a web-based user interface (WebUI) and automated workflow process, the plug-in offers a centralized way to set up, configure, and define backup and restore policies for your Exchange Servers. These policies include policies deployed in a Database Availability Group (DAG). The plug-in offers a detailed level of control that minimizes downtime by allowing you to restore complete Information Stores or individual databases. Through integration with a range of backup devices, your data is protected and stored offsite to meet your disaster-recovery and business-continuity goals.

Key benefits

- **Increases confidence and reduces risk while deploying Exchange**: Plug-in *for Exchange* lets you create backup policies that are flexible enough to account for many recovery scenarios. You do not have to understand Exchange internals.

  Plug-in *for Exchange*'s flexible backup features also include:

  - Protection for standalone and DAG environments
  - Online VSS-based backups
  - Full, Incremental, and Differential Backups while data is online and accessible
  - Copy-Only Backups
  - Protection down to the individual database

  By relying on Plug-in *for Exchange* to implement backup policies, you can focus on more critical tasks without risking your ability to recover what is needed if a failure occurs. In addition, the IT manager's confidence is increased by knowing that email is protected, no matter what.

- **Speeds up restores to reduce downtime**: With Plug-in *for Exchange*, you select what must be restored and the backup set to restore from, and the plug-in automatically performs the restore. To maximize availability, the plug-in is designed for detailed recoveries, which lets you recover complete Information Stores or individual databases. The plug-in uses an automated workflow with point-and-click options for tasks such as defining backups and scheduling jobs.

  Additional Plug-in *for Exchange* restore features include:

    - Full, Incremental, and Differential Restores

    - Restores of complete Information Stores or individual databases

    - Rename of Mailbox Databases during restore

    - Restores to Recovery Databases (RDBs)

    - Restores to alternate Exchange Server

- **Ensure business continuity**: With offsite backups being an important part of the data-protection for business-critical applications, the plug-in takes advantage of NetVault Backup's integration with a range of backup devices. NetVault Backup lets you select which backup device to store the backup on. You can store the backup online in a virtual tape library (VTL). You can also duplicate the job to physical tape libraries shared by multiple Exchange Servers, other proprietary databases, or even general backup purposes.

- **Eliminate backup windows and reduce storage**: Plug-in *for Exchange* gives you the confidence that your email environment is protected and stored offsite for disaster-recovery purposes. At the same time, it frees administrators from having to be available 24x7 because less-experienced personnel can initiate restores, thus reducing downtime and improving business continuity.

Feature summary

- Protection for standalone and DAG environments

- Online VSS-based backups

- Full, Incremental, and Differential Backups while data is online and accessible

- Copy-Only Backups

- Protection down to the individual database

- Support for In-Place Archiving

- Full, Incremental, and Differential Restores

- Restores of complete Information Stores and individual databases

- Rename of Mailbox Databases during restore

- Restores to RDBs

- Restores to alternate Exchange Server

- Rename of Database during restores

- Restore to non-Exchange Server

- Point-and-click WebUI

Target audience

This guide is intended for users who are responsible for the backup and recovery of Exchange Server. Familiarity with Exchange Server administration is assumed. Advanced knowledge of Exchange Server is useful for defining an efficient backup-and-recovery strategy.

Recommended additional reading

Quest recommends that you have the following documentation available for reference when setting up and using this plug-in.

- **Exchange Server 2019:** https://docs.microsoft.com/en-us/Exchange/exchange-server?view=exchserver-2019

- **Exchange Server 2016:** https://docs.microsoft.com/en-us/Exchange/exchange-server?view=exchserver-2016

- **Exchange Server 2013:** http://technet.microsoft.com/en-us/library/bb124558.aspx

- **Exchange Server 2010:** http://technet.microsoft.com/en-us/library/bb124558(v=exchg.141).aspx

The following documentation is also available:

- *Quest NetVault Backup Installation Guide*: This guide provides details on installing the NetVault Backup Server and Client software.

- *Quest NetVault Backup Administrator's Guide*: This guide explains how to use NetVault Backup and describes the functionality common to all plug-ins.

- *Quest NetVault Backup CLI Reference Guide*: This guide provides a description of the command-line utilities.

You can download these guides from https://support.quest.com/technical-documents.

Defining an Exchange data protection strategy

- Defining a strategy
- Examples of backup sequences
- Understanding snapshot-based backups and restores

Defining a strategy

The purpose of creating Exchange Server backups is to recover an Exchange Server that is damaged from media failure or data corruption. Reliable use of backup and restore for recovery requires a strategy that maximizes data availability and minimizes data loss, while accounting for defined business requirements.

A backup-and-restore strategy is divided into two pieces: a backup piece and a restore piece.

- The backup piece defines the type and frequency of backups that are required to meet the goals for availability of the Exchange Server and for minimizing data loss.

- The restore piece defines who is responsible for performing restores and how restores should be performed to recover from the particular type of damage or failure.

Protecting Exchange Server databases

A database is the finest granularity of storage organization within the Exchange Server schema. Mailboxes are assigned to specific databases, which allows segregation of data by mailbox for security or scalability purposes. Occasionally, critical or high-volume mailboxes are segregated to a separate database to improve performance or to perform more frequent backups.

Databases are used to contain mailboxes, messages, folder stores, and various other data objects supported by the Exchange Server. While they can contain a large variety of data objects, databases are typically one of two types: mail stores and public folder stores. In Exchange Server terminology, a *store* is the same as a database.

Microsoft also differentiates editions of the Exchange Server product by increasing support for concurrent databases. The Enterprise Edition is defined as supporting more concurrent databases than the Standard Edition.

Protecting the Exchange Server transaction log

All changes made to the Exchange Server database are first committed to transaction log files. Any time a user modifies data stored in a mailbox or data is added to the mailbox, that change is written to a transaction log file before it is written to the Exchange Server database.

Reviewing the Volume Shadow Copy Service (VSS) backup method

Microsoft supports the ability to create snapshots of Exchange data using VSS. Microsoft provides Exchange-specific VSS writers that coordinate with the Exchange Services, operating on behalf of the Plug-in *for Exchange*, to prepare the Mailbox Database files for backups and freeze the input/output (I/O) activity due to Exchange transactions before backing it up, and then to unfreeze and truncate logs after the backup is complete.

Using VSS, you can:

- Use snapshots to perform consistent backups to disk- or tape-based storage devices.

- Create and store snapshots as backups on NetVault Backup-supported disk arrays.

> **NOTE:** The **Backup Files to Storage** option is supported with any disk-based storage. To use the **Perform Backup Off-Host using Client**, **Retain Array-Based Persistent Snapshot**, and **Discard After** options, the data that you back up must reside on a NetVault Backup-supported disk array. Also, for persistent snapshots, only the metadata is copied to the target.

Managing transaction log files

When defining an Exchange Server backup strategy, transaction log file management is a primary consideration.

> **IMPORTANT:** Important: Because all changes made to the Exchange Server database are first committed to transaction log files, Quest strongly recommends that you perform periodic Full or Incremental Backups to prevent excessive accumulation of transaction log files. Failure to perform regular backups, which reduces the number of transaction log files, might eventually cause backup jobs to fail. Also, accumulating thousands of transaction log files negatively impacts the time required to process a backup job and generate the index. If your environment involves frequent changes to the Exchange Server, Quest also recommends increasing the frequency of backups to reduce the accumulation of transaction log files.

Transaction log file truncation

Transaction log file truncation is used to "clean" the Exchange Server Mailbox Database, thus improving performance, and reducing the disk space requirements and the time required to restore a database.

Quest recommends that backups that perform transaction log file truncation be performed regularly. Often, a weekly or semimonthly backup that performs transaction log file truncation is recommended. The optimum frequency might vary considerably, depending on the use and configuration of each Exchange Server.

If you use a **Backup Type** that supports transaction log file truncation, such as the Full and Incremental Backup types, the truncation of transaction log files is performed by the Exchange Server after the plug-in informs it that the backup completed successfully. Also, when the truncation occurs depends on Exchange Server and whether it still needs the logs for more purposes, such as replication; therefore, truncation might not occur immediately after a successful backup is completed.

In a DAG environment, transaction log file truncation is also determined by the replay lag time and truncation lag time properties of the database. The properties are configurable. Replay lag time defines, in minutes, the amount of time to delay log replay for a database copy. Truncation lag time defines, in minutes, the amount of time to delay log deletion for the database copy, after the log file has been replayed into the database copy.

For Exchange Server to truncate a transaction log file, the following criteria must be met:

- Either Circular Logging is enabled, or Circular Logging is disabled and the log file has been included in a successful Full or Incremental Backup.

- The transaction log file sequence is less than the sequence of the last transaction log file committed to the Exchange Server database. The last sequence committed is recorded in the database checkpoint file (.chk).

- In a replication environment, such as a DAG, all other copies of the database have confirmed reception of the transaction log file and replay of its contents to the database file.

In a DAG environment, each database copy keeps transaction log files until all the database copies have confirmed that the transaction log file has been replayed. If one or more passive copies of the database is suspended or offline, log truncation does not occur, resulting in transaction log file buildup and the consumption of disk space.

> **i** | **NOTE:** If your environment experiences transaction log file buildup and your disk space is diminished, check the online status of the passive database copies. If you must keep one or more database copies suspended or offline for a long period of time, such as for maintenance reasons, consider removing the affected passive database copy instead of suspending it or setting it offline. When maintenance is completed, the passive database copy can be added back.

Loose truncation

To reduce the effects of transaction log file buildup due to suspended or offline database copies, Exchange Server 2013 Service Pack 1 introduced **loose truncation**. With loose truncation, each database copy tracks its own available disk space and applies loose truncation when disk space gets significantly low. With loose truncation applied, each passive database copy independently truncates its own transaction log files. For the active database copy, truncation ignores the passive database copy that is farthest behind replaying logs.

Loose truncation is disabled by default. To enable loose truncation, you must edit the Windows registry on each Exchange Server DAG node. Before enabling loose truncation, make sure that it would benefit your data protection goals. For more information about enabling loose truncation, see your Exchange Server documentation.

Full vs. Transaction Logs-Only Backups

Full Backups back up all files for a database, regardless of the type of file. Transaction Logs-Only Backups back up only the transaction log files for a database.

Full Backups allow all database files to be backed up, providing standalone restore capabilities. Depending on the size of the database, Full Backups can be demanding in terms of storage requirements, and time required to complete the backup. For large databases, storage and time requirements might be a significant consideration.

Transaction Logs-Only Backups are lighter-weight backups that are intended to capture new activity since the last Full Backup was performed. This type of backup can radically reduce the backup time and storage requirements for large databases, but it also introduces dependencies on one or more prior backups to perform a complete restore.

> **i** | **IMPORTANT:** VSS-based backups for non-Continuous Replication environments truncate the transaction logs at the completion of Full or Transaction Logs-Only Backups. Microsoft Exchange Replication Service deletes the backed-up log files both from the active and the passive copy log file paths after it verifies that the to-be-deleted log files have been successfully applied to the passive copy database and both the active and passive copy database check points have passed the log files in question.

Reviewing the backup types

The Plug-in *for Exchange* provides several essential backup types:

- **Full Backups**
- **Copy Backups**
- **Incremental Backups**
- **Differential Backups**

Backup types represent the complete set of options for backing up an Exchange Server with respect to the transaction log management rules.

Full Backups

A Full Backup is the most common type of backup performed. It performs a complete backup of a database or Mailbox Database. A Full Backup does not depend on any other backup and can be restored in a single step.

A Full Backup includes all database files and transaction log files. After a Full Backup is successfully written to a storage device, the Plug-in *for Exchange* informs the Exchange Server that the backup was successful. At that point, the Exchange Server can truncate the transaction log files if you have configured Exchange Server to do so. To keep the Exchange Server in good health and maintain its performance, perform regular Full Backups.

Copy Backups

Sometimes a backup must be taken for a special purpose and should not affect the overall backup and restore procedures for the Exchange Information Store. A Copy Backup is used to back up all database files without performing transaction log file truncation. This backup provides a snapshot of the files as they existed at the time of the backup. A Copy Backup should be used to perform a fast, nonintrusive complete backup of an Exchange Server and is common in scaling and migration scenarios.

Because a Copy Backup does not perform transaction log file truncation, the disk footprint of the Exchange Server does not change as a result of the backup. Therefore, a Copy Backup does not perform any housekeeping measures, and therefore does not maintain the performance of the Exchange Server. Copy Backups should not be considered as a part of a regular backup policy, but as a special-purpose type of backup.

Incremental Backups

Incremental Backups back up the transaction log files that capture the changes made to the database since the most recent Full or Incremental Backup. After the Plug-in *for Exchange* informs the Exchange Server that the backup was successful, the Exchange Server then truncates the transaction logs, which helps with maintenance of the Mailbox Databases. Incremental Backups are fast and correspondingly small.

As with a Transaction Logs-Only Backup, an Incremental Backup must always have a base backup, which is often a Full Backup. Performing an Incremental Backup without first performing a Full Backup might lead to backup failures and prevent you from being able to restore and recover your Exchange Server properly.

Do *not* perform an Incremental Backup in the following conditions:

- **When an Initial Full Backup has not been created**: The transaction log only contains changes made to the database after the last backup was created, which means that a base backup is required.

- **After an Incremental Backup containing multiple Mailbox Databases fails**: In this case, a Full Backup must first be performed before an Incremental Backup can be performed. When an Incremental Backup that contains multiple Mailbox Databases fails, some of the transaction logs are truncated and permanently lost. At that point, the transactions have already been deleted from the Exchange Server logs. Although the data still exists within the Exchange Server, errors would occur if you attempted to restore Incremental Backups that were created after a failed Incremental Backup job. This issue does not apply to a failed Incremental Backup that contains only a single Mailbox Database.

- **After the Transaction Log is manually truncated**: In this case, a Full or Differential Backup must be performed before an Incremental Backup. Microsoft strongly recommends against manually truncating transaction logs.

Differential Backups

Differential Backups back up the transaction log files that capture the changes made since the most recent Full or Incremental Backup. Differential Backups do not perform transaction log truncation, and therefore do not maintain the operating health of the Exchange Server.

Subsequent Differential Backups increase in size and in duration because each Differential Backup includes the transaction log files that were also included in the previous Differential Backup, and the transaction log files that have been generated since the previous Differential Backup. For example, if a Full Backup was taken on Sunday with Differential Backups scheduled for Monday through Saturday, Monday's Differential includes the transaction

log files generated since the Full Backup on Sunday, while Tuesday's Differential includes the transaction log files generated on Monday and those files generated on Tuesday. Wednesday's Differential includes the transaction log files for Monday, Tuesday, and Wednesday, and so on.

As with a Transaction Logs-Only Backup, a Differential Backup must always have a base backup, which is often a Full Backup. Performing a Differential Backup without first performing a Full Backup might lead to backup failures and prevent you from being able to restore and recover your Exchange Server properly.

When used with Full Backups, Differential Backups provide acceptable trade-offs with good backup performance, good restore performance, and reasonable overall database maintenance through the Full Backup.

Incremental vs. Differential Backups

Because Exchange Server truncates transaction log files after they are backed up in Incremental Backups, subsequent Incremental Backups are quicker. They are quicker because only the transaction log files that have been created since the last Incremental Backup are backed up. However, restore sequences that use Incremental Backups require that every Incremental taken between the Full Backup and the point of failure must be restored in succession. This process can result in longer restores due to the increased human intervention required to initiate the multiple restore jobs.

Differential Backups do not truncate the transaction log files after they are backed up. Subsequent Differential Backups take increasing longer because all the transaction log files since the last Full Backup are included in the backup. Nevertheless, restore sequences that use Differential Backups require that only one Differential Backup be restored after the Full Backup is restored. This process results in quicker restores because less human intervention is required during the restore process.

An extra consideration when deciding between Incremental and Differential Backups is the desired frequency of transaction log file truncation. When implementing Incremental Backups, the transaction log files are truncated by the Exchange Server at the same frequency as the Incremental Backups. For example, if Incremental Backups are performed daily, the transaction log files are truncated daily. However, with Differential Backups, the transaction log files are truncated only when a Full Backup is performed. Therefore, if the Full Backup is only performed weekly, the transaction log files are only truncated weekly.

Implementing a Differential Backup strategy leads to faster restores but also requires more frequent Full Backups to maintain the operating health of the Exchange Server.

Examples of backup sequences

- **Full Backups only:** When requirements guarantee backup protection up to the previous day, performing Full Backups nightly should be sufficient under the following conditions:

  - Backup windows are large.

  - After-hours email volume is low.

  - Regular transaction log file truncation is required.

- **Full and Incremental Backups:** When requirements guarantee data protection up to the previous day, *backup time should be as fast as possible*, and regular transaction log file truncation is required, Full Backups coupled with Incremental Backups is the best combination.

  For example, Full Backups are performed every Sunday night at 11:00 p.m. while Incremental Backups are performed Monday through Saturday at 11:00 p.m. Each Incremental Backup includes the transaction log files that were generated since the previous night's backup whether it be the Sunday evening Full Backup or one of the Incremental Backups from a weekday.

  Restoring this type of backup sequence is more time-consuming. For example, if recovery is performed on Tuesday, only Sunday's Full Backup and Monday's Incremental Backup must be restored; whereas, if recovery is performed on Wednesday, Sunday's Full Backup followed by Monday's and Tuesday's Incremental Backups must be restored. Even though the backups are quicker, the restores can take longer due to the intervention that is required to run multiple restore jobs.

- **Full and Differential Backups:** When requirements guarantee data protection up to the previous day, *restore and backup time should be reasonably fast*, and only occasional transaction log file truncation is required, Full Backups coupled with Differential Backups is the best combination.

  For example, Full Backups are performed every Sunday night at 11:00 p.m. while Differential Backups are performed Monday through Saturday at 11:00 p.m. Each Differential Backup includes the transaction log files that were generated since the last Full Backup. This process might require more backup time than an Incremental Backup. Regardless of the point to which recovery is necessary, the same number of restore jobs is required. For example, if recovery is performed on Tuesday, Sunday's Full Backup and Monday's Differential Backup must be restored; whereas, if recovery is performed on Thursday, Sunday's Full Backup followed by Wednesday's Differential Backup must be restored. Even though subsequent Differential Backups increase not only in size but in duration, restores are quicker because fewer restore jobs must be run.

Understanding snapshot-based backups and restores

The plug-in can use a hardware or software VSS provider to create persistent or non-persistent VSS-based snapshots. The plug-in uses a VSS provider to create snapshots on the client, and then copies the selected data from the snapshot, or snapshots, to a storage device.

You can use hardware-based, integrated VSS snapshots with the Dell Compellent storage array. Be aware that if you include a mixture of databases, Compellent and non-Compellent, in the same backup job, the plug-in backs up all included databases to the same storage device.

To create and use persistent snapshots, the data that you want to back up must reside on NetVault Backup-supported disk arrays. If you attempt to create persistent snapshots with data residing on an unsupported disk array, or with data residing in local drives (for which snapshots are taken using the Microsoft Software VSS Provider), taking the required hardware snapshots fails, and the plug-in defaults to using software snapshots. When taking software snapshots the data is only backed up to storage (even if the option **Backup Files to Storage** is not selected), and non-persistent snapshots are created.

The same issue occurs if your backup includes data that resides in different storage that mixes NetVault Backup-supported disk arrays and unsupported disk arrays or local drives. For example, if your backup includes five Exchange Mailbox Databases, for which three of the Mailbox Databases reside in different volumes in a NetVault Backup-supported disk array, and the other two Mailbox Databases reside in two different local drives on the Exchange Server, and you select **Retain Array-based Persistent Snapshot**, taking hardware snapshots fails because all the Mailbox Databases do not reside on NetVault Backup-supported disk arrays. After failing to take the required hardware snapshots, the plug-in attempts to take software snapshots; if successful, the data is saved to NetVault Backup-managed storage, even if the option **Backup Files to Storage** was not selected.

> **i** | **IMPORTANT:** If you are running a DAG environment, ensure that the software-based Microsoft VSS provider is in place on all Exchange Servers included in the DAG and ensure that the software-based provider has preference over any vendor's hardware-based VSS provider.

If your backup target is non-NetVault Backup media, such as a disk array, and you need to complete a restore, be aware of the following: If you restore an Exchange database to a DAG and do not specify a Target Client, the plug-in restores the database to the node that hosts the active copy. If you restore the database from a persistent snapshot, the plug-in maps the volume in the client that took the snapshot during the backup process. Because of these conditions, restoring from a persistent snapshot to a DAG, that is, where the Target Client is a NetVault Backup Virtual Client, requires that the active copy of the database is hosted on the node that takes the snapshots during the backup process.

Example: If **node1** takes a snapshot while **DB1** is being backed up, restoring **DB1** from the snapshot requires that **node1** is the host of the active copy of DB1. Otherwise, the restore fails. For example, if node2 is currently hosting the active copy of DB1, you must move the active role from **node2** back to **node1**.

To change the active-role assignment, use the following PowerShell command:

```
Move-ActiveMailboxDatabase <DatabaseName> -ActivateOnServer <ServerName>
   -MountDialOverride:None
```

If the damage to the database causes validation failures, use the `-SkipClientExperienceChecks` switch to force activation of the database:

```
Move-ActiveMailboxDatabase <DatabaseName> -ActivateOnServer <ServerName>
   -MountDialOverride:None -SkipClientExperienceChecks
```

For more information about NetVault Backup-supported OS versions and disk arrays, see the *Quest NetVault Backup Compatibility Guide*.

Planning your Exchange Server deployment

- Deployment overview
- Standalone deployment
- High-availability deployment using a Single- Location Database Availability Group (DAG)

Deployment overview

Microsoft supports the deployment of Exchange Mailbox Server on either a single server or high-availability DAG environment.

Deploying the Plug-in *for Exchange* in either of these environments is almost identical because the Plug-in *for Exchange* is installed on the server that hosts the Exchange Server Mailbox Databases. The following topics describe how the Plug-in *for Exchange* is deployed in each type of Exchange Mailbox Server deployment.

Standalone deployment

- **Supported Exchange Server Versions**: Exchange Server 2010 and later
- **Supported Exchange Server Editions**: All

In a standalone deployment, there is a single Exchange Mailbox Server. The Plug-in *for Exchange* is installed on the server that is designated the Mailbox Server and performs all backups and restores to and from this server. In this type of deployment, no high-availability solution has been deployed.

While you can install the NetVault Backup Server on the Exchange Mailbox Server, Quest recommends that you use separate machines.

Figure 1. Standalone deployment

- NetVault Backup Server software
- Plug-in for Exchange
- NetVault Backup Client software
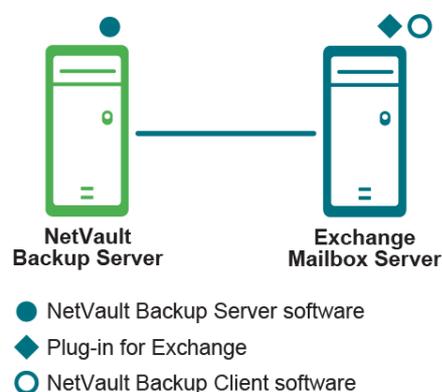
> **i** | **IMPORTANT:** Regardless of the environment, a separate NetVault Backup Server and Exchange Server vs. a single machine configured as both, the Plug-in *for Exchange must* be installed on the host where Exchange Server Mailbox Databases reside. Sample images and procedures throughout this guide assume that this ***two-machine environment*** is in place and that all configuration requirements have been met.

High-availability deployment using a Single-Location Database Availability Group (DAG)

- **Supported Exchange Server Versions**: Exchange Server 2010 and later

- **Supported Exchange Server Editions**: All

- **Plug-in *for Exchange* License Required**: Plug-in *for Exchange* Standard Edition Cluster Support or Plug-in *for Exchange* Enterprise Edition Cluster Support

In an Exchange Single-Location DAG, multiple Exchange Mailbox Servers exist in a single Active Directory Site and are physically located in the same data center. A DAG is created across multiple Mailbox Servers with the Database Copies spread across the multiple Mailbox Servers. In a Single-Location DAG, you create a NetVault Backup Virtual Client where the Plug-in *for Exchange* is installed on each of the Exchange Mailbox Servers.

In the backup process for a DAG environment, the Plug-in *for Exchange* lets you select either all active copies of the databases or a database copy whenever possible. If you choose the latter and there are multiple copies, the copy with the lowest activation preference number is selected. If you choose the latter and a database copy is not available, the active database is selected. For more information about the activation preference number, see *Mailbox Database Copies* at http://technet.microsoft.com/en-us/library/dd979802.aspx.

> **i** | **IMPORTANT:** To support a DAG that does *not* use a cluster name object (CNO) or IP address, create a NetVault Backup Virtual Client that uses a fixed IP address or network name that corresponds to one of the Real Clients in the DAG. If the machine to which the fixed IP address or network name corresponds goes offline for any reason, subsequent backup jobs fail. Either bring the machine back online or update the Virtual Client to use a different IP address or network name.

Figure 2. DAG deployment
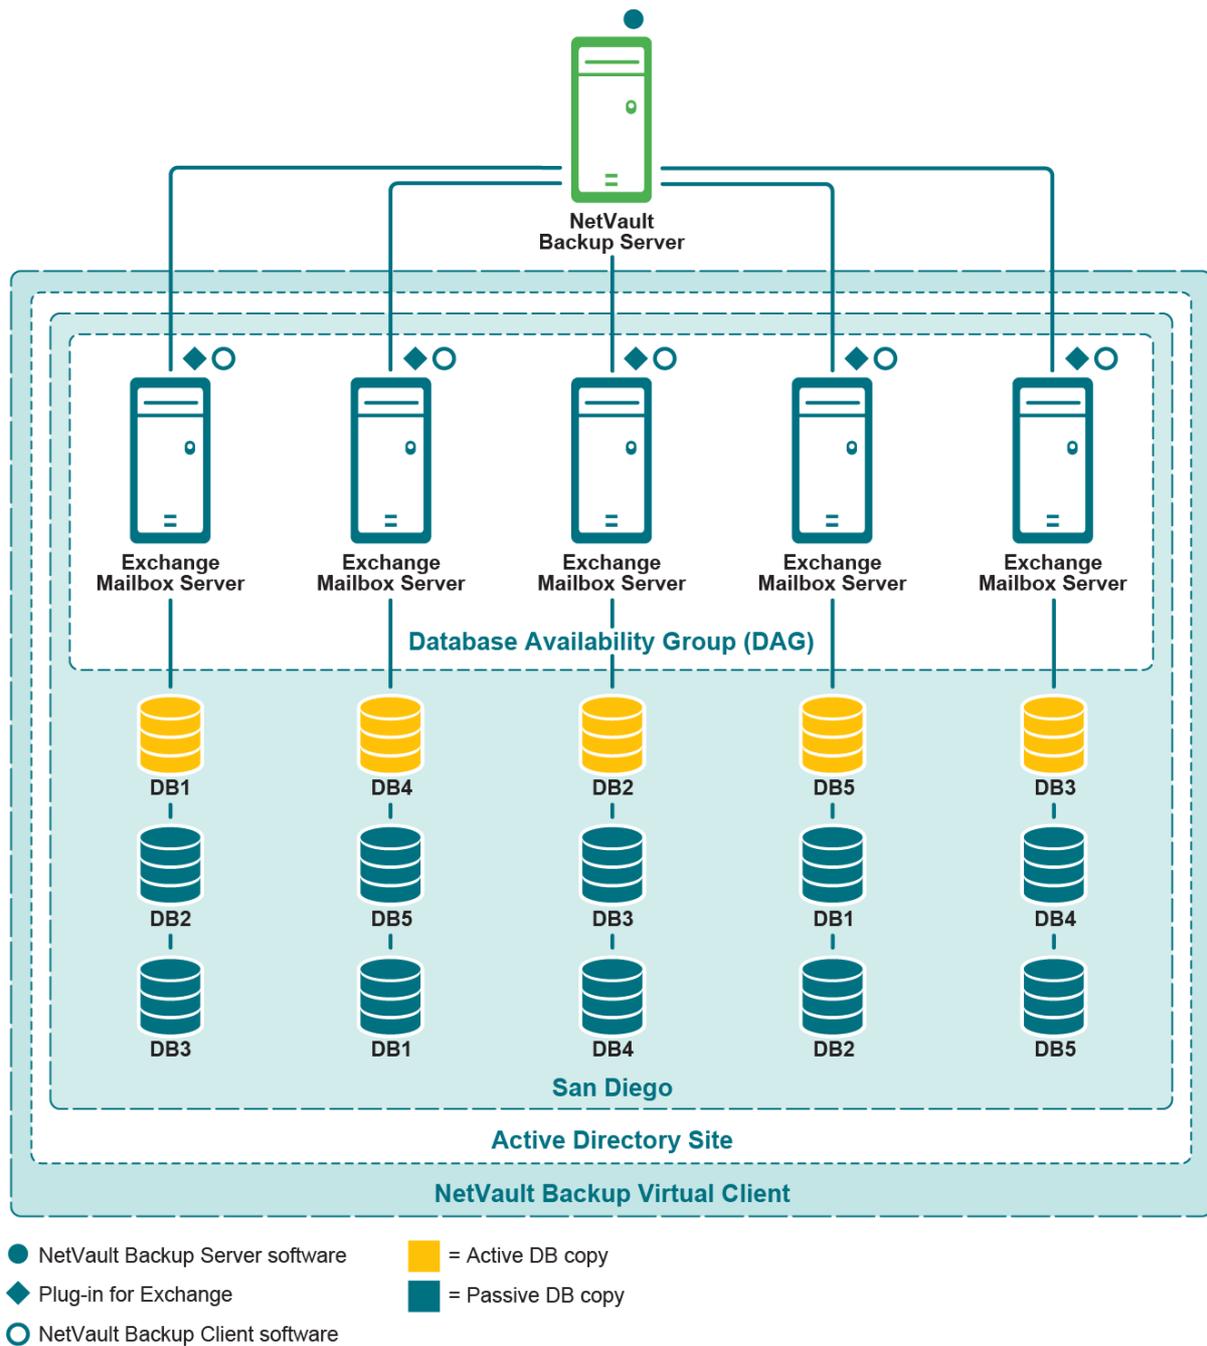
Installing and removing the plug-in

* Installation prerequisites
* Installing or upgrading the plug-in in standalone deployments
* Installing or upgrading the plug-in in high-availability deployments
* Removing the plug-in from a standalone deployment

Installation prerequisites

Before installing the plug-in, ensure that the requirements outlined in the following topics are met on the machine that is to serve as the Exchange Server.

* Disabling Circular Logging
* Enabling services
* Configuring localization settings
* Enabling and using Deleted Item Recovery features
* Additional requirements for DAG deployments

Disabling Circular Logging

If either an Incremental or Differential Backup type is to be created, ensure that **Circular Logging** is disabled. The Exchange Server overwrites files if **Circular Logging** is enabled, making it impossible to reliably restore from the transaction logs. For more information about **Backup Types**, see Reviewing the backup types. To disable **Circular Logging**, perform the steps in the applicable topic:

* Disabling Circular Logging in Exchange 2010
* Disabling Circular Logging in Exchange 2013 and later

Disabling Circular Logging in Exchange 2010

1 Open the **Exchange Management Console**.

2 Navigate to **Organization Configuration > Mailbox**.

3 On the **Database Management** tab, select the applicable database, and click **Properties**.

4 On the **Properties** dialog box, select the **Maintenance** tab, clear the **Enable circular logging** check box, and click **OK**.

5 Close the **Exchange Management Console**, and restart the **Microsoft Exchange Information Store** service for the change to take effect.

For complete instructions, see the relevant Microsoft Exchange documentation. In addition, see http://technet.microsoft.com/en-us/library/dd297937.aspx and http://technet.microsoft.com/en-us/library/dd351151(v=exchg.150).aspx.

Disabling Circular Logging in Exchange 2013 and later

By default, Circular Logging is disabled in Exchange 2013 and later. If you have enabled it, disable it before using the plug-in.

1. In a web browser, open the Exchange Administration Center (EAC) by accessing the Exchange Control Panel (ECP) URL.

2. Log in using an Administrator account.

3. In the feature pane on the left, click **Servers**.

4. In the tab at the top, click **Databases**.

5. Select the applicable database, and click the **Edit** button.

6. In the new browser window, click **maintenance**.

7. Clear the **Enable circular logging** check box, and click **save**.

For complete instructions, see the relevant Microsoft Exchange documentation.

Enabling services

In **Windows Control Panel > Administrative > Services** for the Exchange Mailbox Servers, verify that the following services are enabled and running:

- **Microsoft Exchange Information Store**

- **Microsoft Exchange System Attendant**

- **Microsoft Exchange Replication Service**—required in DAG environments

For VSS backups, also verify that the following services are enabled and running:

- **Microsoft Software Shadow Copy Provider**

- **Volume Shadow Copy**—started automatically by Microsoft Software Shadow Copy Provider

Configuring localization settings

The plug-in uses traditional Windows code page selection to enable localization, and is therefore subject to certain limitations. The plug-in must be used in an environment with correctly configured localizations settings to function properly.

To ensure correct operation in a localized environment, verify that the following conditions are met:

- All Exchange Server entities—Mailbox Databases, Public Folder Databases, and so on—must be named using the same language, with no intermixing of languages.

- The Windows Active Code Page must be set to the same language as used in the Exchange Server.

- The language specified in the NetVault Backup Configurator for both the NetVault Backup Server and Clients must match the language set in the Windows Active Code Page.

- If used in a clustered environment, all nodes in the cluster must be configured to use the same language, also conforming to the previous rules.

Configuring the Windows Active Code Page

1. Open the Windows **Control Panel**.

2. Select **Regional and Language Options**.

3. Select the **Advanced** tab.

4 In the **Select a language to match the language version of the non-Unicode programs you want to use** list, select the correct language.

If the correct language is not listed, install it by following the Microsoft Windows administrative procedures for installing more language support.

Configuring a NetVault Backup Server or Client

1 In the Navigation pane, click **Change Settings**.

2 On the Configuration page, click **Server Settings** or **Client Settings**, whichever is applicable.

3 Under **System and Security**, click **General**.

4 In the **Language Selection** list, select the applicable language, and then click **Apply**.

> **i** | **IMPORTANT:** When languages other than English are used in an environment that is not correctly configured, backup and restore operations may not function as expected. The demonstrated behavior may include, but is not limited to:
>
> • Failure to back up selected items.
>
> • Backing up or restoring items that were not selected, for example, you select the first Mailbox Database, but the second Mailbox Database is affected.
>
> • Failure to back up or restore altogether, often with the indication that no matching items were found.

Enabling and using Deleted Item Recovery features

Quest recommends that you enable the **Deleted Item Recovery** functionality in Exchange to reduce the requirement to perform recovery for individual messages or mailboxes. This feature enables the recovery of recently deleted items by saving messages for a predetermined number of days.

• Enabling Deleted Item Recovery in Exchange

• Recovering items in Outlook 2010 or later

Enabling Deleted Item Recovery in Exchange

1 Open the **Exchange Management Shell**.

2 To enable recoverable items and specify the number of days deleted items are retained, enter the following command:

```
Set-Mailbox -Identity <MailboxServer> -SingleItemRecoveryEnabled
    $True -RetainDeletedItemsFor <dd.hh:mm:ss>
```

For *<MailboxServer>*, you can use the ADObjectID, alias, distinguished name (DN), domain\account, GUID, LegacyExchangeDN, SMTPAddress, or user principal name (UPN). For the length of time, dd = days, hh = hours, mm = minutes, and ss = seconds.

3 To configure single-item recovery and configure recoverable items for a **mailbox**, enter the following command:

```
Set-Mailbox -Identity <MailboxServer> -RecoverableItemsQuota <Limit>
```

For *<Limit>*, enter the maximum amount of space, for example, 15 GB, that can be used before more items cannot be stored in the Recovery Items Folder.

4 To configure single-item recovery and configure recoverable items for a **Mailbox Database**, enter the following command:

```
Set-MailboxDatabase -Identity <MailboxServer> -RecoverableItemsQuota <Limit>
```

Instructions for these procedures can be found in the relevant Microsoft Exchange documentation: http://technet.microsoft.com/en-us/library/ee364752.aspx. In addition, see http://technet.microsoft.com/en-us/library/ee364755.aspx.

Recovering items in Outlook 2010 or later

1   Select the applicable folder, such as the **Inbox** or a folder that you created.

2   Select the **Folder** tab.

3   On the **Ribbon**, click **Recover Deleted Items**, and select the items from the list presented.

4   To have the items restored to your Deleted Items folder, select **Recover Selected Items**.

Complete instructions for this procedure can be found in the relevant Microsoft documentation. In addition, see http://office.microsoft.com/en-us/outlook-help/recover-deleted-items-HA010355039.aspx.

Additional requirements for DAG deployments

Verify that the following prerequisites are met before you install the Plug-in *for Exchange* in a DAG deployment:

- **Microsoft Exchange DAG environment in place**: You must have a properly configured DAG environment.

- **Separate NetVault Backup Server machine**: The machine that is to serve as the NetVault Backup Server must be properly configured. It *must exist outside* the Exchange Server DAG deployment, and it must have network connectivity to the nodes and hosts within the cluster.

- **NetVault Backup Client name must match host name of Exchange Server in a DAG environment**: Verify that the name specified for the NetVault Backup Client matches the host name of the Exchange Server.

- **Review using NetVault Backup in cluster setup details**: Quest strongly recommends that you review the topic dedicated to cluster functionality in the *Quest NetVault Backup Administrator's Guide* to ensure that you understand how the information presented in the following topics works with Exchange Server DAG functionality.

Installing or upgrading the plug-in in standalone deployments

You can install or upgrade the plug-in on individual systems, one at a time. You can also deploy the plug-in to multiple machines by creating a deployment task from the WebUI to push the packages to the specified machines. You can use this method for both new and upgrade installations on Windows-based machines. After the packages are installed on the machines, the new machines are automatically added to the NetVault Backup Server.

- Performing a local installation
- Performing a remote installation

Performing a local installation

1   In the Navigation pane, click **Manage Clients**.

2   On the **Manage Clients** page, select the machine that contains the Exchange Server, and click **Manage**.

3   On the **View Client** page, click the **Install Plugin** button ().

4   Click **Choose Plug-in File**, navigate to the location of the **".npk"** installation file for the plug-in, for example, on the installation CD or the directory to which the file was downloaded from the website.

Based on the OS in use, the path for this software may vary on the installation CD.

5   Select the file entitled **"exs-x-x-x-x-x.npk,"** where **xxxxx** represents the version number and platform, and click **Open**.

6   To begin installation, click **Install Plugin**.

After the plug-in is successfully installed, a message is displayed.

Performing a remote installation

This process is not plug-in-specific. For more information, see the topic on deploying clients and plug-in packages in the *Quest NetVault Backup Administrator's Guide*.

1   In the Navigation pane, click **Guided Configuration**.

2   On the **NetVault Configuration Wizard** page, click **Install Software/Add Clients**.

3   On the **Select Software/Add Clients** page, select **Install NetVault software on remote machines**.

4   In the **Package Store** list, select the applicable repository.

5   Click **Add NetVault plug-in package**.

6   In the **Select packages for deployment** dialog, select the file entitled **"exs-x-x-x-x-x.npk,"** where **xxxxx** represents the version number and platform, and click **Ok**.

The plug-in file is listed in the deployment table.

7   Click **Next**.

8   On the **Machines to Have NetVault Software Installed** page, select the applicable machines.

9   If you are installing an upgrade, select **Allow existing NetVault client installations to be upgraded**.

10 Click **Install Software/Add Clients**.

The status is displayed on the next page.

11 After the plug-in is successfully installed, click **Next** to complete the process.

Installing or upgrading the plug-in in high-availability deployments

Installation of Plug-in *for Exchange* for a clustered environment is different from standalone deployment. This process is completed through the creation of a Virtual Client on the NetVault Backup Server. A Virtual Client is a group of nodes and hosts within the cluster. The NetVault Backup Server sees the cluster as a single client that is created to back up a single clustered resource, that is, an Exchange Server Virtual Server. During the Virtual Client creation process, the plug-in is transferred from the NetVault Backup Server to selected nodes within a cluster and installed on each of the selected nodes.

As described in High-availability deployment using a Single- Location Database Availability Group (DAG), creating the Virtual Client for managing the DAG installs the Plug-in *for Exchange* on each node included in the cluster. During this process, verify that the NetVault Backup Client is installed on each node and that each NetVault Backup Client is added to the Virtual Client that you create for the DAG. This process ensures that the applicable databases are not excluded from the backup process.

Creating a Virtual Client

When you use the WebUI to add a Virtual Client to a NetVault Backup Server, you specify a Virtual Client Address. NetVault Backup Server uses this address to determine which NetVault Backup Real Client (Exchange Server) to contact to initiate the backup job.

The Virtual Client creation process is not plug-in-specific, and you can find complete details in the topic dedicated to cluster functionality in the *Quest NetVault Backup Administrator's Guide*. However, the process is different if you are using any type of high-availability deployment that has an administrative access point versus a DAG deployment that does *not* use an access point.

- Creating a Virtual Client for a deployment that uses an administrative access point
- Creating a Virtual Client for a DAG that does not use an administrative access point

Creating a Virtual Client for a deployment that uses an administrative access point

Consider the following points during the Virtual Client creation process for a deployment that uses an access point:

- **Only include relevant cluster nodes in the Virtual Client**: The hosts that are to be included in the creation of a Virtual Client should *only* be those nodes or Exchange Mailbox Servers that make up the DAG deployment. During the creation of the Virtual Client, you can specify either the IP address of the Exchange virtual cluster or the name of the Exchange virtual cluster. After the creation of the Virtual Client, the plug-in is transferred to all designated cluster nodes and installed locally. The installed Plug-in *for Exchange* can be used through the Virtual Client to back up and restore the DAG deployment.

  > **ℹ️ IMPORTANT:** If you are creating a Virtual Client for a DAG configuration and the DAG is configured to use multiple IP addresses, enter either the network name for the DAG or an IP address that allows network access to the DAG from the machine on which the NetVault Backup Server is installed.

- **Assign a name to the Virtual Client**: Quest strongly recommends that the Virtual Server network name assigned to the Exchange Server be used as the NetVault Backup Virtual Client name. When a Virtual Client is browsed, NetVault Backup locates the node currently in control of the clustered application, and reveals the Exchange Server Instance. With a Virtual Client name set up as the Exchange Server Virtual Server network name, it is easier to recognize the Exchange Server Instance for which the Virtual Client was created.

Creating a Virtual Client for a DAG that does not use an administrative access point

For a high-availability deployment with an administrative access point, the Virtual Client Address that you enter is either a virtual IP address or the network name of the cluster. However, for a DAG without an administrative access point, also known as IP-less DAG or no-CNO DAG, you must create a NetVault Backup Virtual Client that uses a fixed IP address or network name that corresponds to one of the Real Clients (Exchange Servers) in the DAG.

Creating a Virtual Client for a DAG without an administrative access point differs slightly from creating a Virtual Client for a DAG that has an administrative access point. If the DAG does not have an administrative access point, enter the IP address or network name of one of the Exchange Servers that participates in the DAG and on which NetVault Backup Client software is installed. In this case, the Virtual Client Address always resolves to the selected Exchange Server (Real Client) instead of to a DAG.

> **ℹ️ IMPORTANT:** If the Exchange Server (Real Client) to which the fixed IP address or network name corresponds goes offline for any reason, subsequent backup jobs fail. Either bring the Exchange Server back online or update the Virtual Client to use a different IP address or network name for a different Exchange Server (Real Client) that is included in the DAG.

Removing the plug-in from a standalone deployment

For details on removing the Plug-in *for Exchange* in a DAG deployment, see the applicable topic on working with client clusters in the *Quest NetVault Backup Administrator's Guide*.

1   In the Navigation pane, click **Manage Clients**.

2   On the **Manage Clients** page, select the applicable client, and click **Manage**.

3   In the **Installed Software** table on the **View Client** page, select **Plug-in *for Exchange***, and click the **Remove Plugin** button ( ).

4   In the **Confirm** dialog box, click **Remove**.

Configuring the plug-in

- Reviewing Authentication Details: Windows user account permissions
- Configuring the plug-in

Reviewing Authentication Details: Windows user account permissions

Before initiating a backup, configure the plug-in's authentication details and select the Backup Method.

The Plug-in *for Exchange* uses Windows Authentication to provide essential security and access control for sensitive data contained in an Exchange Server's databases. Windows Authentication allows you to log in to a specific Windows user account. You can configure this Windows user account to restrict a user's privileges so that only backups can be performed, or to allow both backups and restores to be performed.

For a user to **back up** Exchange, the Windows user account used must:

- Be a valid domain or local account.
- Be a member of the Power Users group.
- Be a member of the domain's Backup Operators group—for domain accounts.
- Be a member of the Backup Operators group on the machine that Exchange is running on.

To perform a **restore** of Exchange, the Windows user account must be a member of the Administrators group on the machine that Exchange is running on.

The Domain Administrator account has all the permissions necessary to perform backup and restore procedures in the Plug-in *for Exchange* by default.

> **i** | **NOTE:** The **NetVault Process Manager** Service **Log on as** account must be a Local System or an Exchange Administrator Account.

Configuring the plug-in

The plug-in lets you set default options for backup and restore jobs. You can override these options on a per-job basis.

1 In the Navigation Pane, click **Create Backup Job**, and click **Create New** next to the **Selections** list.

2 In the selection tree, open the applicable client node.

3 Click **Plug-in *for Exchange***, and select **Configure** from the context menu.

   The **Exchange Server Configuration** dialog box is displayed. If the selected node is a Virtual Client for a DAG, the **Database Availability Group Name** field displays the name of the DAG. The name of the DAG is followed by a check box that lets you enable editing of the DAG name. For all other nodes, the dialog box shows a **Server Name** field, which displays the machine name of the Exchange Server and cannot be edited.

> **i** | **IMPORTANT:** Ensure that the name in the **Database Availability Group Name** field is the name of the DAG that you want to protect. For a DAG, if the cluster resource **Cluster Name** cannot be accessed or is offline, the plug-in does not retrieve the DAG name and the **Database Availability Group Name** field is blank. To enter a name for the DAG, select the **Edit DAG name** check box, as described in the next step.

4   If you are using a Virtual Client for a DAG that does *not* use a CNO or IP address, do the following:

    a   Select the **Edit DAG name** check box.

    b   In the **Database Availability Group Name** field, enter the name of the DAG.

       The name that NetVault Backup initially provides is tied to a specific machine in the DAG, which is not usually the name of the DAG itself. To prevent backup and restore jobs from failing, enter the correct name of the DAG.

5   In the **Authentication Details** section, edit the following fields:

- **Exchange Administrator User Name**: Enter either a local or domain Windows account that meets the criteria specified in Reviewing Authentication Details: Windows user account permissions. When the **Exchange Server Configuration** dialog box is opened, this account is displayed by default in the Exchange Administrator User Name field.

- **Password**: Enter the password associated with the specified user name. For security reasons, this field is displayed blank by default.

- **Windows Domain**: If a domain account is specified in the **Exchange Administrator User Name** field, enter the name of the domain. The field can be left blank if a local system account is specified. When the **Exchange Server Configuration** dialog box is opened, this domain is displayed by default in the Windows Domain field.

6   In the **Additional Options** section, select the default action for **Incomplete Backup of All Items Selected**.

When multiple items are included in a backup and the plug-in is unable to back up all the selected items, the plug-in lets you specify what action the backup should take. For example, if a job includes multiple Mailbox Databases and the backup of all but one is successful, you can specify what action the backup job should take.

- **Complete with Warnings — Saveset Retained:** The job returns a status of **"Backup Completed with warnings"** and a backup saveset is created that includes the items that were successfully backed up.

- **Complete without Warnings — Saveset Retained:** The job completes and returns a status of **"Backup Completed."** The errors are logged in the NetVault Backup binary logs and ignored on the **Job Status** page. A backup saveset is created that includes the items that were backed up.

- **Fail — Saveset Retained:** The job returns a status of **"Backup Failed."** However, a backup saveset is generated that includes the items that were successfully backed up.

- **Fail — No Saveset Retained:** The job returns a status of **"Backup Failed"** and no saveset of backed-up objects is kept. That is, even if some of the objects were successfully backed up, the saveset is discarded.

> **i** | **IMPORTANT:** You can override the selected default action selected at the individual backup job level.

7   To specify how frequently the nodes in a DAG or cluster send updates to the plug-in, enter the applicable number in the **Monitor Progress Update Interval (in seconds, 0 = continuous)** field.

Depending on the number of nodes residing in a group or cluster and the amount of data processed during a backup, a significant number of messages are sent to the plug-in. Use this option to limit how frequently messages are recorded.

8   If you must instruct the plug-in to ignore warnings resulting from the Exchange integrity check, select the **Ignore Warning when Consistency Check does not Return Operation Results (Not Recommended)** check box.

CAUTION: By default, the plug-in returns a warning if results to the consistency check of the Exchange database are not returned. This issue causes backup jobs to Complete with Warnings. If you know that the consistency check cannot return results because of how your Exchange environment is configured, this option lets you instruct the plug-in to ignore the warning. Quest discourages use of this option.

9   To save the settings, click **OK**.

   With the authentication details properly configured, you can click the **Exchange Server** node to display the available Mailbox Databases.

Backing up data

- Selecting data for a backup
- Setting backup options
- Finalizing and submitting the backup job

Selecting data for a backup

You must use sets—Backup Selection Set, Backup Options Set, Schedule Set, Target Set, and Advanced Options Set—to create a backup job.

Backup Selection Sets are essential for Incremental and Differential Backups. Create the Backup Selection Set during a Full Backup, and use it for Full, Incremental, and Differential Backups. The backup job reports an error if you do not use a Selection Set for the Incremental or Differential Backup. For more information, see the *Quest NetVault Backup Administrator's Guide*.

> **i** | **TIP:** To use an existing set, click **Create Backup Job**, and select the set from the **Selections** list.

1 In the Navigation pane, click **Create Backup Job**.

 You can also start the wizard from the Guided Configuration link. In the Navigation pane, click **Guided Configuration**. On the **NetVault Configuration Wizard** page, click **Create backup jobs**.

2 In **Job Name**, specify a name for the job.

 Assign a descriptive name that lets you easily identify the job when monitoring its progress or restoring data. The job name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. On Windows, there is no length restriction; however, a maximum of 40 characters is recommended.

3 Next to the **Selections** list, click **Create New**.

4 In the list of plug-ins, open **Plug-in** *for Exchange*, and drill down to the name of the Exchange Server or Exchange Virtual Server and its available nodes.

5 Expand the applicable Exchange Server to display all the Mailbox Databases contained within

6 Select the database that you want to back up by clicking the box to the left of that item, that is, to mark it with a green check. You can also select the entire Exchange Server.

 Selecting the **Microsoft Exchange Server** node is equivalent to selecting the **Exchange Server** or **NetVault Backup Virtual Client**, in that all databases included in the DAG or standalone deployment are backed up. However, selecting the **Microsoft Exchange Server** node enables creating a selection set that can be used across multiple installations of Plug-in *for Exchange*.

7 Click **Save**, enter a name in the **Create New Set** dialog box, and click **Save**.

 The name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. On Windows, there is no length restriction. However, a maximum of 40 characters is recommended on all platforms.

Setting backup options

The next step involves creating the Backup Options Set or selecting an existing one.

> **i** | **TIP:** To use an existing set, in the **Plugin Options** list, select the set that you want to use.

1 Next to the **Plugin Options** list, click **Create New**.

2 In the **Exchange Backup Type** section, select the applicable option.

- **Full:** A Full Backup performs a complete backup of the set of items that have been selected. A Full Backup does not depend on any other backup and can be restored in a single step. In Exchange Server, a Full Backup includes all database files and transaction log files.

- **Copy:** A Copy Backup backs up the set of items that have been selected without performing a transaction log file truncation. This backup provides a snapshot of the files as they existed at the time of the backup. A Copy Backup should be used to perform a fast, nonintrusive complete backup of an Exchange Server and is commonly used for offline scaling and migration scenarios.

- **Incremental:** An Incremental Backup backs up the transaction log files of the set of items that have been selected since the most recent Full Backup or Incremental Backup. After the backup is completed, the Exchange Server truncates the transaction logs. This truncation helps maintain the operating health of the Exchange Server.

- **Differential:** A Differential Backup backs up the transaction log files of the set of items that have been selected since the most recent Full Backup or Incremental Backup. Differential Backups do not perform transaction log truncation, and therefore do not maintain the operating health of the Exchange Server.

For more information, see Reviewing the backup types.

3 In the **Snapshot Options** section, complete the following:

- **Backup Files to Storage:** This option, which is selected by default, copies the selected data from snapshots to a storage device. The plug-in uses a VSS provider to create persistent or non-persistent snapshots on the client, and copies the selected data from the snapshot, or snapshots, to the storage device.

- **Perform Backup Off-Host using Client:** Select this option if you want to enable a backup job from a remote NetVault Backup Client that meets the following criteria:

  - You are protecting a standalone deployment; that is, you are not protecting a DAG.

  - The remote client has access to the same array manager as the original client that was configured to run the job.

  - The remote client is not a NetVault Backup Virtual Client, nor was the original client.

  > **i** | **IMPORTANT:** To use this option, you must also select the **Backup Files to Storage** and **Retain Array-based Persistent Snapshot** options.

- **Retain Array-based Persistent Snapshot**: To retain persistent snapshots on a disk array, select this check box. Note the following:

  - If the **Backup Files to Storage** and **Retain Array-based Persistent Snapshot** check boxes are both selected, the plug-in copies the selected files to a storage device and adds the snapshot information to the backup index.

  - If you clear the **Backup Files to Storage** check box and select this check box, the plug-in writes only the index entries to the backup stream and creates persistent snapshots on the disk array.

  > **i** | **NOTE:** This option is visible even if you are not using snapshots that target a disk array.

- **Discard After:** If you selected the **Retain Array-based Persistent Snapshot** check box and you want to delete persistent snapshots from a disk array after a specified period, select this check box and complete the expiration fields.

  Regardless of the expiration settings, a snapshot is automatically deleted when the associated saveset is retired. Therefore, this option is only useful when you want to discard a snapshot *before* the associated saveset is retired. To expire a snapshot immediately, you must log in to the storage array, and manually expire the snapshot. Otherwise, the snapshot expires according to the retention policy selected while creating the snapshot.

> **i** | **NOTE:** To use the **Perform Backup Off-Host using Client**, **Retain Array-based Persistent Snapshot**, and **Discard After** options, the Exchange Mailbox Server that you back up must reside on a NetVault Backup-supported disk array and the vendor's hardware-based VSS provider must be in place. For more information, see Understanding snapshot-based backups and restores.

4   In the **Advanced Options** section, complete the following:

- **Maximum Number of Parallel Streams**: To enable the fastest backups possible, the plug-in uses parallel streams to back up the Mailbox Databases simultaneously. This option lets you specify the maximum number of parallel streams to use during a backup. The maximum number of parallel streams must meet *both* of the following conditions:

  - Less than or equal to (<=) the number of drives for the backup devices specified on the **Target** tab.

  - Less than **Maximum number of Mailbox Databases** that can be created. The following table details the maximum number of Mailbox Databases allowed by each Exchange Server version and edition:

Exchange Server method	Server edition	Maximum number of Mailbox Databases
Exchange Server 2010	Standard	5 Mailbox Databases
	Enterprise	100 Mailbox Databases
Exchange Server 2013	Standard	5 Mailbox Databases
	Enterprise	50 Mailbox Databases
Exchange Server 2016	Standard	5 Mailbox Databases
	Enterprise	100 Mailbox Databases
Exchange Server 2019	Standard	5 Mailbox Databases
	Enterprise	100 Mailbox Databases

For example, if you are targeting a Virtual Tape Library configured with two drives and you have 20 Mailbox Databases, the **Maximum Number of Parallel Streams** should be **2** or less.

> **i** | **NOTE:** In a DAG deployment, the number of parallel streams is allocated on a per-server basis; therefore, the number applies to each server included in the DAG. For example, if your DAG includes three Exchange Servers and you select two streams for backup, it can create up to six simultaneous streams assuming that there are at least two databases on each server that must be backed up. If there are fewer databases, the streams are not created. If there are fewer drives available, the streams wait until the next drive is available.

> **i** | **IMPORTANT:** If you use parallel streams, do not select the **Ensure this backup is the first on the media** option in the **Target Storage** section. If the plug-in generates more than one data stream, each data stream targets a separate piece of media and assumes that it is the first backup on the media.

- **Mount unmounted databases before backup**: This option mounts any unmounted databases before attempting to back up the database.

- **Perform VSS Consistency Check:** This option indicates whether a consistency check of the VSS should be performed on the Exchange Mailbox Databases included in the backup job. Quest recommends that you use this option, which is selected by default.

▪ **Incomplete Backup of All Items Selected:** When multiple items are included in a backup and the plug-in is unable to back up all the selected items, the plug-in lets you specify what action the backup should take. For example, if a job includes multiple Mailbox Databases and the backup of all but one is successful, you can specify what action the backup job should take.

   ▫ **Complete with Warnings — Saveset Retained:** The job returns a status of **"Backup Completed with warnings"** and a backup saveset is created that includes the items that were successfully backed up.

   ▫ **Complete without Warnings — Saveset Retained:** The job completes and returns a status of **"Backup Completed."** The errors are logged in the NetVault Backup binary logs and ignored on the **Job Status** page. A backup saveset is created that includes the items that were backed up.

   ▫ **Fail — Saveset Retained:** The job returns a status of **"Backup Failed."** However, a backup saveset is generated that includes the items that were successfully backed up.

   ▫ **Fail — No Saveset Retained:** The job returns a status of **"Backup Failed"** and no saveset of backed-up objects is kept. That is, even if some of the objects were successfully backed up, the saveset is discarded.

▪ **Backup Algorithm in DAG** (only available in DAG deployments):

   ▫ **All Active**: If you want the Plug-in *for Exchange* on each Mailbox Server to back up only the databases that are active when the backup is performed, select this option.

   ▫ **Backup Database Copies with the Lowest Activation Preference Number**: Select this option if you want the plug-in to select an available database copy. The plug-in selects the copy with the lowest activation preference number for backup. If a database copy is not available for backup or the copy is disabled, the plug-in selects the active database.

   ▫ **Backup Database Copies with the Highest Activation Preference Number**: Select this option if you want the plug-in to select the available database copy with the highest activation preference number. If a database copy is not available for backup or the copy is disabled, the plug-in selects the active database.

   ▫ **Backup Databases from preferred Exchange Server whenever possible**: Select this option if you want to specify a list of preferred Exchange Servers to perform backups from and a list of servers to exclude. The plug-in skips database copies that are in a disabled state.

   This option works with the following fields:

   - **If possible Backup Exchange Databases only from the following Exchange Server**: If you selected **Backup Databases from preferred Exchange Server whenever possible**, enter a comma-separated list of Exchange Servers, in order of preference, that should be used to back up the database. If the database, copy or active, is available in the Exchange Server listed first, it is backed up from that server. If not, the plug-in tries the next server in the list, and so on. If there is not an available copy—including the active database—on any of the listed servers, the plug-in uses the server that contains the active database or the database copy with the lowest activation preference number.

   - **If possible do not Backup Exchange Databases from the following Exchange Server**: If you selected **Backup Databases from preferred Exchange Server whenever possible**, enter a comma-separated list of Exchange Servers, in order of preference, that should *not* be used to back up the database. The server listed first has the lowest preference. If the database, copy or active, is available on any other Exchange Server, the database is not backed up from the server listed. The plug-in tries to exclude all servers listed in this field. If the database is *only* available on one of the excluded servers, the plug-in performs the backup using the last server in the list of exclusions.

   ⓘ **NOTE:** The **If possible Backup Exchange Databases only from the following Exchange Server** option takes precedence over **If possible do not Backup Exchange Databases from the following Exchange Server**. If the backup can be performed using a server from the first list, the plug-in does not process the list of excluded servers.

5   Click **Save** to save the set.

6   In the **Create New Set** dialog box, specify a name for the set, and click **Save**.

    The name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. On Windows, there is no length restriction; however, a maximum of 40 characters is recommended.

Finalizing and submitting the backup job

1   Use the **Schedule**, **Target Storage**, and **Advanced Options** lists to configure any additional required options.

2   Click **Save** or **Save & Submit**, whichever is applicable.

    > **i** | **TIP:** To run a job that you have already created and saved, select **Manage Job Definitions** on the Navigation pane, select the applicable job, and click **Run Now**.

    You can monitor progress on the **Job Status** page and view the logs on the **View Logs** page. For more information, see the *Quest NetVault Backup Administrator's Guide*.

    > **i** | **IMPORTANT:** Restarting services that host Exchange VSS Writers, including the **Microsoft Exchange Information Store**, **Microsoft Exchange Replication Service**, complete OS reboots, and cluster failovers causes running backup jobs to fail. This failure is due to a loss of in-memory 'backup-in-progress' state that Microsoft maintains for VSS backups.

Restoring data

- Restore and recovery overview
- Restoring data using the plug-in
- Using advanced restore procedures

Restore and recovery overview

When preparing to restore an existing backup, several key concepts are useful to understand. The following topics explore these concepts.

- Reviewing the available restore methods
- Reviewing the Phases of the Restore Sequence

Reviewing the available restore methods

The Restore Method is the technology that is selected to perform the restore. The Plug-in *for Exchange* supports several Restore Methods, each with its own specific purpose.

Volume Shadow Copy Service (VSS)

During a VSS-based restore, the plug-in instructs the **Exchange Store Writer** to coordinate with the **Exchange Information Store**—on behalf of the plug-in—to prepare the system for the restore, verify the restore targets, put back the data from the backup device, and then replay the transaction logs as needed.

Copy Files Restore Method

Plug-in *for Exchange* lets you restore the raw files from storage to a specific target directory. This method provides greater freedom in restoring database files, particularly when restoring files for nontraditional purposes, such as data mining or compliance auditing.

The Copy Files Restore Method does not require that you dismount databases from the Exchange Server. However, Database Recovery is not supported when using this restore method.

> **i** | **NOTE:** The Copy File Restore Method is required when creating a Staging Area or Target Directory to perform individual mailbox and mailbox item recovery. For more information, see the *Quest Recovery Manager for Exchange User Guide*.

Reviewing the Phases of the Restore Sequence

Each Exchange Server restore scenario has two distinct phases: the Restore Phase and the Recovery Phase.

Restore phase

The Restore Phase is used to transfer all the files included in a backup from storage to the local disk. This process can be done using one of the Restore Methods defined in Reviewing the available restore methods.

When performing restore sequences that include Incremental or Differential Backups, the Restore Phase is performed for each backup that is restored.

Ensure that adequate disk space is available on the volume or volumes where the databases are stored during a restore. If disk space is exhausted during the Restore Phase, the restore job fails.

Recovery phase

In the Recovery Phase, changes that were logged in the transaction log files are reapplied to the corresponding databases to bring the Exchange Server forward in time. This phase is only available when using the VSS Restore Method.

When restoring sequences of backups (if there are Incremental or Differential Backups), the Recovery Phase is performed one time as the final step in the restore sequence. Recovery is only performed once, regardless of the number of backups included in the restore sequence.

Restoring data using the plug-in

A standard restore procedure includes the steps outlined in the following topics:

- Selecting data for a restore
- Setting restore and recovery options
- Setting security options
- Finalizing and submitting the restore job

Selecting data for a restore

1   On the Navigation pane, click **Create Restore Job**.

2   On the **Create Restore Job—Choose Saveset** page, select **Plug-in** *for Exchange* from the **Plugin Type** list.

3   To filter the items displayed in the saveset table further, use the **Client**, **Date**, and **Job ID** lists.

The table displays the saveset name (job title and saveset ID), creation date and time, and size. By default, the list is sorted by creation date.

The following table outlines the backup-type identifiers for VSS:

Backup type	Backup type identifier
Full	VSS FULL
Copy	VSS COPY
Incremental	VSS INCREMENTAL
Differential	VSS DIFFERENTIAL

4   In the saveset table, select the applicable item.

When you select a saveset, the following details are displayed in the **Saveset Information** area: Job ID, job title, server name, client name, plug-in name, saveset date and time, retirement setting, Incremental Backup or not, Archive or not, saveset size, and snapshot-based backup or not.

5   Click **Next**.

6   On the **Create Selection Set** page, select the data that you want to restore.

NetVault Backup lets you select from the following items for inclusion in a restore job, depending on what version of Exchange you are using:

- Entire Information Store

- Individual Databases—including Mailbox Databases and Public Folder Databases

- Sets of Databases

7   If you are restoring data from a persistent snapshot, review the following:

- If your backup target was NetVault Backup media, skip to Finalizing and submitting the restore job.

  To restore data from a persistent snapshot, no additional steps are required before you submit the job. After you submit the job, the plug-in automatically obtains the snapshot information from the backup index. If the snapshot is available, the plug-in mounts it locally and restores the selected files from the snapshot. If the snapshot is unavailable, the plug-in restores the file data from the storage media.

- If your backup target was non-NetVault Backup media, such as a disk array, review the following information before you skip to Finalizing and submitting the restore job.

  If you restore an Exchange database to a DAG and do not specify a Target Client, the plug-in restores the database to the node that hosts the active copy. If you restore the database from a persistent snapshot, the plug-in maps the volume in the client that took the snapshot during the backup process. Because of these conditions, restoring from a persistent snapshot to a DAG, that is, where the Target Client is a NetVault Backup Virtual Client, requires that the active copy of the database is hosted on the node that takes the snapshots during the backup process. For more information, see Understanding snapshot-based backups and restores.

Setting restore and recovery options

On the **Create Selection Set** page, click **Edit Plugin Options**, and configure the following parameters on the **Restore and Recovery Options** tab:

- **Restore Method:** Select one of the following:

    - **Volume Shadow Copy Service (VSS)**: To perform the restore of selected data using the VSS method, which allows for the recovery of Full, Differential, and Incremental Backup types, select this option. For more information, see Reviewing the Volume Shadow Copy Service (VSS) backup method.

    - **Copy Files to Target Directory Only**: Select this option if you want to restore the raw files from storage to a specific target directory. When choosing this method, database recovery is not available. When restoring using this method, you do not have to dismount databases from the Exchange Server before the start of the restore job.

- **Restore Options:** Select the applicable options:

    - **Target Directory:** If you selected the **Copy Files to Target Directory Only** option, update the path for the target if you do not want to use the default.

        i | **IMPORTANT:** Ensure that the temporary directory is empty. Failing to do so can cause restore jobs to fail with "Restore Environment Information Corrupted" errors.

- **Maximum Number of Parallel Streams**: Specify the maximum number of parallel restore streams to use during the restore. By default, this field displays the maximum number that is allowed based on the version and edition of Exchange Server that was used to generate the backup.

- **Do Not overwrite existing Files:** The option, which is selected by default, applies to the physical files that an Exchange Server database consists of, including database files, transaction logs, and partial transaction logs. This option works with both of the preceding Restore Method options.

  - To complete a soft recovery, leave this option selected to ensure that the plug-in does *not* overwrite a file of the same name in the specified target directory. For example, if you are restoring files to a database directory in which there are transaction logs and partial transaction logs that were created since the last backup was completed, selecting this option retains the existing logs and recovery includes the existing logs as well as additional transaction logs that were restored from the backup.

  - To complete a hard recovery, clear this option to stop the plug-in from checking whether files of the same name exist in the specified target directory. For example, if you are restoring files to an empty database directory or you want to recover files up to the time that the last backup was completed, clearing this option ensures that the files from the last backup are recovered.

  > **i** | **IMPORTANT:** If you are restoring a sequence of Full and Incremental or Differential Backups, use the *same* setting for the **Do Not overwrite existing Files** option for *each* step in the sequence. For example, if you select the option when you restore the base Full Backup, make sure that you select the option for *every* subsequent Incremental or Differential Backup as well.

- **Dismount Databases before Restore**: To dismount the databases that have been selected for restore automatically before actually performing the restore, select this option. *All* databases in selected Mailbox Databases are dismounted when restoring VSS-based backups. Also, databases in RSGs are not dismounted when this option is selected.

- **Set 'This database can be overwritten by a restore' Flag for Databases**: To restore all databases involved in the restore operation automatically and eliminate the requirement to use an Exchange utility for this operation, select this option. If you do not select this option, Quest recommends that you use the Exchange Management Console, Exchange System Manager, or Exchange Management Shell to set the Overwrite Permissions to ensure a successful restore.

  > **i** | **IMPORTANT:** Selecting **This database can be overwritten by a restore** has the following limitations:
  > - When an RDB is configured for a database, the overwrite permission is not set for the databases in the RDB. Quest recommends that you do not select this option if the restore location is redirected to an RDB.
  > - When you rename a Mailbox Database, the overwrite permissions are set for all databases in the target Mailbox Database.
  > - When you rename databases, the overwrite permissions are set for the target databases.

- **Suspend database Replication before Restore**: (only available for DAG deployments) Select this option to disable database copying automatically for all Mailbox Databases that have DatabaseCopy enabled for DAG.

  > **i** | **IMPORTANT:** If this is the first restore that you are performing in the process of recovering a DAG environment, you must select this option. For more information, see Additional guidelines for restoring a DAG.

- **Update database Replication after Recovery (Deletes Existing Files and Resumes Replication)**: (only available for DAG deployments) Select this option to update database copying for DAG automatically. After the existing files of the database copy have been deleted, the database copy is updated by copying the new database and log files. This step deletes the existing database copy and log files. After the restore job has been completed with this option, the status in the **Exchange Management Console** should be **Healthy**. Updating a database copy might take a

significant amount of time depending on the size of the database and the network bandwidth and latency.

> **ⓘ IMPORTANT:** If this is the last restore that you are performing in the process of recovering a DAG environment, you must select this option. For more information, see Additional guidelines for restoring a DAG.

> **ⓘ NOTE:** After the **Update Database Replication after Recovery** has completed successfully, the Mailbox Database **Copy Status** restored in the Exchange Management Console might show **Initializing** instead of **Healthy**. However, this status is not a problem because the Copy Status will change to Healthy after a new log file is generated for the active Mailbox Database.

- **Recovery Options:** Select the applicable options:

  - **Perform Recovery (Enable only on Last Step in Restore Sequence)**: Select this option when performing the last job of a restore sequence. If the restore sequence only includes a Full Backup, select this option for the Full Backup. If the restore sequence includes Incremental or Differential Restores, select this option when restoring the **last backup in the restore sequence**. This option finalizes the restore process after completion of the restore job for which it was selected, that is, the Exchange Server is then alerted to perform any necessary recovery tasks. All other options in the **Recovery Options** section are disabled unless this option is selected.

    > **ⓘ IMPORTANT:** When restoring VSS-based backups, current logs are automatically replayed. If you want to restore up to the point of the last backup, delete the current logs before the VSS Restore is performed.

  - **Mount Databases after Recovery**: (selected by default) This option can be selected to mount the dismounted databases within the databases automatically after the restore completes. Quest recommends that you leave this option selected; otherwise, manually mount the databases after completion of the restore to bring it back online.

Setting security options

To specify job-specific authentication, click the **Security** tab, and configure the following parameters:

- **Exchange Administrator User Name**: By default, this field contains the administrator-level account that was used to back up the original target Exchange Server. You can change this value, but the account named must have the proper backup and restore permissions for the machine that is to serve as the target of the restore.

- **Password**: Enter the password associated with the user name specified in the previous field.

- **Windows Domain**: If a domain account is specified in the **Exchange Administrator User Name** field, enter the name of the domain. The field can be left blank if a local system account is specified.

Finalizing and submitting the restore job

The final steps include setting additional options on the Schedule, Source Options, and Advanced Options pages, submitting the job, and monitoring the progress through the Job Status and View Logs pages. These pages and options are common to all NetVault Backup Plug-ins. For more information, see the *Quest NetVault Backup Administrator's Guide*.

1   To save the settings, click **Ok**, and then click **Next**.

2   In **Job Name**, specify a name for the job if you do not want to use the default setting.

    Assign a descriptive name that lets you easily identify the job when monitoring its progress. The job name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. On Windows, there is no length restriction; however, a maximum of 40 characters is recommended.

3   In the **Target Client** list, select the machine on which you want to restore the data.

> **i** | **TIP:** You can also click **Choose**, and then locate and select the applicable client in the **Choose the Target Client** dialog box.

4 Use the **Schedule**, **Source Options**, and **Advanced Options** lists to configure any additional required options.

5 Click **Save** or **Save & Submit**, whichever is applicable.

You can monitor progress on the **Job Status** page and view the logs on the **View Logs** page. For more information, see the *Quest NetVault Backup Administrator's Guide*.

Performing post-restore steps DAG environments

If you did not use the **Update database Replication after Recovery** option, use the Exchange Management Console or Shell to update the database copies manually. For more information, see http://technet.microsoft.com/en-us/library/dd351100.aspx.

Examples of restore sequences

NetVault Backup sees all restore scenarios using the plug-in as a sequence of restores, from the recovery of a series of 10 Incremental Backups to the recovery of a single Full Backup. Therefore, you must follow a specific set of steps to recover data successfully. The following topics offer examples for various backup types. These examples assume that a noncontinuous replication environment is deployed.

- Restoring a standalone Full Backup

- Restoring Incremental Backup sequence

- Restoring Full and Differential Backup sequence

- Additional guidelines for restoring a DAG

Restoring a standalone Full Backup

Restoring a single Full Backup is seen as a restore sequence by the Plug-in *for Exchange*. Therefore, you must follow these steps to restore a standalone Full Backup and ensure that its recovered data is suitable for use by Exchange.

1 On the Navigation pane, click **Create Restore Job**.

2 On the **Create Restore Job—Choose Saveset** page, select **Plug-in *for Exchange*** from the **Plugin Type** list.

3 In the saveset table, select the applicable item, and then click **Next**.

4 On the **Create Selection Set** page, select the objects to restore.

For example, you can select the entire Information Store or only the databases that have been corrupted as detailed in Selecting data for a restore.

5 Note which items you selected.

6 Click **Edit Plugin Options**, and select the following options on the **Restore and Recovery Options** tab:

- Verify that the **Backup Type** is labeled **FULL**.

- In the **Restore Method** section, select V**olume Shadow Copy Services (VSS)**.

- In the **Restore Options** section, select **Restore Files from Backup**, enter the **Maximum Number of Parallel Streams**, and select **Dismount Databases before Restore**.

- In the **Recovery Options** section, select **Perform Recovery**, and then select **Mount Databases after Recovery**. If you are using the **VSS Restore Method**, verify that **Verify Recovery Completion**, **Replay Current Logs**, and **Delete Temp Log Files after Recovery** are also selected.

7   Click the **Security** tab, and enter the proper **Authentication Details**.

8   Complete and submit the job.

For details, see Finalizing and submitting the restore job.

On job completion, NetVault Backup finalizes the restore and recovery process and all restored data is accessible to the Exchange Server.

Restoring Incremental Backup sequence

This example assumes that Full Backups are performed every Sunday night at 11:00 p.m. and that Incremental Backups are performed Monday through Saturday at 11:00 p.m. Because recovery is being performed on Wednesday, Sunday's Full Backup followed by Monday's and Tuesday's Incremental Backup must be restored.

Recovering a sequence of Incremental Backups includes the steps outlined in the following topics:

- Restoring the original Full Backup
- Restoring all intermediary Incremental Backups
- Restoring the final Incremental Backup

Restoring the original Full Backup

1   On the Navigation pane, click **Create Restore Job**.

2   On the **Create Restore Job—Choose Saveset** page, select **Plug-in** *for Exchange* from the **Plugin Type** list.

3   In the saveset table, select the Full Backup saveset that served as a starting point for the applicable Incremental Backup sequence, and then click **Next**.

4   On the **Create Selection Set** page, select the objects to restore.

For example, you can select the entire Information Store or only the databases that have been corrupted as detailed in Selecting data for a restore.

5   Note which items you selected.

6   Click **Edit Plugin Options**, and select the following options on the **Restore and Recovery Options** tab:

- Verify that the **Backup Type** is labeled **FULL**.
- In the **Restore Method** section, select **Volume Shadow Copy Services (VSS)**.
- In the **Restore Options** section, select **Restore Files from Backup**, enter the **Maximum Number of Parallel Streams**, and select **Dismount Databases before Restore**.
- In the **Recovery Options** section, clear **Perform Recovery**. All other options should be disabled.

7   Click the **Security** tab, and enter the proper **Authentication Details**.

8   Complete and submit the job.

For details, see Finalizing and submitting the restore job.

Restoring all intermediary Incremental Backups

1   On the Navigation pane, click **Create Restore Job**.

2   On the **Create Restore Job—Choose Saveset** page, select **Plug-in** *for Exchange* from the **Plugin Type** list.

3   In the saveset table, select the backup saveset that served as the first Incremental Backup in the sequence, and then click **Next**.

4   On the **Create Selection Set** page, select the objects to restore.

5   Click **Edit Plugin Options**, and select the following options on the **Restore and Recovery Options** tab:

- Verify that the **Backup Type** is labeled **INCREMENTAL**.

- In the **Restore Method** section, select **Volume Shadow Copy Services (VSS)**.

- In the **Restore Options** section, select **Restore Files from Backup**, and enter the **Maximum Number of Parallel Streams**.

- In the **Recovery Options** section, clear **Perform Recovery**. All other options should be disabled.

6   Click the **Security** tab, and enter the proper **Authentication Details**.

7   Complete and submit the job.

For details, see Finalizing and submitting the restore job.

8   Repeat Step 1 through Step 7 for all additional Incremental Backups in the order they were performed except for the *last* Incremental in the sequence.

Specific settings must be made for the last Incremental in the sequence; these settings are covered in the following topic.

Restoring the final Incremental Backup

1   With all previous restores successfully completed, including the original Full and subsequent Incrementals, click **Create Restore Job** on the Navigation pane.

2   On the **Create Restore Job—Choose Saveset** page, select **Plug-in *for Exchange*** from the **Plugin Type** list.

3   In the saveset table, drill down to the last Incremental Backup saveset in the sequence, and then click **Next**.

4   On the **Create Selection Set** page, see the data items that were restored in all previous Incremental Backup restore jobs, and select the *same items* for the restore.

5   Click **Edit Plugin Options**, and select the following options on the **Restore and Recovery Options** tab:

- Verify that the **Backup Type** is labeled **INCREMENTAL**.

- In the **Restore Method** section, **Volume Shadow Copy Services (VSS)**.

- In the **Restore Options** section, select **Restore Files from Backup**, and enter the **Maximum Number of Parallel Streams**.

- In the **Recovery Options** section, select the following:

   □ Select **Perform Recovery**.

   □ If the most recent Differential Backup is selected for restore, select **Replay Current Logs**. If the Differential Backup selected for restore is *not* the most recent, that is, you are restoring an Exchange Server to a specific point before the last available Differential Backup job, clear **Replay Current Logs**.

   If this option is selected, the Exchange Server scans all recent logs and updates all data to its most recent state. The Exchange Server performs this process instead of leaving the restored data at the state it was in at the time the Incremental Backup was completed.

   When restoring VSS-based backups, current logs are automatically replayed. Therefore, if you want to restore up to the point of the last backup, the current logs must be deleted before the VSS Restore is performed.

   □ Select **Mount Databases after Recovery**.

6   Click the **Security** tab, and enter the proper **Authentication Details**.

7   Complete and submit the job.

For details, see Finalizing and submitting the restore job.

Restoring Full and Differential Backup sequence

This example assumes that Full Backups are performed every Sunday night at 11:00 p.m. and that Differential Backups are performed Monday through Saturday at 11:00 p.m. If recovery is performed on Tuesday, Sunday's Full Backup and Monday's Differential Backup must be restored. If recovery is performed on Thursday, Sunday's Full Backup followed Wednesday's Differential Backup must be restored.

Recovering a Full and Differential Backup includes the steps outlined in the following topics:

- Restoring the original Full Backup
- Restoring the applicable Differential Backup

Restoring the original Full Backup

1   On the Navigation pane, click **Create Restore Job**.

2   On the **Create Restore Job—Choose Saveset** page, select **Plug-in *for Exchange*** from the **Plugin Type** list.

3   In the saveset table, select the Full Backup saveset that served as a starting point for the applicable Differential Backup sequence, and then click **Next**.

4   On the **Create Selection Set** page, select the objects to restore.

    For example, you can select the entire Information Store or only the databases that have been corrupted as detailed in Selecting data for a restore.

5   Note which items you selected.

6   Click **Edit Plugin Options**, and select the following options on the **Restore and Recovery Options** tab:

    - Verify that the **Backup Type** is labeled **FULL**.

    - In the **Restore Method** section, select **Volume Shadow Copy Services (VSS)**.

    - In the **Restore Options** section, select **Restore Files from Backup**, enter the **Maximum Number of Parallel Streams**, and select **Dismount Databases before Restore**.

    - In the **Recovery Options** section, clear **Perform Recovery**. All other options should be disabled.

7   Click the **Security** tab, and enter the proper **Authentication Details**.

8   Complete and submit the job.

    For details, see Finalizing and submitting the restore job.

Restoring the applicable Differential Backup

1   On the Navigation pane, click **Create Restore Job**.

2   On the **Create Restore Job—Choose Saveset** page, select **Plug-in *for Exchange*** from the **Plugin Type** list.

3   In the saveset table, select the applicable Differential Backup in the sequence, and then click **Next**.

4   On the **Create Selection Set** page, select the objects to restore.

    For example, you can select the entire Information Store or only the databases that have been corrupted as detailed in Selecting data for a restore.

5   Click **Edit Plugin Options**, and select the following options on the **Restore and Recovery Options** tab:

    - Verify that the **Backup Type** is labeled **DIFFERENTIAL**.

    - In the **Restore Method** section, select **Volume Shadow Copy Services (VSS)**.

    - In the **Restore Options** section, select **Restore Files from Backup**, and enter the **Maximum Number of Parallel Streams**.

- In the **Recovery Options** section, select the following:

  - Select **Perform Recovery**.

  - If the most recent Differential Backup is selected for restore, select **Replay Current Logs**. If the Differential Backup selected for restore is *not* the most recent, that is, you are restoring an Exchange Server to a specific point before the last available Differential Backup job, clear **Replay Current Logs**.

    If this option is selected, the Exchange Server scans all recent logs and updates all data to its most recent state. The Exchange Server performs this process instead of leaving the restored data at the state it was in at the time the Incremental Backup was completed.

    When restoring VSS-based backups, current logs are automatically replayed. Therefore, if you want to restore up to the point of the last backup, the current logs must be deleted before the VSS Restore is performed.

  - Select **Mount Databases after Recovery**.

6 Click the **Security** tab, and enter the proper **Authentication Details**.

7 Complete and submit the job.

For details, see Finalizing and submitting the restore job.

On job completion, NetVault Backup finalizes the restore and recovery process and all restored data is accessible to the Exchange Server.

Additional guidelines for restoring a DAG

Whether you are restoring a Full Backup, Incremental Backup, or Differential Backup for a DAG, there are two options on the **Restore and Recovery Options** tab that you can select or clear depending on where you are in the restore process.

- **Suspend database Replication before Restore**: If this is the *first* restore that you are performing in the process of recovering a DAG environment, you must select this option. This ensures that replication is suspended and recovery continues.

- **Update database Replication after Recovery (Deletes Existing Files and Resumes Replication)**: If this is the *last* restore that you are performing in the process of recovering a DAG environment, you must select this option. This ensures that replication resumes and recovery continues.

Using advanced restore procedures

This topic describes other restore operations that you can perform with the plug-in.

- Renaming Mailbox Databases in Exchange 2010

- Restoring data to a Recovery Database (RDB)

- Performing a disaster recovery of an Exchange Server

- Recovering to an alternate Exchange Server

Renaming Mailbox Databases in Exchange 2010

The Plug-in *for Exchange* allows you to rename a Mailbox Database during VSS-based restores. Renaming is useful if you do not want to overwrite the existing version and want to create a copy of the Mailbox Database. The following prerequisites are met before initiating the restore:

Ensure that the following prerequisites are met before initiating the restore:

- **Target Mailbox Database must exist**: The target Mailbox Database must be created in the Exchange System Manager or Exchange Management Console.

- **Database names must be different from the original**: The name of the target Mailbox Database must be different from the name of the original database.

- **Database filenames must be different from the original**: The database filenames in the target Mailbox Database must be different from the database filenames used in the original Mailbox Database.

> **i** | **NOTE:** Due to a scripting limitation of the Windows PowerShell, the plug-in does not support the use of single or double quotation marks in Mailbox Database names.

To complete the rename process, perform the following steps.

1 After the prerequisites have been met, click **Create Restore Job**.

2 Select **Plug-in *for Exchange*** from the **Plugin Type** list.

3 Select the Mailbox Database to restore from the backup saveset by following the instructions outlined in Selecting data for a restore, and click **Next**.

4 On the **Create Selection Set** page, click the Mailbox Database that you want to rename and select **Rename** from the context menu.

5 In the **Rename/Relocate** dialog box, enter the existing name of the target Mailbox Database in the **Rename** box, and click **Ok**.

   The name of the Mailbox Database is updated to reflect the rename information in parentheses.

6 Continue with the restore procedure as explained in *Setting restore and recovery options* and Finalizing and submitting the restore job.

7 Open the **Exchange Management Shell**.

8 Run the following command:

```
Get-Mailbox -Database <SourceDatabase> |where {$_.ObjectClass
   -NotMatch '(SystemAttendantMailbox|ExOleDbSystemMailbox)'}
   | Set-Mailbox -Database <TargetDatabase>
```

   Replace `<SourceDatabase>` with the name of the original database and `<TargetDatabase>` with the name of the destination.

Restoring data to a Recovery Database (RDB)

An RDB is a special database that is created to hold a restored database where data can be extracted without disturbing user access. The RDB enables the recovery of accidentally deleted mailboxes and other email items, as well supporting data recovery for legal inquiries.

You can find a complete description of RDBs and their usage in the relevant Microsoft Exchange documentation. For more information, see http://technet.microsoft.com/en-us/library/dd876954.aspx.

- **Supported Backup Method**: VSS only

The following procedures illustrate the steps required to establish an RDB, and how to use NetVault Backup to restore backed-up data to it. To create an RDB, use the Exchange Management Shell.

> **i** | **IMPORTANT:** You can only mount one RDB at a time on an Exchange Mailbox Server. If you have already mounted a Recovery Database on your Exchange Mailbox Server, recovery of the new RDB fails.

Creating an RDB

1 Open the **Exchange Management Shell**.

2 Enter the following command:

```
New-MailboxDatabase -Recovery -Name <RDB_Name> -Server <Exchange_Server_Name>
```

> **ⓘ IMPORTANT:** When you create the RDB, do *not* use **Mount-Database**. If the database is mounted, the initial Exchange Server database transaction logs are generated. These logs might interfere with the restore job and cause it to fail.

You can find complete instructions for this procedure in the relevant Microsoft Exchange documentation. For more information, see http://technet.microsoft.com/en-us/library/ee332321.aspx.

Restoring data to an RDB

With the RDB successfully created, you can restore data to it by completing the steps outlined in Renaming Mailbox Databases in Exchange 2010. Amend those steps to including the following steps:

1 After you select the **Rename to** check box in the **Rename/Relocate** dialog box, enter the name of the RDB in the text box to direct the backup to the RDB.

2 On the **Create Selection Set** page, click **Edit Plugin Options**, and select the following options on the **Restore and Recovery Options** tab:

- ▪ **Set 'This database can be overwritten by a restore' Flag for Databases**
- ▪ **Perform Recovery (Enable only on Last Step in Restore Sequence)**
- ▪ **Mount Databases after Recovery**

3 In a DAG, if you created the RDB on an Exchange Mailbox Server that is *not* the active server and you want to select a specific Exchange Mailbox Server on which to perform the restore, select the applicable NetVault Backup Client from the **Target Client** list on the **Create Restore Job** page.

In a DAG, the plug-in automatically directs the restore job for an RDB to the Exchange Mailbox Server that hosts the RDB in the cluster. If applicable and if you created the RDB on an Exchange Mailbox Server that is not the active server, you can select the NetVault Backup Client in the Exchange Mailbox Server that hosts the RDB.

> **ⓘ IMPORTANT:** If the plug-in cannot resolve the address of the Exchange Mailbox Server that hosts the RDB, select the applicable NetVault Backup Client from the **Target Client** list on the **Create Restore Job** page.

Extracting the data from the RDB

After you have restored the database to the RDB, use the **Exchange Management Shell** to extract the data from the RDB to the active mailboxes.

- • Exchange 2010: If you are using the original version of Exchange 2010, use the following command:

  ```
  Restore-Mailbox -Identity <User_Name> -RecoveryDatabase <RDB_Name>
  ```

  The following example restores a mailbox for a user named Scott from the RDB named RDB1:

  ```
  Restore-Mailbox -Identity Scott -RecoveryDatabase RDB1
  ```

- • Exchange 2010 SP1 and later: If you are using service pack 1 or later of Exchange 2010, use the following command:

  ```
  New-MailboxRestoreRequest -SourceDatabase "<RDB_Name>"
    -SourceStoreMailbox "<SourceMailbox_GUID>" -TargetMailbox "<TargetMailBox>"
    -AllowLegacyDNMismatch
  ```

  The following example restores a mailbox with a GUID of 1d20855f-fd54-4681-98e6-e249f7326ddd for a user named Scott from the RDB named RDB1:

  ```
  New-MailboxRestoreRequest -SourceDatabase "RDB1"
    -SourceStoreMailbox "1d20855f-fd54-4681-98e6-e249f7326ddd"
    -TargetMailbox "Scott" -AllowLegacyDNMismatch
  ```

The optional `AllowLegacyDNMismatch` parameter lets you restore the source mailbox to a mailbox that does not have the same legacy DN value.

You can find complete instructions for this procedure in the relevant Microsoft Exchange documentation. For more information, see http://technet.microsoft.com/en-us/library/ee332351.aspx.

Performing a disaster recovery of an Exchange Server

This topic briefly describes how to recover an Exchange system by using backups performed with the Plug-in *for Exchange*. You can use this procedure to rebuild an Exchange Server if the original was destroyed or is otherwise unusable. If you are rebuilding an Exchange Server, the **Active Directory** service should be available in **Disaster Recovery Mode**.

The restore sequence performed should bring the Exchange Server up to the most current point before the failure of the Exchange Server. This sequence could include restoring a Full Backup, a Full Backup plus a Differential Backup, or a Full Backup plus Incremental Backups depending on your Exchange Server backup-and-recovery strategy. However, Quest strongly recommends that you consult the Microsoft documentation for the Exchange version in use for complete details on the proper way to perform this type of recovery. For additional information:

*   **Exchange 2010, Recover Exchange Server** or **Recover a Database Availability Group Member Server:** http://technet.microsoft.com/en-us/library/dd876880(v=exchg.141).aspx and http://technet.microsoft.com/en-us/library/dd638206(v=exchg.141).aspx

*   **Exchange 2013, Recover Exchange Server** or **Recover a Database Availability Group Member Server:** https://technet.microsoft.com/en-us/library/dd876880(v=exchg.150).aspx and https://technet.microsoft.com/en-us/library/dd638206(v=exchg.150).aspx

*   **Exchange 2016 and 2019, Recover an Exchange Server** or **Recover a Database Availability Group Member Server:** http://technet.microsoft.com/en-us/library/dd876880.aspx and http://technet.microsoft.com/en-us/library/dd638206.aspx. Pick the applicable version, 2016 or 2019.

To perform the restore, perform the following steps.

1   **Reinstall the OS.**

    The OS and service pack must be identical to the OS and service pack that was installed when the backup saveset was created.

2   **Reinstall Exchange in Recovery Mode.**

    The Exchange version, edition, and service pack must be identical to the version, edition, and service pack of Exchange that was installed when the backup saveset was created. For complete instructions, see the relevant Microsoft Exchange documentation identified earlier.

3   **Reinstall the NetVault Backup Client and the Plug-in *for Exchange*.**

4   **Add the Exchange Server to the NetVault Backup Domain as a Client.**

    For more information, see the *Quest NetVault Backup Administrator's Guide*.

5   **For details on restoring Exchange backups, see *Restoring data using the plug-in*.**

    > **i** | **IMPORTANT:** Because continuous replication is not enabled before initiating the restore, it is not necessary to disable or resume and update replication when restoring in a disaster-recovery scenario.

6   If applicable, enable continuous replication.

Recovering to an alternate Exchange Server

The Plug-in *for Exchange* also lets you take the backup of a Mailbox Database performed on one Exchange Server and restore it to a secondary Exchange Server. This option can be beneficial when you must recover a Mailbox Database to access its contents while leaving the initial Exchange Server untouched and operational. In addition, restoring to a secondary Exchange Server can be a good method for testing the integrity of backed-up data without interfering with a running Exchange Server. This topic offers instructions on restoring Full, Incremental, and Differential Backup types to a secondary Exchange Server.

Before performing this type of restore, the following actions must be performed on the machine that is to serve as the new target for the restored Exchange Server data:

- **Clean installation of Exchange has been performed on the new target server**: Quest recommends that you complete this step on a newly installed Exchange Server, rather than an existing one.

- **Mailbox Databases set up on the target machine:** The target Mailbox Database names, database names, database filenames, and the database file path or directory cannot match the original Mailbox Databases; they must be different.

  > **i** | **IMPORTANT:** When you use the Exchange Management Console to create the target database in the destination Exchange Server, do *not* select the Mount Database option. If you select this option, Exchange generates the initial set of database transaction logs. These logs might interfere with the restore job and cause it to fail. By not selecting the Mount Database option, Exchange adds the database to the Exchange Server and creates an empty directory for the database. The restore job populates the directory during completion of the job.

- **NetVault Backup and the Plug-in *for Exchange* must be installed on both machines**: For this operation to work, NetVault Backup and the plug-in must be installed on *both* the original Exchange Server and the machine set up as the new target.

- **Both machines must be added to the NetVault Backup Server as Clients**: On the **Manage Clients** page of the NetVault Backup Server, ensure that *both* machines have been successfully added as NetVault Backup Clients. For details on adding Clients to the NetVault Backup Server, see the *Quest NetVault Backup Administrator's Guide*.

To perform the restore, complete the following steps.

1 After the prerequisite operations are finished, use the NetVault Backup WebUI to restore the applicable Mailbox Databases.

   This process is similar to the standard restore process illustrated earlier, except for a few steps.

   This sample procedure assumes that a Full Backup is being restored and not an Incremental or Differential Backup. If you intend to recover a series of Incremental or Differential Backups as well, you would follow this procedure to recover the initial Full Backup first. Subsequent Incremental or Differential Backups would then be recovered as a standard restore of this type. The difference is that you must select the alternate Exchange Server from the **Target Client** list.

2 For details on restoring Exchange backups, see *Restoring data using the plug-in*, and note the following exceptions:

   - Because continuous replication is not enabled before initiating the restore, you do not have to disable or resume and update replication when restoring to an alternate server.

   - If the name of the target database for the restore job is different from the original database that was backed up, that is, you used the Exchange Management Console to create a database, complete the following steps:

     a On the **Create Selection Set** page, click the original database, and select **Rename** from the context menu.

     b In the **Rename/Relocate** dialog box, enter the name of the new target Mailbox Database in the **Rename** box, and click **Ok**.

   - On the **Restore and Recovery Options** tab, select **Set 'This database can be overwritten by a restore' Flag for Databases**.

- If applicable, change the other options on the **Restore and Recovery Options** tab.

  For more information, see Setting restore and recovery options.

- Click the **Security** tab, and complete the applicable **Authentication Details** fields:

  - **Exchange Administrator User Name**: By default, this field contains the administrator-level account that was used to back up the original target Exchange Server. You can change this value, but the account named must have the proper backup and restore permissions for the machine that is to serve as the target of the restore.

  - **Password**: Enter the password associated with the user name specified in the previous field.

  - **Windows Domain**: This field can be left blank, as it is not requirement for recovering to a different Exchange Server.

- In **Job Name**, specify a name for the job if you do not want to use the default setting.

  Assign a descriptive name that lets you easily identify the job when monitoring its progress. The job name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. On Windows, there is no length restriction; however, a maximum of 40 characters is recommended.

- In the **Target Client** list, select the Exchange Server on which the target database was created.

- Click **Save** or **Save & Submit**, whichever is applicable.

  You can monitor progress on the **Job Status** page and view the logs on the **View Logs** page. For more information, see the *Quest NetVault Backup Administrator's Guide*.

3   To point to the new database, open the **Exchange Management Shell**, and run the following command

```
Get-Mailbox -Database <SourceDatabase> |where {$_.ObjectClass
   -NotMatch '(SystemAttendantMailbox|ExOleDbSystemMailbox)'}
   | Set-Mailbox -Database <TargetDatabase>
```

Replace *<SourceDatabase>* with the name of the original database and *<TargetDatabase>* with the name of the destination.

Troubleshooting

- Upgrading the plug-in fails
- Diagnosing and resolving VSS-related issues
- Addressing cluster-related issues
- Troubleshooting other issues

Upgrading the plug-in fails

Symptom

If the **nvexchangeserver.exe** process is running when you install an upgrade, the upgrade might fail and display the following message:

```
Failed to install software. A key could not be installed.
```

Solution

To work around this issue, do either of the following on the NetVault Backup Client where Exchange Server is running:

- Use Windows Task Manager to locate and stop the **nvexchangeserver.exe** process.
- Stop and start the **NetVault Backup Service**.

For more information on stopping and starting the NetVault Backup Service, see the *Quest NetVault Backup Administrator's Guide*.

Diagnosing and resolving VSS-related issues

When a VSS-backup job is run, the plug-in might generate one or more shadow copies (snapshots). After completing the job, the plug-in sends the VSS subsystem a notification indicating that the shadow copy can be deleted. Also, if the VSS job is canceled for any reason and the plug-in detects the event, the plug-in sends the same type of notice. However, there are instances in which shadow copies are not deleted, which means that the VSS inadvertently stores obsolete shadow copies.

If the plug-in is unable to create a VSS shadow copy, or if the backup job fails during the processing of taking or managing snapshots, it might indicate that obsolete shadow copies are interfering with the VSS subsystem.

In addition to the NetVault Backup logs and the individual job logs, the following resources can expedite the diagnoses of VSS-related issues with backup and restore jobs.

- **Windows Application Log**: In this log, VSS Writers record detailed error messages that are not logged in the NetVault Backup logs. For more information, see http://msdn.microsoft.com/en-us/library/ms157312.aspx.

- **Vssadmin and DiskShadow:** You can use Vssadmin to list the VSS writers and providers, display their status, and ensure that the VSS Writers that the plug-in depends on are enabled and available. Vssadmin also indicates whether it has encountered any errors. DiskShadow, which is a newer utility, also provides information that might be useful. For more information, see http://technet.microsoft.com/en-us/library/bb491031.aspx and http://technet.microsoft.com/en-us/library/cc772172(v=WS.10).aspx.

  You can also obtain additional information from a command prompt by typing **vssadmin** or **diskshadow**.

Quest recommends that you use the Vssadmin and DiskShadow utilities to maintain the VSS subsystem, and to address specific issues, such as deleting obsolete shadow copies. You can use these utilities to list the shadow copies stored by the VSS, remove shadow copies, and manage the space allocated in a specific volume for the storage of shadow copies.

> **i** | **NOTE:** In addition to using these utilities to remove the shadow copies, Quest recommends that you restart the Volume Shadow Copy and the Microsoft Software Shadow Copy Provider services before resuming a VSS-backup job. In some instances, you might also have to restart the Microsoft Exchange Information Store service.

Addressing cluster-related issues

Clustered Exchange environments, such as DAG, present more challenges, especially when the plug-in tries to detect all the applicable NetVault Backup Clients or when the clients must communicate with each other. Usually, NetVault Backup scans the network to locate known NetVault Backup Clients. However, under certain conditions, such as when clients are in the same network but different segments or when clients are located in geographically distant networks, the NetVault Backup Clients might not be aware of each other. This issue can cause errors, such as the display of a "Failed to launch Proxy Master in *<NameOfNetVaultBackupClient>*" message in the binary log during a backup job. This message might indicate that the NetVault Backup Client that is running the backup job is not aware of the network location of the client that is specified in the message.

To work around this issue, you can use the **"machines.dat"** file to identify all clients that the NetVault Backup Client must be aware of. Each NetVault Backup Client contains this list of known NetVault Backup Clients in the NetVault Backup installation's **etc** directory. For example, in a typical Windows installation, the path is: **C:\Program Files (x86)\Quest\NetVault Backup\etc\machines.dat**. Adding a known NetVault Backup Client name as an IP address to this file lets the local NetVault Backup Client locate the remote NetVault Backup Client in the network.

If you encounter the "Failed to launch Proxy Master in *<NameOfNetVaultBackupClient>*" message, Quest recommends that you add all NetVault Backup Clients in an Exchange cluster to the specific NetVault Backup Client that runs the backups. In the following steps, the DAG contains three NetVault Backup Clients: Client_1, Client_2, and Client_3. The NetVault Backup Client that manages backups is Client_1, so the "machines.dat" file on Client_1 is the one that you would modify.

To add a client, perform the following steps.

1 Open the NetVault Backup Configurator, and select the **Service** tab.

2 On the **Service** tab, click **Stop Service**.

  NetVault Backup Services are stopped.

3 In Windows, navigate to the **etc** directory where NetVault Backup is installed, for example, C:\Program Files (x86)\Quest\NetVault Backup\….

4 To ensure that you can restore the original settings if necessary, make a backup copy of the **"machines.dat"** file, for example, "machines.dat_saved."

5 Use a text editor to open the original **"machines.dat"** file.

6 Verify that the local NetVault Backup Client, for example, Client_1, is already listed in the file.

```
[Client_1]
Networks=<IPaddressOfClient_1>
Preferred Address=<IPaddressOfClient_1>
Contact Address=<IPaddressOfClient_1>
```

7 At the end of the file, add a section for each additional client that the local NetVault Backup Client, for example, Client_1, must be aware of—replace *<IPaddressOfClient_x>* with the applicable IP address.

```
[Client_2]
Networks=<<IPaddressOfClient_2>
Preferred Address=<IPaddressOfClient_2>
Contact Address=<IPaddressOfClient_2>

[Client_3]
Networks=<IPaddressOfClient_3>
Preferred Address=<IPaddressOfClient_3>
Contact Address=<IPaddressOfClient_3>
```

8 Save the file.

9 Return to the **Service** tab in the NetVault Backup Configurator, and click **Start Service**.

The Current State changes to Running.

10 If the **Current State** remains **Running**, the changes have been implemented; otherwise, if the state changes to **Stopped** after a few seconds, complete the following steps:

    a Open the Windows Task Manager, and select the **Processes** tab.

    b To display the processes in alphabetical order, click the **Image Name** column heading.

    c To determine whether the **nvstatsmngr.exe** or **nvstatsmngr.exe*32** process is listed, check the list.

    d If either process is listed, select it, and click **End Process**.

    e If a confirmation prompt is displayed, click **End Process**.

    f Close the Task Manager, return to the NetVault Backup Configurator, and click **Start Service** again.

After these steps are completed, NetVault Backup automatically updates the information in the "machines.dat" file with additional information for each NetVault Backup client that you added.

For example:

```
[Client_2]
Type=xxxx
UDP Fragment Size=xxxx
Server=xxxx
Description=xxxx
NVVersion=xxxx
NVBuildLevel=xx
Networks=nnn.nnn.nnn.nnn
Preferred Address=nnn.nnn.nnn.nnn
Contact Address=nnn.nnn.nnn.nnn
Fixed entry=xxxx
Id=xx
Version=xxxxxxxx
Outside Firewall=xxxx
```

Troubleshooting other issues

This topic describes some common errors and their solutions. In cases where an error occurs and is not described in this table, obtain the Microsoft Exchange Server error number from the NetVault Backup logs, and then see the relevant Microsoft Exchange Server documentation for help.

Table 1. Troubleshooting

Error message	Explanation
Failed to Connect or Failed to log in as user *<userName>*	Ensure that all relevant Exchange Services are started before initiating backup or restore jobs. These services include: • **Microsoft Exchange Information Store** • **Microsoft Exchange System Attendant** • **Microsoft Software Shadow Copy Provider** (VSS backups) • **Volume Shadow Copy** (VSS backups)—started automatically by Microsoft Software Shadow Copy Provider
Failed to get Backup Components	Ensure that Exchange Server databases are mounted.
No Suitable Media for Job	This error occurs when multiple Mailbox Databases are backed up simultaneously by using parallel streams and the number of media drives is less than the number of Mailbox Databases.
An error occurred. Please check that databases are dismounted.	Even if a specific database has been dismounted, if the **This database can be overwritten by a restore** option is not selected on the **Database** tab of the **Store Properties** dialog box, this error occurs.
• Failed to add backup record • Failed to write index of backup to the database These messages indicate that the selected data was backed up, but the job's index information was not properly added by NetVault Backup to its database. Without this index information, the data cannot be properly restored.	**Method 1:** Open the **Manage Devices** page, select the backup media, and click **Scan**. NetVault Backup stores index information for backup jobs in two locations: in the NetVault Database and on the media targeted by the backup. When you scan the backup media, the index information is added to the NetVault Database. To verify that the information was added, open the **Manage Job Definitions** page, and locate the specific job. If you can run the job now, the scan process has corrected the problem. **Method 2:** If the scan has failed, run the backup job again.

Table 1. Troubleshooting

Error message	Explanation
Proxy Master on DAG member server *'<ServerName>'* died unexpectedly	NetVault Backup uses a threshold for NetVault Backup messages between NetVault Backup Client machines. By default, the threshold is 8 MB. For large Plug-in *for Exchange* backups that include several Exchange Server transaction log files, this activity might exceed the threshold. Exceeding the threshold causes the backup to fail and adds the specified error message in the NetVault Backup binary logs.

To increase the NetVault Backup messaging threshold:

1 For each NetVault Backup Client that is part of the DAG, access the NetVault Backup **config** directory, and use a text editor to modify the **"nvnmgr.cfg"** file.

2 At the end of the file, add the following entry and specify the value in bytes.

```
[Network:Message Validation Threshold]
Value=16777216
```

For example, to increase the level to 16 MB, add the following entry:

```
[Network:Message Validation Threshold]
Value=16777216
```

To increase the level to 32 MB, add the following entry:

```
[Network:Message Validation Threshold]
Value=33554432
```

3 Save the updated file.

4 Make the same change to the **"nvnmgr.cfg"** file for every NetVault Backup Client that is part of the DAG.

About us

Quest provides software solutions for the rapidly-changing world of enterprise IT. We help simplify the challenges caused by data explosion, cloud expansion, hybrid datacenters, security threats, and regulatory requirements. We are a global provider to 130,000 companies across 100 countries, including 95% of the Fortune 500 and 90% of the Global 1000. Since 1987, we have built a portfolio of solutions that now includes database management, data protection, identity and access management, Microsoft platform management, and unified endpoint management. With Quest, organizations spend less time on IT administration and more time on business innovation. For more information, visit www.quest.com.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at https://support.quest.com.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request.
- View Knowledge Base articles.
- Sign up for product notifications.
- Download software and technical documentation.
- View how-to-videos.
- Engage in community discussions.
- Chat with support engineers online.
- View services to assist you with your product.