



**SnapManager® 7.2 for Microsoft® Exchange Server**

# **Administration Guide**

January 2019 | 215-10626\_CO  
doccomments@netapp.com



# Contents

<b>Product overview .....</b>	<b>8</b>
<b>Backing up and verifying your databases .....</b>	<b>10</b>
SnapManager backup overview .....	10
SnapManager naming conventions .....	10
Defining a backup strategy .....	11
Backing up your databases for the first time .....	14
Verifying the initial backup set .....	15
Scheduling recurring backups .....	16
Scheduling frequent recovery point backups .....	17
Backup planning for databases in DAGs .....	17
Backing up using the Backup wizard .....	19
Creating backups using the Backup and Verify window .....	20
Using the remote additional copy backup feature .....	22
Configuring gapless DAG backups using remote additional copy .....	23
Managing backups using Remote Additional Copy Backup .....	24
control.xml file settings for remote additional copy backup .....	25
Using Database Filter to display databases to back up in a DAG .....	26
What the Activation Preference Number of a mailbox database is .....	27
Managing backup retention .....	27
Maximum number of Snapshot copies per volume .....	27
Criteria for deleting backups .....	28
Managing multiple backup management groups .....	29
Automatically deleting backups .....	30
Automatic deletion of orphaned SnapInfo folders .....	30
Explicitly deleting backups .....	31
Managing database verification .....	35
LUN requirements for verifying databases in a backup set .....	36
Database verification workload management .....	36
Backup verification status reporting .....	38
Where to run database and transaction log integrity verification .....	38
When to verify the databases in a backup set .....	38
Managing transaction log backups .....	39
Backup options .....	39
Triggering UTM retention .....	39
Option to back up transaction logs that Exchange will truncate .....	43
Frequent Recovery Point backup and restore operations .....	43
Changing the backup management group of an existing backup set .....	45
Reasons that a SnapManager backup might fail .....	46
Problem: cluster failover during backup .....	46
Problem: Snapshot copy limit reached .....	46
Problem: SnapInfo directory being accessed .....	46

Problem: SnapInfo directory out of space .....	47
Problem: data does not match .....	47
Problem: busy Snapshot copy .....	47
Problem: Snapshot copy already exists .....	47
Problem: out of disk space .....	47
Problem: SnapManager server initialization failed .....	47
Troubleshooting integrity verification failure on SnapMirror destination volumes .....	47
<b>Restoring databases .....</b>	<b>49</b>
When to choose SnapManager restore to recover Exchange mailbox databases .....	49
What SnapManager restore does .....	49
Types of SnapManager Restore operations .....	50
Snapshot copies created during a restore process .....	50
Methods that can decrease restore process time .....	50
Transaction log sequence verification options .....	51
LUN Clone Split Restore method .....	51
SnapManager restore operations in DAG environments .....	52
Guidelines for using SnapManager restore .....	53
Which type of restore operation to perform .....	53
Guidelines for restoring from a SnapManager backup copy .....	53
Restoring from a SnapManager backup copy .....	54
Decisions to make before restoring from a SnapManager backup copy .....	54
Restoring databases using the Restore wizard .....	55
Restoring databases using the Restore window .....	56
Restoring Exchange databases in a DAG .....	59
Backups available for restore in a DAG .....	59
Restoring a backup copy created on a different server .....	60
Reseeding a failed database copy in a DAG .....	61
Using a Recovery Database .....	62
Limitations of using a Recovery Database .....	62
Restoring a mailbox database to the Recovery Database in Exchange Server .....	63
When to delete a Recovery Database .....	64
Restoring mailboxes from unmanaged media .....	64
Troubleshooting restore operations .....	65
<b>Mailbox restore operation using Single Mailbox Recovery .....</b>	<b>66</b>
Recovering mailbox data .....	66
Cleaning up after Single Mailbox Recovery .....	67
<b>What SnapManager reports contain .....</b>	<b>68</b>
SnapManager reports in a DAG .....	68
Reasons to change the report directory location .....	68
Changing the SnapManager report directory .....	68
Locating the report directory in a DAG .....	69
Viewing SnapManager reports .....	69
Printing SnapManager reports .....	69

Deleting SnapManager reports .....	70
Control of report auto-archiving .....	70
<b>Modifying your database configuration on NetApp storage .....</b>	<b>71</b>
Moving Exchange data with the Configuration wizard .....	71
Moving an individual database to a LUN .....	72
Changing the location of a database .....	72
Viewing the full path for database files .....	73
Dataset protection policy .....	73
Migration of transport database paths for Exchange Server .....	74
Moving transaction logs to a LUN .....	74
Changing the location of transaction log files .....	75
Viewing the full path for transaction log files .....	75
Configuration using the SnapManager control file .....	75
Configuration of SnapManager in a DAG using a control file .....	76
Importing Exchange Server configuration information using the control file .....	76
Exporting Exchange Server configuration information using the control file .....	77
SnapManager control file XML schema .....	77
<b>Configuring SnapManager application settings .....</b>	<b>83</b>
Where to access SnapManager application settings .....	83
Adding Exchange servers to be managed .....	84
Enabling database migration back to local disks .....	84
Disabling database migration back to local disks .....	85
Considerations for selecting the database verification server .....	85
Configuring the verification server .....	85
Remote verification prerequisites .....	86
Selecting the Snapshot copy access method for database verification .....	87
Database verification throttling .....	88
Database verification throttling options .....	88
How database verification throttling works .....	88
Calculating the verification throttling sleep interval .....	89
Configuring database verification throttling .....	90
Throttling entries in the SnapManager backup and verification report .....	90
Impact of database verification on performance .....	90
Database verification override during restore operation .....	91
Configuring the database verification override option .....	91
Verification override entry in the SnapManager restore report .....	92
Configuring default settings for the Run Command After Operation option .....	92
Fractional space reservation .....	93
What can happen with a fractional-space-reserved volume .....	93
Fractional space reservation policies .....	94
Fractional space reservation policies to manage Exchange data .....	96
Viewing current fractional space reservation data for a LUN .....	98
Event notification options .....	99


Configuring automatic event notification settings .....	100
<b>Advanced administration .....</b>	<b>103</b>
Maximum configurations supported by SnapManager .....	103
Service account requirements for archiving backup sets with SnapVault (7- Mode environments only) .....	103
Concurrent backup verification .....	104
Managing integrity verification jobs .....	105
<b>Upgrading SnapManager .....</b>	<b>107</b>
Upgrading SnapManager interactively .....	107
Upgrading SnapManager from a command line .....	107
<b>Repairing, reinstalling, and uninstalling SnapManager .....</b>	<b>109</b>
Repairing SnapManager .....	109
Reinstalling SnapManager .....	109
Moving Exchange data to a local disk .....	110
Uninstalling SnapManager in interactive mode .....	110
Uninstalling SnapManager using the CLI .....	111
<b>SnapManager backup archiving .....</b>	<b>113</b>
Why organizations archive data .....	113
Guidelines for archiving SnapManager backups .....	113
Methods of archiving SnapManager backups .....	114
Archives created with NDMP or the dump command .....	114
Evaluation of the NDMP and dump command method of archiving .....	115
Example using Run Command After Operation feature to archive SnapManager backups .....	115
Archives created using an industry-standard backup utility .....	116
Evaluation of an industry-standard backup utility method of archiving ....	117
Example: Using an industry-standard backup utility to archive SnapManager backups .....	118
How to use the Exchange Backup Agent to archive Exchange backups .....	119
Evaluation of the Exchange Backup Agent method of archiving .....	119
Example: Using Exchange Backup Agent to archive Exchange backup copies .....	119
If you use a centralized backup model .....	120
Automatic backup archiving using the Run Command After Operation feature ....	120
Command arguments supported by the Run Command After Operation feature .....	120
Specifying the command to be run by the Run Command After Operation feature .....	121
Enabling the launch of SnapManager scripts from a UNC path .....	122
<b>SnapManager command-line reference .....</b>	<b>124</b>
Guidelines for using the SnapManager for Exchange PowerShell command-line interface .....	124
Launching SnapManager for Exchange PowerShell .....	124
new-backup .....	125
verify-backup .....	132

delete-backup .....	136
get-backup .....	138
restore-backup .....	140
reseed-database .....	145
Get-JobStatus .....	147
Change-JobPriority .....	148
Cancel-Job .....	149
Export-config .....	150
Import-config .....	151
<b>Copyright .....</b>	<b>154</b>
<b>Trademark .....</b>	<b>155</b>
<b>How to send comments about documentation and receive update</b>	
<b>notifications .....</b>	<b>156</b>
<b>Index .....</b>	<b>157</b>

## Product overview

SnapManager for Microsoft Exchange Server is a host-side component of the NetApp integrated storage solution for Microsoft Exchange, offering application-aware primary Snapshot copies of Exchange databases. You can use SnapManager with Data ONTAP SnapMirror technology to create mirror copies of backup sets on another volume, and with Data ONTAP SnapVault technology to archive backups efficiently to disk.

Together these tools offer a complete Snapshot-based data protection scheme that is as scalable, reliable, and highly available as the underlying storage system. The following illustration shows the components in a SnapManager deployment with clustered Data ONTAP:



### SnapManager highlights

SnapManager features seamless integration with Microsoft products on the Windows host and with NetApp Snapshot technology on the back end. It offers an easy-to-use, wizard-based administrative interface.

- *Integration with the Microsoft Volume Shadow Copy Service (VSS)* ensures that write requests are frozen and write caches are flushed before backups are taken. SnapManager supports Windows Volume Manager, Windows Server Failover Clustering, Microsoft Multipath I/O (MPIO), and Exchange Database Availability Groups.
- *Fast, nondisruptive Snapshot technology* using NetApp SnapDrive for Windows software enables you back up databases in seconds and restore them in minutes without taking Exchange Servers or databases offline. Snapshot copies consume minimal storage space. You can store up to 255 copies per volume.
- *Automated central administration* offers hands-off, worry-free data management. You can schedule routine Exchange Server database backups, configure policy-based backup retention, set up point-in-time and up-to-the-minute restore operations and proactively monitor your Exchange Server environment with periodic email alerts. PowerShell cmdlets are available for easy scripting of backup and restore operations.

In addition to these major features, SnapManager offers the following:

- Integrated Single Mailbox Recovery enables you to restore individual mailboxes, email messages or attachments, calendar items, deleted items, drafts, or contacts (Single Mailbox Recovery must be installed separately)
- Simplified migration of existing databases to NetApp storage with an easy-to-use Configuration wizard

- Nondisruptive, automated backup verification
- Fast reseeding of databases in a Database Availability Group
- Support for physical and virtualized infrastructures
- Support for iSCSI, Fibre Channel, and FCoE
- Support for service-level Role-Based Access Control (RBAC)

SnapManager also offers support for Data Availability Groups (DAGs) without a Cluster Administrative Access Point (AAP) running on Exchange Server 2016/2013 environments.

Microsoft Exchange Server 2013 SP1 (and later) on Windows Server 2012 R2 allows you to create a DAG without an AAP, also known as an IP-less DAG or Active Directory-detached cluster. This simplifies the creation and management of DAGs by reducing resource requirements and removing the need for Exchange to create objects in the Active Directory. The operational benefits to creating a DAG without an AAP are that the cluster group contains only the File Share Witness resource and there is no need to provision permissions. This can be beneficial when delegating administration to different groups.

A DAG without an AAP is a Windows cluster with the following features:

- No IP address assigned to the cluster or DAG, and therefore no IP address resource is assigned in the cluster core resource group
- No network name assigned to the cluster, and therefore no network name resource is assigned in the cluster core resource group
- No cluster name object (CNO) in the Active Directory

The name of the cluster or DAG is not registered in the DNS and is not resolvable over the network. A DAG without an AAP cannot be managed using the Failover Cluster Management tool. It must be managed using Windows PowerShell, and PowerShell cmdlets must be run against individual clusters.

## Backing up and verifying your databases

---

You should back up your databases as soon as they are available in NetApp storage. You can then verify the initial backups and schedule recurring backups and recurring backup verifications.

**Note:** For information on how to install SnapManager on Windows hosts and how to set up NetApp storage for SnapManager usage, see the *SnapManager for Microsoft Exchange Server Installation and Setup Guide*.

### Related information

[\*SnapManager 7.2 for Microsoft Exchange Server Installation and Setup Guide for Data ONTAP operating in 7-Mode\*](#)

[\*SnapManager 7.2 for Microsoft Exchange Server Installation and Setup Guide For Clustered Data ONTAP\*](#)

## SnapManager backup overview

SnapManager uses NetApp Snapshot technology to create online, read-only copies of databases. It uses an Exchange Server tool to verify the integrity of the backups.

SnapManager backs up a database by creating Snapshot copies of the volumes in which the following reside:

- Database data files
- SnapInfo directories and transaction logs

Together these Snapshot copies comprise a *backup set*. SnapManager uses a backup set to restore a database.

After SnapManager backs up your databases, it can perform an integrity verification of the backup sets. SnapManager uses the Exchange System Management Tools to check the database and transaction log files for physical and logical corruption. Verification ensures that you can use backup sets to restore databases as needed.

**Note:** Database verification is disabled by default if you have a Database Availability Group (DAG). Verification is not required for DAG databases that have at least two copies, each of which has a viable backup. For more information, see [\*NetApp KB Article 3013600: Does verification in SME need to occur in an Exchange 2010 DAG deployment?\*](#)

**Important:** SnapManager cannot restore databases from Snapshot copies created by Data ONTAP or SnapDrive. You should perform backups using SnapManager only.

## SnapManager naming conventions

SnapManager offers both unique and generic naming conventions for naming Snapshot copies. The unique naming convention contains the variable *date\_time* in the name, and the generic naming convention includes the string *recent* in the name of the most recent Snapshot copy.

When you use the unique naming convention, the most recent Snapshot copy is identified by the most recent date and time. SnapManager does not rename this Snapshot copy when the next Snapshot copy is created.

When you use the generic backup naming convention, the most recent Snapshot copy is identified by the Snapshot copy name that includes the string `recent`. This is the naming convention older versions of SnapManager use and is the default setting to enable backward compatibility.

When you have datasets configured in your system and you choose the generic naming convention, no archives are created. To create archives, apply the unique naming convention with the archival process enabled. If you archive the backups using PowerShell, SnapManager changes the generic naming convention to the unique naming convention. In the GUI, you can select the backup naming convention in the Backup Settings dialog box.

**Note:** Select the unique naming convention unless you have legacy scripts that need a Snapshot copy with `recent` in its name.

## Defining a backup strategy

Defining a backup strategy before you create your backup jobs helps ensure that you have the backups that you need to successfully restore your databases. Your Service Level Agreement (SLA) and Recovery Point Objective (RPO) largely determine your backup strategy.

**Note:** For SnapManager best practices, see [NetApp Technical Report 4221: Microsoft Exchange Server 2013 and SnapManager for Exchange Best Practices Guide for Clustered Data ONTAP](#) or [NetApp Technical Report 4224: Microsoft Exchange Server 2013 and SnapManager for Exchange Best Practices Guide for Data ONTAP Operating in 7-Mode](#).

### What type of SnapManager backup do you need?

SnapManager supports two types of backups:

Backup type	Description
Database backup	<p>You can choose from two database backups:</p> <ul style="list-style-type: none"> <li>• Full backup Backs up database files and truncated transaction logs. Exchange Server truncates transaction logs by removing entries already committed to the database. This is the most common backup type.</li> <li>• Copy backup Backs up database files and transaction logs that have not been truncated. Use this backup type when you are backing up your databases with another backup application. Keeping transaction logs intact ensures that any backup application can restore the databases.</li> </ul>
Frequent recovery point backup (FRP)	<p>Backs up truncated transaction logs, copying only transactions committed since the most recent full backup or FRP backup.</p> <p>If you schedule frequent recovery point backups alongside database backups, SnapManager can restore databases to a specific recovery point more quickly. For example, you might schedule database backups at the start and end of the day and frequent recovery point backups every hour.</p> <p><b>Note:</b> SnapManager does not verify transaction logs when it creates frequent recovery point backups. SnapManager verifies the backups when it verifies the backup sets created from database backups.</p>

### When should you back up your databases?

The most critical factor for determining a database backup schedule is the rate of change for the database. You might back up a heavily used database every hour, while you might back up a rarely

used database once a day. Other factors include the importance of the database to your organization, your Service Level Agreement (SLA) and your Recover Point Objective (RPO).

Even for a heavily used database, there is no requirement to run a full backup more than once or twice a day. Regular transaction log backups are usually sufficient to ensure that you have the backups you need.

**Tip:** The more often you back up your databases, the fewer transaction logs SnapManager has to play forward at restore time, which can result in faster restores.

**Important:** SnapManager can perform one operation at a time. Do not schedule overlapping SnapManager operations.

### When should you verify backup copies?

Although SnapManager can verify backup sets immediately after it creates them, doing so can significantly increase the time required to complete the backup job. It is almost always best to schedule verification in a separate job at a later time. For example, if you back up a database at 5:00 p.m. every day, you might schedule verification to occur an hour later at 6:00 p.m.

For the same reason, it is usually not necessary to run backup set verification every time you perform a backup. Performing verification at regular but less frequent intervals is usually sufficient to ensure the integrity of the backup set. A single verification job can verify multiple backup sets at the same time.

**Note:** Verification is not required for DAG databases that have at least two copies, each of which has a viable backup. For more information, see [NetApp KB Article 3013600: Does verification in SME need to occur in an Exchange 2010 DAG deployment?](#)

**Important:** SnapManager can perform one operation at a time. Do not schedule overlapping SnapManager operations.



### How many backup jobs do you need?

You can back up your databases using one backup job or several. The number of backup jobs that you choose typically mirrors the number of volumes on which you placed your databases. For example, if you placed a group of small databases on one volume and a large database on another volume, you might create one backup job for the small databases and one backup job for the large database.

Other factors that determine the number of backup jobs that you need include the size of the database, its rate of change, and your Service Level Agreement (SLA).

### Which backup naming convention do you want to use?

A backup naming convention adds a string to Snapshot copy names. The string helps you identify when the copies were created. There are two naming conventions:

Naming convention	Description
Unique	<p>Adds a time stamp to all Snapshot copy names.</p> <p>This is the default option.</p> <p>Example:</p> 
Generic	<p>Adds the string "recent" to the name of the most recent Snapshot copy.</p> <p>All other Snapshot copies include a time stamp.</p> <p>Example:</p> 

The selected naming convention applies to all backups. You should use the unique naming convention unless you have a script that requires the constant string "recent."

You can change the naming convention in the **Backup Settings** dialog box.

### Which backup management group do you want to assign to the backup job?

You select a backup management group to apply a labeling convention to Snapshot copies. When you back up a database, you can choose from three management groups:

Management group	Description
Standard	Does not include the name of the management group in Snapshot copy names.  Example:  exchsnap_E13SP1WK8S1_07-02-2014_10.45.13
Daily	Adds "Daily" to Snapshot copy names.  Example:  exchsnap_E13SP1WK8S1_07-02-2014_10.53.59_Daily
Weekly	Adds "Weekly" to Snapshot copy names.  Example:  exchsnap_E13SP1WK8S1_07-02-2014_10.59.15_Weekly

For example, if you schedule daily and weekly backups, you should assign the backups to the Daily and Weekly management groups, respectively.

**Note:** Management groups do not enforce a backup schedule.

### How long do you want to retain backup copies on the source storage system and the SnapMirror destination?

You can choose either the number of days you want to retain backup copies, or specify the number of backup copies you want to retain, up to 255. For example, your organization might require that you retain 10 days worth of backup copies.

If you set up SnapMirror replication, the retention policy is mirrored on the destination volume.

**Note:** For long-term retention of backup copies, you should use SnapVault.

### How long do you want to retain transaction log backups on the source storage system?

SnapManager needs transaction log backups to perform *up-to-the-minute restores*, which restore your database to a time between two full backups. For example, if SnapManager took a full backup at 8:00 a.m. and another full backup at 5:00 p.m., it could use the latest transaction log backup to restore the database to any time between 8:00 a.m. and 5:00 p.m. If transaction logs are not available, SnapManager can perform *point-in-time restores* only, which restore a database to the time that SnapManager completed a full backup.

Typically, you might require up-to-the-minute restores for only a day or two, which means you would retain transaction log backups for one or two days.

### Do you want to verify backup copies using the source volume or a destination volume?

If you use SnapMirror or SnapVault, you can verify backup copies using the Snapshot copy on the SnapMirror or SnapVault destination volume, rather than the Snapshot copy on the primary storage system. Verification using a destination volume reduces load on the primary storage system.

### If you need to create backups using another tool, what backup type should you use?

If you need to create backups using another backup tool, create copy or differential backups only with that tool. Normal (full) and incremental backups truncate transaction logs, effectively disabling SnapManager up-to-the-minute restores.

## Backing up your databases for the first time

After you migrate your databases to NetApp storage, you should back them up immediately. You can schedule recurring backups after the initial backup and verification.

### About this task

These steps show you how to quickly back up your databases using the Backup and Verify option. You can use the Backup wizard if you prefer.


If you have a Database Availability Group (DAG), you can back up all databases in the group by running the backup job from the DAG. You cannot back up at the DAG level if either of the following are true:

- You are archiving databases to SnapVault secondary volumes that use Data ONTAP operating in 7-Mode.
- There are nodes in the DAG that use non-NetApp storage.

You must back up the databases from each DAG member instead.

### Steps

1. In the **Console Root** tree, expand the server on which the databases reside and click **Backup**.



2. In the **Backup** pane, select the databases that you want to backup.



3. In the **Actions** pane, click **Backup and Verify**.



4. In the **Backup** dialog box, keep **Create backup** selected and define the properties for the backup job:

For this field...	Do this...
Delete backup	Specify a retention policy for backup copies on the source storage system by defining the number of backup copies to retain or the number of days to retain backup copies.
Up-to-minute Restore options	Click this field and then specify a retention policy for transaction logs.
Verify backed up databases and transaction logs	Clear this field because it is best to verify databases in a separate operation.
Backup management group	Select a management group.
Run command after operation	If you want to run a command after the backup operation, select this field. You specify the command after you click <b>Backup</b> .
Backup archiving options	If you set up a SnapVault destination volume, select the option to archive the backup copy to the destination volume.
SnapMirror options	If you set up a SnapMirror destination volume, select the option to replicate the backup copy to the destination volume.
Advanced Options	If you back up your databases using another backup application, click this field, then click <b>Backup Type</b> , and finally select <b>Copy backup</b> . Keep the default selections for other fields in the advanced options.

5. Click **Backup**.

6. If you chose to run a command after the operation, specify where to run the command, the path to the program or script, the SnapManager variables to execute, and the command arguments, and then click **OK**.

7. In the Backup Status dialog box, click **Start Now**.

You can view details of the operation in the Backup Task List and Backup Report tabs.

## Verifying the initial backup set

You should verify an initial backup set to confirm the integrity of the databases.

### Steps

1. In the **Backup** pane, select the databases that you want to include in the backup verification schedule.
2. In the **Actions** pane, click **Backup and Verify**.
3. In the **Backup** dialog box, select **Verify backup sets** and then define the properties for the backup verification:

For this field...	Do this...
Select the number of recent unverified backups to be verified	Keep the default. You should have only one backup set at this point.
Backup management group	Select a management group.
Run command after operation	If you want to run a command after the backup operation, enable the field. You specify the command after you click <b>Verify</b> .

For this field...	Do this...
Backup archiving options	If you archived the backup set to a SnapVault destination volume and you want to verify the backup set on the destination storage system to reduce load on the primary storage system, click <b>Verify on previously archived backup</b> .
SnapMirror options	If you replicated the backup set to a SnapMirror destination volume and you want to verify the backup set on the destination storage system to reduce load on the primary storage system, click <b>Verify on available SnapMirror destination volumes</b> .

4. Click **Verify**.
5. If you chose to run a command after the operation, specify where to run the command, the path to the program or script, the SnapManager variables to execute, and the command arguments, and then click **OK**.
6. In the Backup Status dialog box, click **Start Now**.

You can view details of the operation in the Backup Task List and Backup Report tabs.

## Scheduling recurring backups

You can schedule recurring backup jobs using Windows Scheduled Tasks.

### Steps

1. In the **Backup** pane, select the databases that you want to include in the backup schedule.
2. In the **Actions** pane, click **Backup and Verify**.
3. In the **Backup** dialog box, keep **Create backup** selected and define the properties for the backup schedule, as described in [Backing up your databases for the first time](#) on page 14.
4. Click **Schedule**.
5. In the **Schedule** dialog box, enter a job name, enter your credentials, and click **OK**.
6. Create the schedule using Windows Scheduled Tasks:
  - a. Click **Schedule**.
  - b. Specify the schedule.
  - c. Click **OK**.
  - d. Click **Yes** to save the job.

### After you finish

You can view details about the backup job in the SnapManager Scheduled Jobs pane.

## Scheduling frequent recovery point backups

Frequent recovery point backups help you preserve transaction logs. You should schedule transaction log backups alongside full database backups at a frequency that allow you to meet your Recovery Point Objective (RPO).

### Steps

1. In the **Console Root** tree, select the server on which the databases reside.
2. In the **Actions** pane, click **Frequent Recovery Point Backup**.
3. In the **Frequent Recovery Point Backup** dialog box, define the properties for the transaction log backup schedule:
 

For this field...	Do this...
Selected for backup	Select the databases that you want to include in the backup schedule.
Backup frequency	Specify how often you want to back up the transaction logs.
Run command after operation	If you want to run a command after the backup operation, select this field. You specify the command after you click <b>Create Job</b> .
Update SnapMirror after operation	If you set up a SnapMirror destination volume, select this field to replicate the backup copy to the destination volume.
4. Click **Create Job**.
5. Create the schedule using **Windows Scheduled Tasks**:
  - a. Click **Schedule**.
  - b. Specify the schedule.
  - c. Click **OK**.
  - d. Click **Yes** to save the job.

## Backup planning for databases in DAGs

By planning the backup strategy for your Exchange databases, you can minimize the chances of losing data if a restore operation is necessary, while still controlling the resources needed to create and maintain the backups.

The strategy you choose depends, in part, on the number of Database Availability Group nodes, the number of databases, the size of the databases, the network links between the DAG nodes, the service level agreements, and how quickly you must be able to reseed a database:

- If your DAG resides on nine nodes, you can choose a gapless backup strategy based on the SnapManager for Exchange remote additional copy backup feature, which greatly simplifies backup administration.
- If the DAG is larger than nine nodes, choose a server-oriented backup strategy; depending on your recovery requirements, you might add frequent recovery point (FRP) backups to the server-oriented backups.

### Planning backups for nine-node DAGs

For a DAG consisting of nine nodes, you can use SnapManager remote additional copy backup to simplify backup planning. These plans still permit up-to-the-minute database restore and reseed

operations for any database in the DAG, independent of the availability of any specific node in the group. However, the database layout plays a major role in the backup completion time. Make sure that the databases are placed in an optimal fashion to achieve the best results while using gapless backup.

The following table provides an example of an optimal layout:

DAG1									
Server1	Server2	Server3	Server4	Server5	Server6		Server7	Server8	Server9
Active							Passive		
DB1	DB10	DB19	DB28	DB37	DB46		DB1	DB19	DB37
DB2	DB11	DB20	DB29	DB38	DB47		DB2	DB20	DB38
DB3	DB12	DB21	DB30	DB39	DB48		DB3	DB21	DB39
DB4	DB13	DB22	DB31	DB40	DB49		DB4	DB22	DB40
DB5	DB14	DB23	DB32	DB41	DB50		DB5	DB23	DB41
DB6	DB15	DB24	DB33	DB42	DB51		DB6	DB24	DB42
DB7	DB16	DB25	DB34	DB43	DB52		DB7	DB25	DB43
DB8	DB17	DB26	DB35	DB44	DB53		DB8	DB26	DB44
DB9	DB18	DB27	DB36	DB45	DB54		DB9	DB27	DB45
Passive							Passive		
DB10	DB1	DB28	DB19	DB46	DB37		DB10	DB28	DB46
DB11	DB2	DB29	DB20	DB47	DB38		DB11	DB29	DB47
DB12	DB3	DB30	DB21	DB48	DB39		DB12	DB30	DB48
DB13	DB4	DB31	DB22	DB49	DB40		DB13	DB31	DB49
DB14	DB5	DB32	DB23	DB50	DB41		DB14	DB32	DB50
DB15	DB6	DB33	DB24	DB51	DB42		DB15	DB33	DB51
DB16	DB7	DB34	DB25	DB52	DB43		DB16	DB34	DB52
DB17	DB8	DB35	DB26	DB53	DB44		DB17	DB35	DB53
DB18	DB9	DB36	DB27	DB54	DB45		DB18	DB36	DB54

For more information about using remote additional copy backup for a DAG, see [Remote additional backup copy](#) on page 22.

If “cluster group” failover happens and if DAG-level gapless backup was started before failover, then backup can possibly hang and time out after 18 hours (default). Best practice is to not start cluster failover when backups are running, but only after or before the backup is completed.

**Note:** Exchange 2013 CU1 supports only 50 databases per node and Exchange 2013 CU2 supports 100 databases.

### Planning backups for DAGs with more than nine nodes

Alternatively, you can take a node level approach, as follows:

1. 00:30 to 02:00 full backup of Copy 4
2. 02:00 to 03:00 copy backup of Copy 3

3. 03:00 to 04:00 copy backup of Copy 1 and, simultaneously, copy backup of Copy 2

Depending on your service level requirements, this schedule might be sufficient. If it is not, rather than repeat this schedule, you can add FRP backups. FRP backups are very quick, because they are Snapshot copies of the transaction logs. Adding the FRP backups, the modified schedule might be something like the following:

1. 00:30 to 02:00 full backup of Copy 4
2. 02:00 to 03:00 copy backup of Copy 3
3. 03:00 to 04:00 copy backup of Copy 1 and copy backup of Copy 2
4. 03:00 to 04:00 FRP backup of Copy 3 and FRP backup of Copy 4
5. 05:00 to 00:00 FRP backups of all copies every 30 minutes

For more information about configuring frequent recovery point backups, see [Frequent Recovery backup operations](#) on page 44.

## Backing up using the Backup wizard

You can back up databases by using the SnapManager Backup wizard. Creating a backup enables you to store and recover your database.

### About this task

When performing a backup, consider the following points:

- Do not move the cluster group to another node in the DAG while a backup is running.
- If OnCommand Unified Manager Core Package integration is required, you must place the LUNs for your database, logs, and SnapInfo in a qtree.
- Each mailbox database in Exchange Server requires a dedicated LUN to be a valid target for SnapManager Backup.

### Steps

1. In the **Scope** pane, select the Exchange server node that you want to back up.
2. Click **Backup**.

SnapManager displays the list of databases in the Backup view in the Result pane, and the corresponding actions that you can perform for SnapManager backup in the Actions pane.

**Note:** If the SnapManager MMC snap-in is connected to a DAG, you must use the Database Filter to specify the criteria to display the Exchange databases for backup.

3. Select the databases that you want to back up.

The Result pane shows whether the database is dataset-enabled, the name of the enabled dataset, the SnapMirror status, the SnapVault status, and other details about the database.

4. Click **Backup Wizard** in the **Actions** pane.

The Welcome window appears.

5. Follow the instructions in the **Backup** wizard to initiate a backup process.
6. In the **Completing the Backup wizard** dialog box, click **Finish** after you verify that all the settings in the window are what you want.

The Backup Status window appears.

7. Click **Start Now** to start the backup.

The backup is performed and the Snapshot copy is written to the volume. SnapManager Backup completes each task and checks it off on the list shown in the **Backup Task List** view. You can alternate between the task check-off list and the progress report. If the backup operation is successful, the Task view shows the check-off list with the tasks completed.

**Note:** If you enable the **Notification** option, an email message is sent and the event is posted to the Windows event log.

## Creating backups using the Backup and Verify window

You can use the Backup and Verify window to create a backup copy more quickly than by using the Backup wizard.

### About this task

You can choose to archive your database to a secondary storage system and when to verify the archived backup.

When performing a backup, consider the following points:

- Do not move the cluster group to another node in the DAG while a backup is running.
- If OnCommand Unified Manager Core Package integration is required, you must place the LUNs for your database, logs, and SnapInfo in a qtree.
- Each mailbox database in Exchange Server requires a dedicated LUN to be a valid target for SnapManager Backup.

### Steps

1. In the **Actions** pane, select **Backup**.
2. In the **Backup and Verify** window, select the databases that you want to back up.
 

**Note:** If you decide to back up a database with an invalid status, SnapManager does not allow you to proceed with the backup operation.
3. In the **Backup Management Group** list, select a management group for the backup copy that you want to create.
4. If you set up a SnapVault destination volume, select a backup archiving option under **Backup archiving options** to archive the database at the secondary storage system. If you set up a SnapVault destination volume,
5. If you set up a SnapVault destination volume, select a backup retention group under **Backup archiving options** to determine the retention time of the dataset on the archived secondary storage system.
6. If you want to delete older backup copies of this backup management group automatically, select one of the **Delete Backups** options.
 

**Note:** Ensure to select one of the **Delete Backups** options to manage your current Snapshot copies.
7. Click **Advanced** to view or change the setting to retain up-to-the-minute restore capability.

If you want to...	Then do this...
<b>Avoid creating a break in the continuity of transaction logs (between the previous backup and the present time)</b>	Select the <b>Retain up-to-the-minute restore ability for older backups in other management groups</b> option.  Snapshot copies of LUNs that contain transaction logs related to any management group that is not selected are not deleted from the SnapInfo directory.
<b>Reduce the space consumption of the storage system by transaction logs (by allowing more transaction logs to be deleted)</b>	Clear the <b>Retain up-to-the-minute restore ability for older backups in other management groups</b> option.

8. Do the following if you want to view or change backup transaction log settings:

If you want to...	Then do this...
<b>Retain up-to-the-minute restore ability from a backup that is not the most recent</b>	Select <b>Back up transaction logs that will be truncated by Exchange at the end of the backup</b> .  This option is automatically selected each time you launch SnapManager.
<b>Reduce the amount of disk space consumed by transaction log backups that Exchange truncates after the backup finishes</b>	Clear the <b>Back up transaction logs that will be truncated by Exchange at the end of the backup</b> option. The option remains disabled only until you exit SnapManager.  <b>Note:</b> Clearing this option causes SnapManager Backup to also remove the option to defer database verification.

9. Select **Verify Backed Up Databases and Transaction Logs** to verify the backup immediately after the backup is complete.

**Note:** Although it is possible to restore from an unverified backup, you should not do so.

10. If you want to run a command after the backup process is complete, select the **Run Command After Operation** check box.

You specify where to run the command, the path to the program or script, the SnapManager variables to execute, and the command arguments in the Run Command After Operation dialog. The dialog opens after you click **Backup Now** or **Schedule**.

11. Depending on your requirements for SnapMirror updates, do either of the following:

If...	Then do this...
<b>Your volume is a SnapMirror source volume and you do not want the destination volume to be updated after the backup process is complete</b>	Clear the <b>Update SnapMirror After Operation</b> check box.
<b>Your volume is a SnapMirror source volume and you want the destination volume to be updated after the backup process is complete</b>	Select the <b>Update SnapMirror After Operation</b> check box.

12. Select **Verify on available SnapMirror destination volumes** to run integrity verification on the Exchange databases and transaction logs that are stored on the destination volume.

The **Verify on available SnapMirror destination volumes** is available and checked by default, only when the selected databases have at least one SnapMirror relationship.

13. Depending on when you want to run the backup process, perform one of the following:

If you want to...	Then do this...
<b>Run your backup process immediately</b>	Click <b>Backup Now</b> , and then proceed to step to back up the selected databases.  The Backup Status window appears, showing the Task list.
<b>Defer your backup process</b>	Click <b>Schedule</b> to use the Windows Task Manager to schedule your backup process, and then proceed to the next step.

14. In the **Schedule Jobs** window, name your backup job, provide the user ID and password for the job, and then click **OK**.

**Note:** If this name already exists as a Windows scheduled task and you want to replace it with a new job, select the **Replace if it Exists** check box, and then click **OK**.

The Schedule jobs window appears.

15. In the **Schedule Jobs** window, use the **Schedule** tab to specify when the job is to run and, if you want the job to repeat, at what frequency.

16. After you schedule your job, click **OK**.

You can use Control Panel to modify the schedule or cancel the scheduled job.

The backup job runs at the times you specified in the **Schedule Task** view.

17. In the **Backup Status** window, click **Start Now** to back up the selected databases.

The backup operation finishes and the Snapshot copy is written to the volume. SnapManager Backup completes each task and checks it off on the list shown in the **Backup Task List** view.

You can alternate between the task check-off list and the progress report.

If the backup process is successful, the Task view shows the check off list with the tasks completed.

**Note:** If you enable Notification, an email message is sent and the event is posted to the Windows Application event log.

## Using the remote additional copy backup feature

For DAGs, you can create a simpler, easier backup strategy than is needed for larger systems by using the SnapManager for Exchange remote additional copy backup feature. The simplified plan still permits up-to-the-minute restore operations, even if a node is unavailable.

With remote additional copy backup enabled, you can do the following:

- Set up gapless Database Availability Group backups by requesting a copy of the databases to be backed up using the full backup procedure.

You should not select duplicated copies of the same database for gapless backups. Although duplicate gapless backups work, typically they take longer than necessary to create several redundant gapless backups. You can select a single copy of each database for the gapless backup using one of the following methods:

- Using the tree view, select a single copy of each database.
- Name each single database copy using a PowerShell `-database` parameter.

Alternatively, you can optimize the list of copies with the following database filter options:

- **Only active databases**
- **Only passive databases**
- **With specific activation preference for backup**
- Enable gapless DAG backup by selecting **Create a copy database backup on other DAG nodes along with current backup**.
- Select on which nodes the copy-only backups are to be performed.
- For a copy-only backup on other nodes, select the backup retention policy and request that the policy be executed after the copy backup is created.

Advantages of using gapless DAG backup include:

- If you perform a gapless DAG backup, do not also perform a primary backup of each copy of the database; otherwise, redundant and unnecessary backups are created.
- You can perform a copy backup on active as well as passive databases.

SnapManager performs a remote additional copy backup by performing the following:

1. Creates copy-only backups on those nodes selected for copy backup
2. Creates a full backup on the selected databases
3. Backs up the logs on the selected nodes
4. Requests that Microsoft Exchange truncate the log on the primary backup node
5. Replicates the truncation to the remote copies of the database

## Configuring gapless DAG backups using remote additional copy

On DAGs, you can use a simplified procedure to back up the databases while preserving your ability to perform up-to-the-minute restore operations, even if a node becomes unavailable.

### Steps

1. Connect directly to the DAG server.  
You cannot use a remote additional backup copy if you directly connect to a member server.
2. Select the databases for which you want to create a backup.  
Do not select multiple copies of the same database for backup.
3. Using the **Advanced Options** window, click the **Remote Additional Copy Backup** tab.
4. Select the **Create additional backup (copy based) on DAG remote server(s)** check box.  
**Note:** If you schedule a gapless backup with the verification option selected, only the cluster owner node copy is verified. The other copies are unverified.
5. (Optional) Select either one or both of following options:
  - **Delete backups during additional backups on remote servers**
  - **If the dataset is enabled, archive backups created on remote servers to secondary storage**
6. Using the **Advanced Backup Options** dialog box, select on which server you want to run copy-based additional remote copy backups.  
For best results, select all DAG nodes (this is the default).

**Note:** Delete backups during additional backups on remote servers can be achieved by using either of the following methods:

- Using In Excess of numbers
- Using Older than days

Remote backup inherits the same Management Group that was selected in the Backup wizard for local backups.

For example, in a gapless backup at DAG level, if you select **Daily** as the Management Group, then this setting applies to all of the remote additional copy backup databases (passive databases) available on the other nodes.

Management Group	<b>N1</b> Active database	<b>N2</b> Passive database - Additional copy backup	<b>N3</b> Passive database - Additional copy backup
Backup with <b>Daily</b> selected	Daily	Daily	Daily
Backup with <b>Weekly</b> selected	Weekly	Weekly	Weekly
Backup with <b>Standard</b> selected	Standard	Standard	Standard

#### Backup grouping (sub-jobs)

When gapless backup is scheduled from the GUI and before the backup runs, the databases are grouped before arriving at the number of backup groups created based on the following criteria:

- How many databases are existing per node based on the database type and in what order (either Active or Passive copy, but not both)?
- How many Passive database copies are selected for the remote additional copy backup feature?
- How many backup target servers are selected for the remote additional copy backup feature?
- Only limited set of LUNs can be added to take VSS-based snapshot.

For example, for the following database layout, the number of Backup groups created for gapless backup = 9 (Backup sub-jobs).

DAG Nodes		
<b>N1</b>	<b>N2</b>	<b>N3</b>
DB1 (Active)	DB1*	DB1*
DB2*	DB2 (Active)	DB2*
DB3*	DB3*	DB3 (Active)

\* = Passive

#### Managing backups using Remote Additional Copy Backup

You can use the PowerShell `new-backup` cmdlet with various parameters to create and manage backups, using the remote additional copy backup feature. For example, you can specify the number

of backup copies to keep on remote nodes, how long backup copies are kept, and whether they are archived.

For details of the `new-backup` command and the following parameters associated the remote additional backup feature, see the topic *new-backup* in the *SnapManager command-line reference* section:

- `RemoteAdditionalCopyBackup`
- `AdditionalCopyBackupDAGNode`
- `RetainRemoteAdditionalCopyBackup`
- `RetainRemoteAdditionalCopyBackupDays`
- `ArchiveRemoteAdditionalCopyBackup`

## control.xml file settings for remote additional copy backup

You can control various aspects of remote additional copy backup by including certain settings in the `control.xml` file. For example, you can control the number of backup copies kept on a remote node, and how long they are kept.

If you want to do this...	Then use this setting...
Control the deletion of backups on remote nodes	<pre>&lt;DELETE_BACKUPS_OPTION_REMCOPYBACKUP&gt;XXXX&lt;/ DELETE_BACKUPS_OPTION_REMCOPYBACKUP&gt;</pre> <p>integer</p> <p>0=no backups are deleted on the remote nodes, 1=backups are deleted on the remote nodes</p>
Control the number of backups to keep on the remote node after a gapless backup is done	<pre>&lt;BACKUP_SET_TO_KEEP_REMCOPYBACKUP&gt; XXXX&lt;/ BACKUP_SET_TO_KEEP_REMCOPYBACKUP&gt;</pre> <p>integer</p> <p>Equals number of backups to be kept after a gapless backup is done</p>
Control how long (days or part days) to keep backups after a gapless backup is done	<pre>&lt;BACKUP_SET_TO_KEEP_IN_DAYS_REMCOPYBACKUP&gt; XXXXX.XXXX &lt;/ BACKUP_SET_TO_KEEP_IN_DAYS_REMCOPYBACKUP&gt;</pre> <p>real-float</p> <p>Equals days or partial days to keep backups after a gapless backup is done</p>
Define the member servers on which remote copy backups can be done	<pre>&lt;HOSTS_REMCOPYBACKUP&gt; SERVERNAME1, SERVERNAME2, ...&lt;/ HOSTS_REMCOPYBACKUP&gt;</pre> <p>string</p> <p>Comma-delimited list of member servers on which to perform remote copy backup; if none is specified, all cluster nodes are processed</p>
Control whether additional remote additional copy backups are done	<pre>&lt;REQUEST_REMCOPYBACKUP&gt;xxxx&lt;/REQUEST_REMCOPYBACKUP&gt;</pre> <p>integer</p> <p>0=no remote additional copy backup is done, 1=remote additional copy backup is done</p>

### Example control.xml fragment for remote additional copy backup control

```

<BACKUP>
  <BACKUP_CLIENT_SETTING>
    <DELETE_BACKUPS_OPTION_REMCOPYBACKUP>1</
    <DELETE_BACKUPS_OPTION_REMCOPYBACKUP>
      <BACKUP_SET_TO_KEEP_REMCOPYBACKUP>8</
    <BACKUP_SET_TO_KEEP_REMCOPYBACKUP>
      <BACKUP_SET_TO_KEEP_IN_DAYS_REMCOPYBACKUP>7.0308
        </BACKUP_SET_TO_KEEP_IN_DAYS_REMCOPYBACKUP>
        <HOSTS_REMCOPYBACKUP>SNAPMGR-11-VM30,SNAPMGR-11-VM31</
      <HOSTS_REMCOPYBACKUP>
        <REQUEST_REMCOPYBACKUP>1</REQUEST_REMCOPYBACKUP></
    <BACKUP_CLIENT_SETTING>
  </BACKUP>

```

## Using Database Filter to display databases to back up in a DAG

You can use Database Filter to specify the criteria to display the Exchange Server databases to be backed up in a DAG. Databases that satisfy the filter criteria are displayed, enabling you to select them for backup and verification. The Database Filter also affects how the backup job is scheduled.

### About this task

A DAG is a group of up to 16 servers, each with many databases. Database Filter enables you to narrow the number of databases that are displayed. You can access Database Filter when you connect to the DAG from the following locations within the SnapManager user interface:

- Backup wizard
- Actions pane for backup

The SnapManager for Exchange Backup wizard does not prepopulate the names of the DAG database copies and names of the servers when scheduling a backup. However, the wizard prepopulates the server names along with the DAG database names when a backup filter is applied.

When a database failover occurs, a passive copy becomes active and the schedule fails. You can change the SnapManager job backup command to remove the server names from the backup command, as in the following example:

```

new-backup -Server 'DAGEXCH' -ClusterAware -ManagementGroup 'Standard' -
  RetainDays 15 -RetainUtmDays 15 -Database
  'MDB01','MDB25','MDB02','MDB26','MDB27'
  -UseMountPoint -ActiveDatabaseOnly -RemoteAdditionalCopyBackup $True -
  RetainRemoteAdditionalCopyBackupDays 15

```

When you specify criteria that is different from the default criteria “any,” the word *Filtered* is displayed in italics in the **Results** pane of the Backup window.

### Steps

1. Depending on the databases that you want to display, select one of the following:

If you want to...	Then do this...
Display databases on a member server of the DAG	Select the server from the <b>Show databases located on an Exchange member server</b> list.

If you want to...	Then do this...
Display all databases in the DAG	Select Any from the <b>Show databases owned by a particular server</b> list.
2. To display the databases based on database activation attributes, select one of the following from the <b>Show databases based on the copy specified</b> box:	
If you want to...	Then do this...
Narrow the display to active primary databases on the member server selected in Step 1 or in the DAG	Select the <b>Active Primary Database</b> option from the <b>Copy</b> list.
Narrow the display to passive copy databases on the member server selected in Step 1 or in the DAG	Select the <b>Passive Copy Database</b> option from the <b>Copy</b> list.
Display both active and passive databases on the member server selected in Step 1 or in the DAG	Select the <b>Any</b> option from the <b>Copy</b> list.
3. To narrow the display of databases based on their Exchange activation order, select one of the following from <b>Show databases by their Exchange activation order</b> box:	
If you want to...	Then do this...
Display the databases irrespective of their Activation Preference numbers	Select Any from the <b>Activation Preference</b> list.
Display the databases based on their Activation Preference number	Select the Activation Preference number from the <b>Activation Preference</b> list.

## What the Activation Preference Number of a mailbox database is

The Activation Preference Number of a mailbox database in Exchange Server is a number specifying the preference value of that database copy. You can select a particular database to back up by specifying its Activation Preference Number through the Database Filter.

The original database in a DAG always has an Activation Preference Number of 1. During database activation, when multiple database copies satisfy the criteria for activating, the Activation Preference Number is used to decide which database copy is to be activated.

## Managing backup retention

Your backup retention strategy needs to balance storage efficiency with restore needs. You can specify that SnapManager automatically delete older backups or transaction logs, or you can delete these items explicitly.

**Note:** You should not use SnapDrive or storage system administration tools to delete Snapshot copies created by SnapManager. Doing so leaves behind unwanted data that cannot be removed.

## Maximum number of Snapshot copies per volume

The Data ONTAP software used with SnapManager supports a maximum of 255 Snapshot copies per volume, including copies not created by SnapManager. Because each SnapManager backup operation

creates Snapshot copies, a SnapManager backup operation fails if the volume that contains the database LUN exceeds the 255 Snapshot copy capacity.

**Note:** The total number of Snapshot copies on a volume might exceed the number of retained backups. For example, if a single volume contains both the SnapInfo directory and the databases, each backup operation generates two Snapshot copies on the volume.

## Criteria for deleting backups

To avoid reaching the limit of 255 Snapshot copies per volume, you should delete the backups that you no longer need. You can delete backups by using either the age-based criteria or the quantity-based criteria.

### Age-based Snapshot copy deletion

You can delete backups that are assigned to a particular backup management group and that are older than a specified number of days. You can specify the deletion criteria by selecting the `Delete backups older than` option. Enter the number of days for which you want to keep the most recent backups for the backup management group.

### Quantity-based Snapshot copy deletion

You can delete backups that are assigned to a particular backup management group so that only a specified number of backups remain for that backup management group. Specify the deletion criteria by selecting the `Delete backups in excess of` option, which will delete the oldest backups. Enter the number of the most recent backups that you want to keep for the backup management group.

**Note:** Automatic deletion of older backups in a management group is the recommended method for managing the number of Snapshot copies stored on your system.

When SnapManager counts the number of stored backups, it also counts the backups that are shared by multiple databases. Therefore, more backups than you specify in the `Delete backups older than` or `Delete backups in excess of` box are retained.

For example, assume that you back up databases A and B. Both reside on the same volume and contain the following backup sets.

Backup set	Description
Database A	
exchsnap__orbit3_01-23-2012_16.21.07	Old backup—good
exchsnap__orbit3__recent	Recent backup—good
Database B	
exchsnap__orbit3_01-23-2012_16.21.07	Old backup—good
exchsnap__orbit3__recent	Recent backup—inconsistent

You set the “Delete backups in excess of” box to 2 to preserve the most recent two backup sets.

To preserve two good backups for Database B, SnapManager does not delete the Snapshot copy `exchsnap__orbit3_01-23-2012_16.21.07`, which is old and good. Therefore, because both databases reside on the same volume, three backups for Database A remain instead of two.

## Managing multiple backup management groups

You need to understand how SnapManager handles backup retention when multiple management groups are used in the backup process.


SnapManager handles backup retention and marks each Snapshot copy as either UTM or PIT based on the individual management group Snapshot copy and retention settings. When multiple management groups (Standard, Daily, and Weekly) are used on a single host or per DAG, this process becomes multi-layered.

For example, before starting the backup retention job, the management groups must have the following settings for their Snapshot copies:

- Standard: 8
- Daily: 4
- Weekly: 2

The Standard group has the following retention settings:

- Backups to retain: 8
- Backups to retain with UTM logs: 5



SnapManager handles this sample scenario as follows:

1. When the next Standard Snapshot copy (S9) is taken, SnapManager deletes the first copy (S1) because the Standard retention setting allows for only eight backups on the host.
2. SnapManager retains the most recent five Standard backups with logs (UTM) (S9, S8, S7, S6, S5) as stated by the retention policy, and then deletes the remaining three (S4, S3, S2) with logs and marks them as PIT.

3. When SnapManager marks any backup of a particular management group to PIT (S4, S3, and S2), SnapManager checks in the same operation for any older backups that exist in the other management groups (Daily and Weekly) and marks all of them as PIT (W1, D3, D2, and D1). Once the chain of transaction logs is broken, the Snapshot copy preceding that particular backup cannot be used for up-to-minute restore operations.

**Note:** Any FRPs that exist between two consecutive PIT backups will be removed as there is no need of additional transaction logs for PIT backup to be restored.

In the case of remote backup retention, the same logic applies to local backups, and the management groups for local backup are used for all remote backups.

UTM retention will remain the same for both local and remote backup retention. In the example, you should replace the backup retention value with remote retention value and keeping the other values the same will get the know how many remote retention will be UTM and PIT.

## Automatically deleting backups

You can manage the number of Snapshot copies you store by configuring SnapManager to delete backup copies automatically, based on how old the backup copies are or how many of them are stored.

Automatic deletion deletes a backup copy only if the backup copy has the following characteristics:

- The backup copy is in the same management group as the management group of the backup copies that you just created.
- The backup copy is the oldest backup copy of the database.
- The number of backup copies exceeds the backup retention level that you specified in the `Delete backups older than` option or the `Delete backups in excess of` option.

**Note:** If a database is no longer available in the Exchange server, backup copies associated with it cannot be removed by the delete backup module.

If you do not select automatic backup deletion, backup copies that are created after the current backup job are retained. This requires manual removal of backup copies or enough storage capacity for all backup copies and transaction logs. You can delete the retained backup copies by selecting automatic backup deletion in the next backup job.

## Automatic deletion of orphaned SnapInfo folders

SnapManager for Exchange checks for orphaned SnapInfo (SIF) folders left behind from an automatic deletion operation and deletes these folders during the next backup operation, which saves volume space.

When either database or Log LUN backup copies are automatically deleted, SIF folders are left behind. These SIF folders contain physical exchange transaction log files that consume a large amount of volume space and should be deleted. During the backup operation, SnapManager runs in the background and performs the following process to delete these orphaned folders:

1. SnapManager reviews each SIF folder from each database.
2. SnapManager determines whether the database or LUN Snapshot copy is missing for each SIF folder.

**Note:** Frequent Recovery Point (FRP) folders are not considered during this operation since there are no associated Snapshot copies.

3. If the database or LUN Snapshot copy is missing, then SnapManager designates the SIF folder as an orphaned folder.

4. SnapManager deletes the orphaned SIF folder, including the Log folder, during the next backup operation.

## Explicitly deleting backups

You can explicitly delete any backups or Snapshot copies of LUNs created during a restore operation by selecting them. You can delete either an individual backup or multiple backups.

When you explicitly delete backups, you can also ensure that the deletion is limited so as not to create a break in the continuity of the transaction logs between the previous backup and the present time. Doing so allows you to retain up-to-the-minute restore ability for older backups from other management groups.

**Attention:** You should not use SnapDrive or the storage system administration tools to delete Snapshot copies created by SnapManager. Doing so leaves unwanted data that you cannot remove.

You can delete individual selected Snapshot copies for either full database backups or transaction logs. You can also select a database to be deleted, the types of backup set components to be deleted (full database backups or SnapInfo directory backups), and the type of backup management group to be deleted for multiple backups.

When you explicitly delete backup sets (or related SnapInfo Snapshot copies) that contain selected databases, you can expand the scope of the deletion to include backups that contain only some of the selected databases (in addition to backups that contain all the selected databases).

### Option to retain up-to-the-minute restore ability

If you delete backups and transaction logs that are not the oldest backups in your backup list, the oldest backups are no longer available for up-to-the-minute restore. Ensure that you do not break the continuity of transaction logs between the previous backup and the present time.

When you delete backups of a particular backup management group, limiting the backup deletion allows you retain the ability to use the older backups in an up-to-the-minute restore operation.

To limit the backup deletion, use the `Retain up-to-the-minute restore ability for older backups in other backup management groups` option in the Advanced Options dialog box. Selecting this option consumes more space on your storage system for the transaction logs that are not deleted, and it is selected by default.

If all backups have the same management group designation, there is no effect if you clear the option. SnapManager deletes the SnapInfo directory transaction logs according to the Delete backups setting, and ignores the Retain up-to-the-minute restore ability option.

By default, SnapManager backs up selected database files and all associated transaction logs, so that up-to-the-minute restores are possible from all backups. If you do not need to perform an up-to-the-minute restore operation from the older backups, delete the transaction logs to free space on the storage system that contains the backups. The `Retain up-to-the-minute restore ability for older backups in other backup management groups` option is automatically selected the next time you start SnapManager.

As an example, assume that you have the following backups stored:

- exchnsnap\_WINSRV3\_01\_05\_2004\_19.05.37\_Daily (Oldest)
- exchnsnap\_WINSRV3\_01\_06\_2004\_08.23.14
- exchnsnap\_WINSRV3\_01\_06\_2004\_10.22.19
- exchnsnap\_WINSRV3\_01\_06\_2004\_12.18.58
- exchnsnap\_WINSRV3\_01\_06\_2004\_14.28.03 (Newest)

If you perform a Standard backup, and specify that you want to retain only the four most recent Standard backups the oldest Standard backup is deleted, and the daily backup becomes point-in-time only:

- exchsnap\_WINSRV3\_01\_05\_2004\_19.05.37\_Daily (Oldest) (Daily backup becomes point-in-time only)
- exchsnap\_WINSRV3\_01\_06\_2004\_08.23.14 (Oldest Standard backup deleted)
- exchsnap\_WINSRV3\_01\_06\_2004\_10.22.19
- exchsnap\_WINSRV3\_01\_06\_2004\_12.18.58
- exchsnap\_WINSRV3\_01\_06\_2004\_14.28.03
- exchsnap\_WINSRV3\_01\_06\_2004\_16.23.48 (Newest) (New Standard backup created)

To avoid breaking the continuity of transaction logs, you can enable the `Retain up-to-the-minute restore ability for older backups in other backup management groups` option.

## Explicitly deleting individual backup copies

You can delete an individual backup copy.

### About this task

If you select the `Retain up-to-the-minute restore ability for older backups in other backup management groups` option, transaction logs are deleted only from this backup management group; transaction logs are not deleted from other backup management groups. If you clear this option, transaction logs are deleted from other backup management groups as well.

### Steps

1. In the Scope pane, select **Restore**.
2. Select the database in which the backup copy resides.
3. Click **Delete** in the **Actions** pane.
4. In the **Delete Single Backup** window, select the backup copies that you want to delete.
5. Review the list of backup copies that share the Snapshot copy you want to delete, because all of these backup copies are deleted simultaneously.
6. If you want to retain up-to-the-minute restore ability for older backups in other backup management groups, select `Retain up-to-the-minute restore ability for older backups in other backup management groups`.
7. Click **OK**.

## Explicitly deleting backup sets or SnapInfo Snapshot copies

You can delete backup sets (or only the SnapInfo Snapshot copies related to those backup sets) that contain one or more selected databases. If you need to free some space or reduce your Snapshot copy

count without deleting a SnapManager Backup, you can delete the SnapInfo Snapshot copies without affecting the associated backup.

#### About this task

The Backup Component selection is set to Backup Data Sets by default, so that you delete the entire backup set, including related transaction logs and SnapInfo directories. You can narrow this selection further by selecting **SnapInfo Snapshot Copies Only**.

The Management Group selection further narrows the scope of the deletion by specifying the backup management group of the backups that you want to delete.

If you select the `Retain up-to-the-minute restore ability for older backups in other backup management groups` option, transaction logs are deleted only from the specified backup management group; transaction logs are not deleted from other backup management groups. On the other hand, if you clear this option, transaction logs are deleted from other backup management groups as well.

#### Steps

1. In the **Actions** pane, click **Delete Backup**.

**Note:** You can access **Delete Backup** when you connect to a member server, but not when you connect through the DAG.

2. If you want to display a list of the backup components on your SnapManager system, select **Backup Component**.

**Note:** You cannot select an individual backup component for deletion by using this list.

3. Select one or more databases whose backups are to be deleted.

4. To delete backups that contain some of the selected databases (in addition to the backups that contain all the selected databases), click **Advanced > Delete backups that contain databases belonging to one or more storage groups**.

**Note:** This option is useful only if you created backups that did not contain all databases in a set.

For this backup deletion operation only, multiple backup deletions delete backups that contain any one or more of the selected databases.

5. Select the Backup Component.

6. Select the Management Group.

**Note:** By default, this is set to **Standard**; this can be changed to either **Daily** or **Weekly**.

7. If you want to retain up-to-the-minute restore ability for older backups in other backup management groups, click **Advanced > Retain up-to-the-minute restore ability for older backups in other backup management groups**.

**Note:** If you selected All for the backup management group, this option has no effect and is disabled.

8. In the specified database and backup management groups, specify those that you want to delete.

If you want to delete...	Then...
The oldest backups	In the <b>Delete oldest backups in excess of</b> box, specify how many of the newest backups in the specified backup management group you want to preserve.

If you want to delete...	Then...
All the backups in the specified backup management group	Select the <b>Delete all backups in the specified management group</b> option.
Only the backups that are older than a specified number of days	In the <b>Delete backups older than</b> box, specify the number of days for which you want to keep the most recent backups.

9. Perform *one* of the following actions:

- Review the backup components that would be deleted with the parameters entered without deleting them.
  - a. Click **Delete Preview**.  
The Delete Backup Data Set Preview window is displayed. A list of the backup components identified for deletion are displayed.
  - b. After previewing the deletion, if you want to delete the backup components listed in the preview, click **Delete**.  
The backup components listed in the preview window are deleted.
- Delete the backup components without previewing the deletion.
  - Click **Delete**.  
When the deletion is complete, a status pop-up window is displayed.

### Explicitly deleting Snapshot copies created by SnapManager Restore

You can delete Snapshot copies created during previous restore operations.

#### Steps

1. In the **Actions** pane, click **Delete Backup**.
2. Select **Delete snapshot of LUNs created during restore**.
3. If you want to display a list of the backup components on your SnapManager system, select **Backup Component**.
 

**Note:** You cannot select an individual backup component for deletion by using this list.
4. Specify the restore Snapshot copies that you want to delete.

If you want to delete...	Then...
The oldest restore Snapshot copies	In the <b>Delete oldest snapshots in excess of</b> box, specify how many of the newest restore Snapshot copies you want to preserve.
All the restore Snapshot copies	Select the <b>Delete all snapshots created during restore</b> option.
Only the restore Snapshot copies that are older than a specified number of days	In the <b>Delete backups older than</b> box, specify the number of days for which you want to keep the most recent restore Snapshot copies.

5. Perform *one* of the following actions:

- Review the restore Snapshot copies that would be deleted, without deleting those copies.
  - a. Click **Delete Preview**.  
The Delete Backup Data Set Preview window is displayed. A list of the restore Snapshot copies identified for deletion are displayed.

- b. After previewing the deletion, if you want to delete the restore Snapshot copies listed in the preview, click **Delete**.

The restore Snapshot copies listed in the preview window are deleted.

- Delete the restore Snapshot copies without previewing the deletion.

- Click **Delete**.

When the deletion is complete, a status pop-up window is displayed.

### Problem deleting backups due to busy Snapshot copy error

If you delete a backup copy of a LUN that was already backed up by another Snapshot copy, you get an error message saying that the Snapshot copy is busy and cannot be deleted. In this case, you must delete the most recent backup copy before the older backup can be deleted.

You can view Snapshot copy status in OnCommand System Manager or use the storage system `snap list` command. For more information about deleting busy Snapshot copies, see the *System Administration Guide* for your version of Data ONTAP.

**Note:** To avoid this situation, ensure that you do not make backup copies of LUNs that are already backed up by Snapshot copies (for example, during a verification or while archiving from a LUN backed by a Snapshot copy).

## Managing database verification

SnapManager uses `ChkSgFiles`, a Microsoft Exchange utility, to verify the Exchange Server database files. It uses `ChkSgFiles.dll`, an integrity verification library, to verify the Exchange Server databases and transaction logs. `ChkSgFiles` also verifies the page-level integrity of the databases.

The `ChkSgFiles.dll` library checks the databases and transaction log files for physical and logical corruption by using checksum verification. You must install Exchange Server Management Tools on the server that performs integrity verification.

**Note:** By default, SnapManager Restore requires that you restore only from verified backups, but you can override this requirement.

If there are at least two database copies in a DAG, database integrity checking is not required. Microsoft recommends verification of the transaction logs. By default, when performing a DAG backup with SnapManager for Exchange, verification is off.

You can monitor both locally running and remote verification jobs through the main SnapManager for Exchange Admin Center.

Note the following with regard to verification servers and verification jobs:

- A single Exchange mailbox server can run only one verification process at a time on a particular verification server.
- A verification server can simultaneously run one verification job from each Exchange mailbox server.
- More than one verification server can be used to simultaneously verify more than one backup job on a single Exchange mailbox server.
- A verification server can utilize a virtual machine to off-load the verification.

Optionally, you can use the `Eseutil` command-line tool to check the database and the transaction logs for physical-level corruption. `Eseutil` is installed automatically as part of Exchange Server. For more information about physical-level corruption, see your Microsoft Exchange documentation.

## LUN requirements for verifying databases in a backup set

SnapManager requires that all databases be mounted before verification. The SnapManager verification server must have enough drive letters or mount points to mount the LUNs that store the backup sets that you verify, depending on your verification scenario.

SnapManager mounts the LUNs that contain the backup sets that you select for verification with SnapDrive commands. Each mounted LUN requires one available drive letter or mount point.

To verify backups that are stored on a single LUN or across multiple LUNs, the verification server must have at least one drive letter or mount point available.

If you are using a remote verification server, you need two unassigned drive letters or mount points to verify multiple backup sets in a single job.

In other situations, you need one unassigned drive letter or mount point.

The one unassigned drive letter or mount point is to mount the LUN in a Snapshot copy. The Snapshot copy stores the transaction log directory for multiple databases for transaction log verification. When SnapManager verifies the transaction logs for the first backup set, it does not unmount the LUN or release the assigned drive letter. SnapManager reuses the same LUN for the transaction log verification for the second backup set.

If a second drive letter or mount point is not available, schedule the backup or verification jobs to verify one backup set at a time.

**Note:** If you verify the destination volumes that contain more than one database sharing the same LUN for the transaction logs, ensure that a second drive letter or mount point is available for the Exchange server.

## Database verification workload management

Because running database verification on a production server can place a significant load on both the server and the storage systems, you must manage this workload to avoid degraded Exchange response, particularly during peak work hours.

You can use any of the following methods to manage your database verification workload:

- Deferred database verification
- Remote database verification
- Remote verification on a SnapVault destination
- Verification throttling
- Integrity verification on destination SnapMirror volume

## Deferred database verification

You can distribute your system workload by disabling automatic integrity verification, which takes place immediate after your backup is created, and then performing a separate verification later.

## Remote database verification

You can distribute your system workload by disabling automatic integrity verification, which takes place immediate after your backup is created, and then run verification from another Exchange server.

## Remote verification on a SnapVault destination

You can distribute your system workload by offloading the verification process to the SnapVault destination.

After installing or upgrading to this release, you must specify SnapMirror and SnapVault destination volumes in the Database Verification Settings dialog before you can schedule backup set verification on those volumes. In the Database Verification Settings dialog, click **Verification Server > Verification on Destination Volumes**, then choose the destination volumes.

## Verification throttling

You can slow down the integrity verification to decrease the load on the Windows host and the storage systems by using verification throttling when using Microsoft Exchange Server.

**Note:** Database verification is not a support requirement for databases with at least two copies in a Database Availability Group. By default, when performing a DAG backup with SnapManager for Exchange, verification will be off.

SnapManager can throttle integrity verification (the Microsoft Exchange consistency checker utility) in checksum verification mode to reduce the load on the host CPU, and on the storage subsystem. This .dll file checks the databases and transaction log files for integrity verification.

## Integrity verification on the destination SnapMirror volume

SnapManager supports integrity verification on the destination SnapMirror volume.

Integrity verification on the destination SnapMirror volume is available through the following:

- Backup and Verification window
- Backup wizard

**Note:** For SnapMirror databases, in Backup view, the **Verify on SnapMirror Destination** will be enabled only at the Node level. For this to be enabled at DAG level, the node hosting the SnapMirror database should be the DAG owner.

- Deferred verification
- **Test Restore** button and Restore wizard

**Note:** When you perform a restore operation from the Restore window, SnapManager prompts for integrity verification on the source or destination volume, if there are any SnapMirror volume relationships associated with it.

The Choose SnapMirror Destination Volumes for Integrity Verification window displays each SnapMirror volume as a tree showing the relationship among the volume, the LUN, and the databases contained in it. For each source volume, there is a list of destination volumes, and each destination volume displays a SnapMirror state. You can select a SnapMirror destination volume for each SnapMirror source volume for which you want verification.

After installing or upgrading to this release, you must specify SnapMirror and SnapVault destination volumes in the Database Verification Settings dialog before you can schedule backup set verification on those volumes. In the Database Verification Settings dialog, click **Verification Server > Verification on Destination Volumes**, then choose the destination volumes.

## Backup verification status reporting

You can use SnapManager Restore to determine the verification status of each backup and of the transaction logs for each backup.

SnapManager Restore shows you a list of the backups that have been created and indicates the verification status of each backup. For each listed backup, the date and time of the backup operation is displayed, as well as an icon that indicates the verification status of the backup.

You can also determine whether the transaction logs for a specific backup have been verified by selecting the backup in the Restore window. The transaction log verification status is included with the backup information shown in the Result pane.

## Backup verification status icons

The following verification status icons represent the verification status of the databases.

Verification status icon	Database verification status
	Verified databases.
	Unverified databases.
	Databases failed verification and cannot be restored.
	Verified databases, but unverified transaction logs. If you need to restore from backups with this symbol, contact technical support.

## Where to run database and transaction log integrity verification

Regardless of when you verify the databases in a backup set, you can perform the verification either on the production server or on a remote verification system.

In the simplest SnapManager configuration, verification is run from your production server; however, this type of verification is CPU and disk I/O intensive, so you might want to avoid performing verification on the production server during peak usage as it can affect Exchange performance.

Performing integrity verification on a remote system minimizes the negative performance impact on Exchange system resources and the backup schedule.

## When to verify the databases in a backup set

You can verify the databases in your SnapManager backup sets immediately after creation, at a later scheduled time, or when you restore them.

SnapManager can automatically verify full backup sets at the time the backup is created if the **Verify backed up databases and transaction logs** option is selected. (By default, this option is OFF at the DAG; when doing backups at the node level this option is ON.) This is simple and ensures that each database in the backup set is verified. However, this method significantly increases the time required to complete the backup.

You can start an operation to verify the databases contained in one or more backup sets that have already been created. You can start the verification immediately, or you can schedule the verification to occur later when it does not affect performance or delay later backups.

If you attempt to restore from a backup set on which a database consistency check has not been run successfully, SnapManager prompts (but does not require) you to first verify the databases in that backup set.

**Note:** When you perform verification on a LUN clone and you make a Snapshot copy of the volume while a LUN clone exists, a “busy Snapshot copy” is created, which might cause problems when you attempt to delete Snapshot copies. To avoid this, do not schedule backups while a verification is in progress.

## Managing transaction log backups

SnapManager provides two ways for you to start or schedule a transaction log backup by using the SnapManager Backup wizard or the Backup and Verify option.

### Backup options

There are several options when backing up transaction logs.

The following are the backup choices for the `Retain up-to-minute restore ability for older backups` option:

- Backups generated in the last (days)
- The most recent (number of backups)
- All the older backups (do not delete any logs)

### Where to specify the backup options

You can specify the backup options in one of the following locations:

- **Backup wizard, Retain up-to-the-minute restore ability for older backups**
- **Backup and Verify, Up-to-minute Restore Options**
- **Delete Backup, Advanced Options**

After you set the options, other backups outside the specified backups in the same management group become point-in-time only.

PowerShell continues to support the `NoUtmRestore` switch; however, the switch is disabled as soon as you specify values for `RetainUtmBackups` or `RetainUtmDays`.

### Triggering UTM retention

To reduce disk space issues with backups, up-to-the-minute (UTM) retention is triggered independently of overall backup retention.

Prior to SnapManager 6.1 for Microsoft Exchange, UTM retention was dependent on Backup Retention, and did not run if there were no backups or Snapshot copies to delete per the Backup Retention settings. The UTM Restore Retention did not enforce the retention policy until the Backup Retention policy was triggered.

For example, if there is a retention of 30 days for SnapManager for Exchange backups and a period of 4 days of UTM Restore Retention, UTM restore policy does not delete backed up transaction logs and mark the older backup sets “Point in Time” until the Backup Retention policy starts deleting the excess backup sets because its specified threshold is reached (when more than 30 days of backups are created). This situation has the potential to create disk space issues.

To address this concern, UTM retention is now separate from the backup retention, and can be triggered independently after the backup. Those two types of retention coexist within SnapManager for Exchange.

### Valid input for UTM-related switches

The following switches can have the following values:

#### `RetainUtmBackups`

Values: an integer between 0 and 255

#### `RetainUtmDays`

Values: a float value  $\geq 0$

A negative number generates an error message.

If you specify values for both `RetainUtmBackups` and `RetainUtmDays`, an error message is generated. However, there is one exception: in PowerShell, if you specify -1 for either `RetainUtmBackups` or `RetainUtmDays`, then SnapManager ignores the switch and does not generate an error. Instead, it uses the prior backup retention method and takes the default value (0) for `NoUtmRestore`.

If none of the three UTM-related switches is specified, the prior backup retention method is used with the `NoUtmRestore` default value (0).

### Registry keys

The following registry keys are added under `HKEY_LOCAL_MACHINE->Software->Network Appliance->SnapManager for Exchange->Client`:

- `KeepLastNSnapshotUtm` REG\_DWORD
- `KeepLastNSnapshotOlderDayUtm` REG\_SZ
- `KeepLastNSnapshotOptUtm` REG\_DWORD

When both `KeepLastNSnapshotUtm` and `KeepLastNSnapshotOlderDayUtm` have a valid value, SnapManager uses the following values of `KeepLastNSnapshotOptUtm` to determine if the operation is either number-based or day-based:

- 0: Number-based
- 1: Day-based
- 2: Retain up-to-the-minute restore ability for all the older backups

There are three new entries added to the control file under `<BACKUP_CLIENT_SETTING>`:

- `<BACKUP_SET_TO_KEEP_UTM>`
- `<BACKUP_SET_TO_KEEP_IN_DAYS_UTM>`
- `<DELETE_BACKUPS_OPTION_UTM>`

When both `BACKUP_SET_TO_KEEP_UTM` and `BACKUP_SET_TO_KEEP_IN_DAYS_UTM` have a valid value, SnapManager uses the following values of `DELETE_BACKUPS_OPTION_UTM` to determine if the operation is either number-based or day-based:

- 0: Number-based
- 1: Day-based
- 2: Retain up-to-the-minute restore ability for all the older backups

### UTM retention scenario 1: backups in a stand-alone system

This scenario assumes four backups exist in the stand-alone system with backup retention set to 8 and UTM to 3.

The result of this scenario is after the backup is complete, there will be 5 backups with the first 3 of backup with UTM restorability and the last 2 with PIT restorability.

- Before:

N1 <- old	=> UTM
N2 <-	=> UTM
N3 <-	=> UTM
N4 <-	=> UTM

- After:

N1 <- old	=> PIT
N2 <-	=> PIT
N3 <-	=> UTM
N4 <-	=> UTM
N5 <- new	=> UTM

### UTM retention scenario 2: days of backup in a stand-alone system

This scenario assumes 4 days of backup exist in the stand-alone system with backup retention as 8 days and UTM as 3 days.

The result of this scenario is after the backup is complete, there will be 5 backups with the first 3 days of the backup with UTM restorability and the last 2 days with PIT restorability. Although backup retention is set to 8, UTM is still enforced and deletes the extra transaction logs from specified management group.

- Before:

D1 <- old	=> UTM
D2 <-	=> UTM
D3 <-	=> UTM
D4 <-	=> UTM

- After:

D1 <- old	=> PIT
D2 <-	=> PIT
D3 <-	=> UTM
D4 <-	=> UTM
D5 <- new	=> UTM

### UTM retention scenario 3: backups on a DAG Active node system

This scenario assumes 4 backups exist on a DAG Active node system with backup retention set to 8 and remote backup retention is also set to 8 and UTM to 3.

Although backup retention is set to 8, UTM still gets enforced and deletes the extra transaction logs from specified management group.

The result of this scenario after backup is complete is as follows:

- Active node:

There will be 5 backups with the first 3 of the backups with UTM restorability and the last 2 with PIT restorability.

N1 <- old	=> PIT
N2 <-	=> PIT
N3 <-	=> UTM
N4 <-	=> UTM
N5 <- new	=> UTM

- Passive node:

There will be 5 backups with the first 3 of the backups with UTM restorability and the last 2 with PIT restorability.

N1 <- old	=> PIT
N2 <-	=> PIT
N3 <-	=> UTM
N4 <-	=> UTM
N5 <- new	=> UTM

### UTM retention scenario 4: days of backup on a DAG Active node system

This scenario assumes 4 days of backups exist on a DAG Active node system with backup retention set to 8 days and remote backup retention is also set to 8 days and UTM to 3 days.

Although backup retention is set to 8, UTM is still enforced and deletes the extra transaction logs from the specified management group.

The result of this scenario after backup is complete is as follows:

- Active node:

There will be 5 days of backups with the first 3 days of backup with UTM restorability and the last 2 days with PIT restorability.

D1 <- old	=> PIT
D2 <-	=> PIT
D3 <-	=> UTM
D4 <-	=> UTM
D5 <- new	=> UTM

- Passive node:

There will be 5 days of backups with the first 3 days of backup with UTM restorability and the last 2 days with PIT restorability.

D1 <- old	=> PIT
D2 <-	=> PIT
D3 <-	=> UTM
D4 <-	=> UTM
D5 <- new	=> UTM

## Option to back up transaction logs that Exchange will truncate

If you do not need to retain up-to-the-minute restore ability from a backup that is not the most recent, you can reduce the amount of disk space required by omitting the backup of transaction logs that Exchange truncates after the SnapManager Backup operation finishes.

By default, SnapManager backs up database and the transaction logs. SnapManager creates a database Snapshot copy and a SnapInfo Snapshot copy, and truncates transaction logs by removing any entries already committed to the database.

Manage this feature using the **Back up transaction logs that will be truncated by Exchange at the end of the backup** option in the Advanced Options dialog box.

The option to back up all transaction logs (including those that Exchange truncates after the backup operation finishes) is selected by default. When you clear the option to back up all transaction logs, the database and transaction log verification cannot be deferred, and the option to back up all transaction logs is automatically enabled the next time you start SnapManager.

## Frequent Recovery Point backup and restore operations

A Frequent Recovery Point (FRP) backup operation backs up new transaction logs and is performed after a full backup operation. You can restore data up to a recovery point that you select.

SnapManager combines the restore operation of a full backup copy and the required transaction logs to restore to the selected recovery point.

### How the Frequent Recovery Point feature works

FRP backs up transaction logs at a frequency that you determine to meet the Recovery Point Objective (RPO), and backs up the most recent transaction logs created after the most recent full backup copy or the previous FRP backup copy.

**Note:** If FRP backup is started when the full backup is in progress, the FRP backup waits for the full backup to finish. The same is true in reverse; if full backup is started when FRP backup is in progress, the full backup waits until the FRP backup is complete.

The Frequent Recovery Point feature creates backup copies at a specific frequency and names them using a unique naming convention. The lowest value for the interval between any two Frequent Recovery Point backup copies is 10 minutes. The default value is 15 minutes.

You can create a Frequent Recovery Point backup copy after you create a full backup copy. You can trigger a SnapMirror update of the SnapInfo volume after the Frequent Recovery Point backup is complete.

A Frequent Recovery Point backup operation creates a Snapshot copy on the SnapInfo volume and names it using the following convention:

`efrpinfo_<exchserver name>_date_time`

SnapManager retains only one Snapshot copy with this name.

You can use “Run command after operation” with the Frequent Recovery Point feature.

## Frequent Recovery Point backup operations

There are various operations that you can perform with Frequent Recovery Point backups.

- Create a Frequent Recovery Point backup
- Schedule a new Frequent Recovery Point backup job with the name FRPBackup
- Modify an existing scheduled Frequent Recovery Point backup job
- Delete an existing Frequent Recovery Point backup job

## Frequent Recovery Point backup operation in a DAG

You can perform a Frequent Recovery Point (FRP) backup operation in a Database Availability Group (DAG) by connecting either to the DAG or to a member server.

When you connect to the DAG, the Frequent Recovery Point backup operation creates FRP backups on all member servers in the DAG. You can create FRP backup for either active or passive databases, or for all databases on a specified member server, or on all member servers in the DAG.

Each server has its own FRP backup, independent of other member servers in the DAG.

## Verification of Frequent Recovery Point backup copies

SnapManager does not perform transaction log verification when it creates Frequent Recovery Point backup copies. When you verify a full backup copy, SnapManager verifies all its previous Frequent Recovery Point backup copies.

## Deletion of Frequent Recovery Point backup copies

When you delete a full backup without enabling up-to-the-minute restore functionality, SnapManager deletes all the subsequent Frequent Recovery Point (FRP) backup copies up to the next full backup. You cannot delete FRP backup copies independent of full backups.

## Frequent Recovery Point backup reports

SnapManager creates a separate folder in which to save Frequent Recovery Point backup reports and names it using the naming convention "FRP Backup [server]". The naming convention for the file names of reports does not change.

## Performing a Frequent Recovery Point backup operation

You can either create a Frequent Recovery Point backup or schedule Frequent Recovery Point backups.

### Steps

1. In the Scope pane, select an Exchange server.
2. In the **Actions** pane, click **Frequent Recovery Point Backup**.

If...	Then...
<b>You have not scheduled the backup job of Frequent Recovery Point</b>	SnapManager does not display any job details. You can use the <b>Create Job</b> button to create a new job specification.
<b>You have scheduled the backup job of Frequent Recovery Point</b>	SnapManager displays the job details. Use <b>Update Job</b> to update the job specifications.

**Note:** You can disable the current job by selecting the **Disable scheduled job** check box.

3. Under **Selected for backup**, click **Select All** to select all the databases at one time, or click **Unselect All** to clear all your selections.
4. Under **Backup Frequency**, specify the maximum frequency at which the recovery point backup needs to be made:
  - a. In the **Every** list, select a number of minutes or hours from the list box to specify the interval between the Frequent Recovery Point backups.
  - b. In the **Start at** list, select the time to run the Frequent Recovery Point backup.
5. Under **Operation options**, select **Run command after operation** if you want to run a command or script of your choice after the backup operation completes.
6. Do one of the following:
 

<b>If you want to...</b>	<b>Then do this...</b>
<b>Create one Frequent Recovery Point backup</b>	Click <b>Create Recovery Point</b> . The Backup Status window is displayed with the list of tasks.
<b>Run Frequent Recovery Point backups at a specific interval</b>	Click <b>Create Job</b> . A scheduled job named FRPBackup is created under Windows Scheduled Tasks.

7. When you click **Create Recovery Point** and the backup status is displayed, click **Start Now** to create one new recovery point backup.
8. In the **Backup Status** window, click **Close** when the backup operation completes.  
SnapManager displays a message with the result of the Frequent Recovery Point backup.

## Changing the backup management group of an existing backup set

You can use the Change Backup Management Group dialog box to change the backup management group to which the selected backup set belongs.

### About this task

You cannot change the backup management group of the most recent backup sets that were created using the generic naming convention.

### Steps

1. In the SnapManager console root, click **Restore**.
2. In the **Restore** panel, locate the backup set whose management group you want to change:
  - Database Snapshots (Standard group)
    - exchsnap\_exchservername\_\_date\_time
    - exchsnap\_exchservername\_\_ recent
  - Database Snapshots (Daily or Weekly group)
    - exchsnap\_exchservername\_\_date\_time\_\_backupmgmtgroup
    - exchsnap\_exchservername\_\_backupmgmtgroup\_\_recent

3. Right-click the name of the backup set to open a context menu and then select **Change Management Group**.
4. Review the backups listed in the **Backups sharing this Snapshot list**.

The backup management group for all these backups is changed if you complete this operation because they share a common backup set.
5. In the **New Management Group** list, select the backup management group you want to change to.

When you change a backup's backup management group, you also change that backup's name because the name includes the backup management group.
6. Click **OK**.

The backup management group for this backup and all backups listed in the All Backups Sharing This Snapshots list is changed.  
The report for the backup management group change is in the `Miscellaneous` report directory.

## Reasons that a SnapManager backup might fail

If your SnapManager backup fails, check the backup report to determine the cause of the failure.

### Problem: cluster failover during backup

If a cluster failover or a Windows cluster move group happens during a backup operation, the backup fails, and you need to restart the backup operation.

### Problem: Snapshot copy limit reached

You get an error message if you try to back up a LUN that contains more than 255 Snapshot copies. The backup operation fails irrespective of whether SnapManager created the Snapshot copies or not.

**Note:** Automatic backup deletion is performed only after a successful backup process is complete. Therefore, you must be able to create new Snapshot copies before you begin a new backup procedure.

### Problem: SnapInfo directory being accessed

You get an error message if you access the SnapInfo directory while performing a backup operation.

A SnapManager backup operation might include renaming a SnapInfo subdirectory, and Windows does not allow you to change a directory name while it is being accessed. Accessing the SnapInfo directory with Windows Explorer might cause the backup to fail. Ensure that you do not hold exclusive access to the SnapInfo directory on the Exchange host when a backup process is performed.

## Problem: SnapInfo directory out of space

You get an error message if your SnapInfo directory runs out of space. Expand the LUN that contains the SnapInfo directory and ensure that enough space remains in the volume for Snapshot copy creation.

## Problem: data does not match

You get an error message if you make changes to your Exchange Server databases after SnapManager started if you do not refresh your view. You can refresh your view by pressing F5, or you can restart SnapManager.

## Problem: busy Snapshot copy

You get an error message if you back up a LUN when a Snapshot copy of the LUN already exists and then you try to delete a Snapshot copy of the LUN. Event 249 is logged by SnapDrive and SnapManager backups fail.

## Problem: Snapshot copy already exists

You get an error message either if the system clock on the SnapManager host is not synchronized with the storage system clock or if SnapMirror replication is running when you try to start a backup operation.

Synchronize the system clock on the SnapManager host and the storage system clock so that SnapDrive functions properly. Also, ensure that any SnapMirror replications have enough time to complete before you initiate another SnapManager backup process.

## Problem: out of disk space

You get an error message when the database or the transaction logs use all of the available disk space in a volume. Resize the volume or expand the LUN.

## Problem: SnapManager server initialization failed

You get an error message either if the SnapManager server account or the server account permissions have changed; or if you exit SnapManager when the smesrvr.exe process is running.

Ensure that you use the correct server account and the correct server account permissions and that you terminate any orphaned SnapManager processes that run when you exit SnapManager.

## Troubleshooting integrity verification failure on SnapMirror destination volumes

If you have changed the volume SnapMirror relationship by mirroring data to a new destination storage system, SnapManager is not able to find the destination storage system and issues an error. You can set the new path as the volume SnapMirror destination for verification purposes.

### About this task

SnapManager for Exchange integrity verification on volume SnapMirror destination fails when the destination storage system is no longer available or it has changed. If the volume SnapMirror destination storage system is no longer available, SnapManager for Exchange provides the following error message:

**The destination volume Filer2:Vol1 of the source Logical Disk E does not have valid SnapMirrored state.**

SnapMirror State is Unknown. Please check the SnapMirror relationship Filer1:Vol1 Filer2:Vol1\_dest. Unable to mount snapshot, abort verification.

If you changed the volume SnapMirror relationship by mirroring data to a new destination storage system (for example, Filer3), SnapManager might still try to contact Filer2 which might not necessarily exist. This is very likely to happen when you perform a storage system hardware upgrade and replace an old storage system head with a new one. When you run the SnapMirror update command on new destination storage system it will work, but SnapManager for Exchange might not be up to date with the new destination storage system.

To resolve this problem, ensure that the following requirements are met:

- SnapMirror status on both the volume SnapMirror source and the volume SnapMirror destination storage system is updated with the new destination and that you have run at least one successful manual update (from storage system console) after the volume SnapMirror destination has changed.
- The new destination storage system holds the SnapManager for Exchange user in its built-in Administrators group.

Complete the following steps to set the new path as a volume SnapMirror destination for verification purposes.

**Steps**

1. Click SnapManager for Exchange Backup verification settings.
2. Click the **Verification server** tab.
3. Click the **Verification on destination volumes** button.

## Restoring databases

---

You can restore a DAG copy, an archived backup set, or databases created on different Exchange servers. You can also restore a database to a Recovery Database.

You can restore your databases either from an earlier backup copy or from an archive version. You can then optionally restore the databases up to their current state by replaying the transaction logs. Replying the transaction logs is necessary if your Exchange data becomes corrupted or unavailable.

You can restore the production backup copies to either of the following locations:

- The same database
- The Recovery Databases

Additionally, you can also restore database copies to member servers in the DAG.

## When to choose SnapManager restore to recover Exchange mailbox databases

You should know when to use SnapManager to recover an Exchange mailbox from database failure or data corruption.

The database switchover to a different server in the Database Availability Group (DAG) can be done using the EAC or the Exchange management shell.

Use SnapManager in the following situations:

- When using Exchange Server and DAG cannot recover the failed database.
- When using Exchange Server and DAG to recover the database will cause data loss.
- When using Exchange Server and DAG cannot recover the database to a particular point in time; for example, when a virus has spread to all copies of databases.
- When using Exchange Server and DAG is confusing to the administrator, and using SnapManager restore is simple, fast and reliable.

## What SnapManager restore does

SnapManager restore performs a sequence of steps when it restores a backup copy. If you are restoring in a DAG, do not attempt to manage any cluster resources while the restore operation is running.

1. Checks that the backup copy is verified.

If not, SnapManager checks if you have enabled the **Database Verification Override** option and continues.

**Attention:** You must restore only from verified backup copies to ensure a successful restore operation.

2. Creates a Snapshot copy of the volume to be restored so that the state before the restore operation is preserved.
3. Restores the data from the Snapshot copy of the LUN to be restored.
4. Restores the transaction logs that you need to replay to restore the data.

5. Uses the Exchange recovery process to play the transaction logs to the restored databases. On Windows, the soft recovery process is used, which is a transaction-log replay process that occurs when a database is remounted after an unexpected stop, or when transaction logs are replayed into an offline backup of a database.

## Types of SnapManager Restore operations

There are two types of SnapManager Restore operations. In a point-in-time restore operation, the transaction logs in the active file system not committed to the databases at the time of backup operation are replayed. In an up-to-the-minute restore operation, all uncommitted transaction logs are replayed and are applied to the databases. This option is selected by default.

### Point-in-time restore operation

In this case, only the uncommitted transaction logs that existed in the active file system at the time the backup copy was created are replayed. All the transaction logs beyond the point-in-time (after the backup copy was created) that exist in the transaction log directory and that belong to the restored database are removed.

You can use this method to restore a database back to a time before a corruption occurred.

### Up-to-the-minute restore operation

There are some transaction logs that are not committed to the databases. In an up-to-the-minute restore operation, all uncommitted transaction logs, from the time the backup set was created up to the most current time, are played forward and applied to the databases. This includes transaction logs from the backup sets, in addition to the transaction logs in the transaction log directory. A contiguous set of transaction logs is required for an up-to-the-minute restore operation to succeed.

#### Example: Up-to-the-minute restore operation

Suppose you run SnapManager Backup every day at noon, and on Wednesday at 4 p.m. you decide to restore from a previous backup set. For some reason, the backup set from Wednesday noon failed verification, so you decide to restore from the Tuesday noon backup set. After that backup set is restored, all the transaction logs are played forward and applied to the restored databases. This starts with those that were not committed when you created Tuesday's backup set and continues through the latest transaction log written on Wednesday at 4 p.m.

## Snapshot copies created during a restore process

Every time you perform a restore process, a Snapshot copy is created on the volume that contains the transaction logs.

The name of the Snapshot copy that is created during a restore operation contains the prefix `rstrsnap_`. After you verify that a restore process is completed successfully and you are satisfied with the results, you can delete the restored Snapshot copy using the **Delete** option in the Actions pane.

## Methods that can decrease restore process time

You can decrease the time that SnapManager requires to complete a restore operation by making more backup copies and ensuring that any LUN clone split operations in progress are complete.

The more backup operations you perform, the fewer Exchange transaction logs need to be played forward at restore time. At a minimum, you should perform one SnapManager full database backup operation every 24 hours.

Before proceeding with failback to the production site, ensure that any LUN clone split operations that are in progress are complete. The LUN clone split functionality supports significantly faster

online Snapshot copy restore operation when using SnapManager or SnapDrive to restore databases. By default, this functionality is enabled.

## Transaction log sequence verification options

You can make a restore operation faster by omitting the pre-restore transaction log sequence and database metadata checks. However, if you omit these checks, and a problem exists with your transaction log sequencing or your database metadata, the restore operation fails.

To perform an exhaustive check of the metadata, select the **Verify Transaction Log Sequence...** and the **Exhaustive Verification** check boxes in **Restore > Advanced Options**. For backward compatibility among SnapManager versions, the exhaustive check of metadata is selected by default.

To perform a quicker, less comprehensive check of the transaction logs, select **Verify Transaction Log Sequence...** and clear **Exhaustive Verification** in **Restore > Advanced Options**.

To omit metadata verification, clear the **Verify Transaction Log Sequence...** and the **Exhaustive Verification** check boxes in **Restore > Advanced Options**.

You have the following pre-restore transaction log sequence options available:

- Exhaustive transaction log sequence checks
- Non-exhaustive transaction log sequence checks

### Exhaustive transaction log sequence

The exhaustive transaction log sequence checks the transaction log headers and mounts the Snapshot copies for further verification.

SnapManager follows this sequence of steps:

1. Analyzes the transaction log headers.
2. Verifies the following conditions:
  - The transaction logs are all present and sequenced correctly.
  - The signatures of all the transaction logs match.
  - Each transaction log's create time matches the next log's previous time.
3. Mounts a Snapshot copy of the databases in this backup copy and analyzes the database headers.
4. Checks if the database signature matches the log database signature.

### Non-exhaustive transaction log sequence

The non-exhaustive transaction log sequence uses only the file names of the transaction log for verification. SnapManager verifies that the transaction log sequence is correct and that all log files are present.

This option is enabled if you clear the **Exhaustive Verification** check box in **Restore > Advanced Options**.

## LUN Clone Split Restore method

SnapManager uses the LUN Clone Split Restore method to restore only the affected LUNs in a volume.

In a LUN Clone Split Restore operation, the only data blocks duplicated during the restore operation are those currently in use in the active file system. The restore destination must exist in the active file system.

SnapManager follows these steps to restore a LUN to its LUN clone:

1. At backup time, SnapManager creates a LUN clone.  
The LUN clone is unavailable for data I/O. The LUN clone inherits the presentation of the LUN in the active file system.
2. At restore time, SnapManager renames the LUN clone with the name of the original LUN.
3. LUN Clone Split starts in the background.  
The renamed LUN clone becomes available for data I/O after the LUN Clone Split operation completes.

### Verifying the LUN clone split status

You can improve SnapManager restore performance by ensuring that any LUN clone split operation already in progress has been completed. You can determine the LUN clone split status by using the storage system console or SnapDrive.

#### Step

1. Use the `Data ONTAP lun clone split status` command on the storage system console.  
Additionally, SnapDrive displays the LUN Clone Split status in the **Operation Status** column of the SnapDrive user interface. For more information, see the relevant SnapDrive and Data ONTAP documentation.

### SnapManager restore operations in DAG environments

In a DAG, SnapManager does not take the entire Exchange DAG instance offline; instead, it takes the online database that is to be restored offline before performing a restore operation for the active copy of the database.

SnapManager performs the following steps when reseeding a Passive database:

1. Suspend the replication of the passive copy before starting the restore operation by using the `suspend-database` PowerShell command.
2. Restore the latest local Snapshot copy based on the volume layout.  
**Note:** An archived copy cannot be used.
3. Resume replication of the passive copy after the restore operation is completed using the `resume-database` PowerShell command.
4. The restore operation completes with passive database copy in a healthy state.

**Note:** The active copy of the database is not dismounted during the restore operation for the passive database copy.

The purpose of a passive database restore operation is to reseed the database more quickly using a Snapshot copy, based on the following reasons:

- The passive copy is corrupt, the exchange re-synchronization fails and it is unrecoverable without a reseed.
- Some environments with a slow WAN link across DAG nodes might require a long time to reseed a database copy.
- Restoring from a backup of a currently passive copy will recover the passive copy without changing the (newer) content of the active copy.

## Guidelines for using SnapManager restore

You must follow some guidelines when deciding whether to use an up-to-the-minute restore operation or a point-in-time restore operation. Guidelines also apply to the restore destination and the naming of the databases.

### Which type of restore operation to perform

You should consider your reasons for performing a restore operation so that you can decide which type of restore operation to perform: point-in-time restore operation or up-to-the-minute restore operation.

#### When to use an up-to-the-minute restore operation

- You want to play forward all the transactions up to the most current time.
- You want to restore individual databases.
- You want to restore backup copies after a point-in-time restore operation of a backup copy that is not the most recent.
- You want to perform an up-to-the-minute restore operation from any backup copy, including the most recent backup copy after a point-in-time restore operation of the most recent backup operation.

You might lose all transactions that occurred between the time when you created the last backup copy and when you performed the point-in-time restore operation.

**Note:** For an up-to-the-minute restore operation to succeed, a contiguous set of all required transaction logs must be in the SnapInfo folder and the transaction log directory of the database.

#### When to use a point-in-time restore operation

- You want to recover the databases as they were at a particular point in time: for example, when the most recent backup copy was created.
- You want to restore databases to a Recovery Database .
- You want to restore all existing backup copies after a point-in-time restoration of a backup that is not the most recent one.

## Guidelines for restoring from a SnapManager backup copy

You must follow some guidelines about the type of restore operation, the restore destination, and naming of databases before you begin to restore from a backup.

- Use the most recent backup copy for an up-to-the-minute restore process; otherwise, more transaction logs need to be replayed, and the restore process takes more time.  
The most recent backup copy's name contains `recent` if you did not use the unique naming convention; otherwise, it uses the most recent timestamp.
- If you rename a database after backing it up, you cannot recover it from that backup copy.  
If you must recover a database from a backup copy made before the database was renamed, return the database to its former name, then recover it from the backup copy by the same name.
- If you rename a database, make sure that you back it up as soon as possible.

**Note:** When you rename a database, unmount all of the databases. After they are all unmounted, you can remount the databases. At this point, Exchange recognizes the new database name, and SnapManager can make new backup copies.

- If you restore the database to the Recovery Database, you can perform only a point-in-time restore process.
- Store all databases that reside on the same LUN together.

## Restoring from a SnapManager backup copy

You can use SnapManager to restore databases from a backup copy in two ways.

- SnapManager Restore wizard
- Restore window accessed from the Actions pane

### Decisions to make before restoring from a SnapManager backup copy

When you perform a restore operation, SnapManager prompts you to provide some information like the restore object and destination, type of restore operation, and verification options that you need to keep ready with you before you make decisions during the operation.

#### Restore object and destination

- What items do you want to restore?  
You can select the production databases copies that you want to restore.
- Which database do you want to restore?
- To what location do you want to restore the backup sets?  
You can restore the backup sets to either one of the following locations:
  - To the same database or same server
  - To the Recovery Database

#### Type of restore operation

- Do you want to perform an up-to-the-minute restore operation (the default) or a point-in-time restore operation?
- Do you want to perform an actual restore operation or only a test restore operation?  
A test restore operation gives you a preview of the actual restore process to help you decide if you want to go ahead with your current settings.
- If you want to perform an actual restore operation, do you want to mount your databases after the restore operation is complete?

#### Verification options

- Do you want to verify the transaction log sequence and database metadata before the restore operation?  
SnapManager recovers the databases if you select **Recover and mount databases after restore**. If you clear this option, SnapManager recovers only the databases, but does not mount them.
- If you want to perform a test up-to-the-minute restore operation, do you want to verify the current transaction logs?

The option to verify the current transaction logs confirms that all the required transaction logs are present.

- If the backup copy is unverified, do you want to verify it before the restore operation?
- Which verification settings do you want to apply for restoring backup sets?
- Do you want to select SnapMirror destination volumes for integrity verification?

If yes, click the **Verification Server** tab in the Database Verification Settings dialog box to select SnapMirror destination volumes.

## Mount options

When you perform a restore operation, you can specify the mount options after restore. You can select the option to enable SnapManager to mount the databases immediately after completing the restore operation or you can mount the databases manually at a later time.

When you are restoring backups that were created on other Exchange servers or restoring backup sets from an archive, you have the additional option to associate all mailboxes in the newly restored databases to the currently connected Exchange server.

## Restoring databases using the Restore wizard

You can restore a database from a backup set by using the SnapManager Restore wizard that guides you through the restore process.

### Before you begin

- Close all windows on the Exchange server that runs SnapManager.
- Disable any SnapManager operations that are scheduled to run against the Exchange data that you want to restore, including any jobs scheduled on the remote administration server or the remote verification server.

### About this task

You can restore only one database at a time in a Database Availability Group.

You can optionally perform a SnapManager backup and verification operation after restoration to verify that your restored database is free of physical-level corruption.

Do not perform a restore operation while a backup operation is in progress. If you cancel a current backup operation in progress, SnapManager pauses all the active scheduled backup jobs on the Exchange server, or on all nodes in the DAG environment, and cancels the current backup copy before performing the restore operation. On completing the restore operation, SnapManager reenables the paused scheduled backup jobs. All the other inactive jobs are not changed.

If you want to restore after the current backup operation completes, SnapManager pauses all the active scheduled backup jobs on the current Exchange server (on all nodes in the DAG environment) but waits for the completion of the current backup operation before performing the restore operation. On completing the restore operation, SnapManager reenables the paused scheduled backup jobs. All the other inactive jobs are not changed.

If you select a backup that is or has any archived backups, the Results pane displays the following details about them:

- EDB primary storage
- EDB primary LUN path
- EDB secondary storage
- EDB secondary qtree path

- EDB secondary Snapshot copy name

The EDB secondary storage, EDB secondary qtree path, and EDB secondary Snapshot copy name are used to mount the Snapshot copies using the SnapDrive CLI. To recover a mailbox from the archived backup, use the Single Mailbox Recovery tool.

### Steps

1. In the **Scope** pane, select the Exchange server node; then click **Restore**.

SnapManager displays the Exchange Server databases. When the SnapManager connects to the DAG, all backups in the DAG are displayed and when it connects to a member server of the DAG, only the backups on that server are displayed.

2. Double-click the database that you want to restore.
3. In the **Actions** pane, click **Restore wizard**.
4. Follow the instructions in the Restore wizard and go to the **Restore Status** window.
5. After the restore process is complete, click **OK**.

Your Exchange server comes back online.

## Restoring databases using the Restore window

Unlike using the Restore wizard to restore your databases, using the Restore window enables you to configure the restore process on your own.

### Before you begin

- Close all windows on the Exchange server that runs SnapManager.
- Disable any SnapManager operations that are scheduled to run against the Exchange data that you want to restore, including any jobs scheduled on remote management or remote verification servers.

### About this task

Do not perform a restore operation while a backup operation is in progress. If you cancel a current backup operation in progress, SnapManager pauses all the active scheduled backup jobs on the Exchange server, or on all nodes in the Database Availability Group environment, and cancels the current backup copy before performing the restore operation. After completing the restore operation, SnapManager re-enables the paused scheduled backup jobs. All the other inactive jobs are not changed.

If you want to restore after the current backup operation completes, SnapManager pauses all the active scheduled backup jobs on the current Exchange server (on all nodes in the DAG environment) but waits for the completion of the current backup operation before performing the restore operation. On completing the restore operation, SnapManager re-enables the paused scheduled backup jobs. All the other inactive jobs are not changed.

After an actual restore process, you can optionally perform a backup and verification operation to verify that your restored database is free of physical-level corruption.

### Steps

1. In the Scope pane, click **Restore**.
2. Select the database that you want to restore.

If you want to restore from an archived backup copy and have datasets configured in your system, select the SnapVault and SnapMirror-enabled storage systems from the listed archived backup copies.

3. Double-click the backup copy under the database that you want to restore.
4. In the **Actions** pane, select **Restore**.
5. In the **Choose databases to restore** pane, select the databases that you want to restore.
6. Under **Type of Restore**, select if you want an up-to-the-minute restore or a point-in-time restore operation.
7. If you do not want to perform exhaustive verification of the transaction log sequence and database metadata before the restore process, click **Advanced Options**.
8. Under **Operation Options**, clear the **Verify Transaction Log Sequence and Database Metadata Before Restore** and **Exhaustive Verification** check boxes.
9. Coordinate the backup and restore processes using the **Job Control Options** pane:

If you want to...	Then do this...
<b>Cancel the current backup operation in progress</b>	Select <b>Cancel conflicting backup that is in progress</b> .
<b>Restore after the current backup operation is complete</b>	Select <b>Wait for running backup to complete</b> .
<b>End the restore operation</b>	Select <b>Abort restore if conflicting operation is running</b> .

10. If the restore server does not have access to the archived backup storage, click **Advanced options**, and then under **Archived Backup Access**, select **Restore server does not have access to the archived backup storage**.
11. If you are performing a test restore operation, click **Test Restore**.
  - a. To check the current logs (for up-to-the-minute restore operations only) and run database verification, in the **Select Test Restore Options** window, select **Check current Logs**.
  - b. To verify databases and transaction logs on available destination volumes, select **Run Verification on Computer**.

SnapManager displays the Restore Status window, showing the tasks that are performed for the restore process.

12. Click **Start Now**.

### Restoring databases to a specified FRP

You can use the Restore window to restore a point-in-time Frequent Recovery Point (FRP) database backup that is created after a full backup or copy backup operation using the Restore window. SnapManager combines the restore operation of a full database backup or copy backup and the required transaction logs to restore to the selected recovery point. You can also do a test restore to check the current logs (for up-to-the-minute restore operations only) and run database verification.

#### Before you begin

You must know the FRPs of the backup copies that you want to restore.

No backup operations can be running.

### About this task

Do not perform a restore operation while a backup operation is in progress. If you cancel a current backup operation in progress, SnapManager pauses all the active scheduled backup jobs on the Exchange server, or on all nodes in the DAG environment, and cancels the current backup operation before performing the restore operation. On completing the restore operation, SnapManager reenables the paused scheduled backup jobs. All the other inactive jobs are not changed.

If you want to restore after the current backup operation finishes, SnapManager pauses all the active scheduled backup jobs on the current Exchange server (on all nodes in the DAG environment) but waits for the completion of the current backup operation before performing the restore operation. On completing the restore operation, SnapManager reenables the paused scheduled backup jobs. All the other inactive jobs are not changed.

After a restore process, you can optionally perform a backup and verification operation to verify that your restored database is free of physical-level corruption.

### Steps

1. In the **Scope** pane, click **Restore**.
2. Select the databases that you want to restore, and then double-click the backup sets that you want to restore.
3. In the **Actions** pane, select **Restore**.
4. In the **Choose databases to restore** pane, select the databases that you want to restore.
5. Under **Type of Restore**, select **Point in time**.
6. Click **Select Recovery Point**.
7. In the **Scope** area of the **Recovery Point Selection** list, select the date, start time, and end time.
8. Click **Refresh**.
9. Select a recovery point from the **Available Recovery Points** list.

SnapManager restores the transaction logs up to the selected recovery point. These logs are later rolled forward by Exchange during recovery.

10. Click **OK**.
11. To coordinate the backup and restore processes, click **Advanced Options**.

If you want to...	Then do this...
<b>Cancel the current backup operation in progress</b>	Select <b>Cancel conflicting backup that is in progress</b> .
<b>Restore after the current backup operation is complete</b>	Select <b>Wait for running backup to complete</b> .
<b>End the restore operation</b>	Select <b>Abort restore if conflicting operation is running</b> .

12. In the **Operation Options** pane, select one of the following to configure exhaustive verification and mount settings:

If you want...	Then do this...
To perform exhaustive verification of the transaction log sequence and database metadata before the restore process	Select <b>Verify Transaction Log Sequence and Database Metadata Before Restore and Exhaustive Verification.</b>
SnapManager to automatically remount your restored databases immediately after the restore operation	Select <b>Recover and mount database after restore.</b>

13. If you are performing a test restore operation, click **Test Restore**.
  - a. To check the current logs (for up-to-the-minute restore operations only) and run database verification, in the **Select Test Restore Options** window, select **Check current Logs**.
  - b. To verify databases and transaction logs on available destination volumes, select **Run Verification on Computer**.

SnapManager displays the Restore Status window, showing the tasks that are performed for the restore process.
14. Click **Start Now**.

#### Recovery of a database to an FRP

When you restore a full or copy backup, you can specify the recovery point from the subsequent FRP backups. SnapManager performs a point-in-time restore operation and recovers the database up to the specified recovery point.

## Restoring Exchange databases in a DAG

You can restore Exchange Server databases in a DAG from backup copies created on the same Exchange Server or from backup sets created on multiple Exchange Servers in the DAG. You can restore the database to a number of different destinations in addition to its original location.

You can restore a database in a DAG from the following sources:

- SnapManager backup copies created on the same Exchange Server
- Backup sets created on multiple Exchange Servers in the DAG

You can restore data to a specified Frequent Recovery Point (FRP). You can also restore a database to the Recovery Database.

#### Backups available for restore in a DAG

When the SnapManager MMC snap-in connects to a Database Availability Group (DAG), you get a consolidated view of all backups in the DAG. You can then select a backup copy to restore.

Backups are created on the server where the database is located. When the SnapManager MMC snap-in connects to a member server of the DAG, you can view backups specific to that server. For each database that is backed up, its backups are grouped under two headings: **Local Backups** and

**Archived Backups**. Because the backup name includes the name of the Exchange server, it indicates where the backup is located.

**Note:** The database restore takes place on the server where the backup is located.

## Restoring a backup copy created on a different server

You can use the SnapManager Restore Wizard to restore backup copies that were created on Exchange to another Exchange mailbox server in the same organization.

### About this task

If an Exchange Server (mailbox server) was down due to physical server crash, you must disconnect the LUNs, remap the LUNs to another standalone Exchange Server, then run the restore operation from the Restore Wizard to recover the database.

**Note:** This is applicable only for standalone mailbox servers.

If the backup set is not verified, SnapManager gives you the option to verify the backup set before the restore operation.

After the restore process is completed, you can optionally perform a backup and verification operation to verify that your restored database is free of physical-level corruption.

### Steps

1. In the **Scope** pane, select an Exchange server.
2. In the **Actions** pane, click **Restore wizard**.
3. In the **Which Exchange Server Created the Backups** window, select the **Restore backup created on a different server** option.
4. In the **Exchange Server Where Backups were Created** window, specify the **Exchange server where the databases were backed up** and the **SnapInfo Directory Path**.
5. Select a backup set that you want to restore.
6. In the **Where to restore** panel, select a restore option.
7. Follow the instructions in the Restore Wizard, and go to the **Mount Options** page.
8. To recover and mount a database after the restore process, select **Recover and mount database after restore**.
9. Do not select **Update user accounts associated with mailboxes in restored databases to point to mailbox server with the new name**.

**Note:** If you restore the backup created on a different server, execute the `Set-Mailbox` cmdlet outside of SnapManager for Exchange to associate all of the mailboxes in the newly restored databases to the currently connected Exchange server. Use the following command:

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

## Reseeding a failed database copy in a DAG

You can reseed a passive database copy that is in a failed state to restore it to the latest local Snapshot copy and place the passive database copy in a healthy state.

### Before you begin

- The passive database copy must be in a failed state (Failed, FailedandSuspended, or Suspended) before running the reseed operation.
- The database copy that is being reseeded must be a passive copy only.
- The Snapshot backups must be available.


### About this task

The reseeding process implemented in SnapManager restores an existing passive copy that is in a failed state (Failed, FailedandSuspended, or Suspended) to the latest local Snapshot copy (an archived copy cannot be used) to make it healthy.

Reseeding a passive database copy can occur in the following database configurations:


- Databases residing in different LUNs on different volumes

In this case, the reseed workflow performs a volume restore operation that uses the snapshot-restore-volume ZAPI:



- Databases residing in different LUNs on the same volume

In this case, the reseed workflow performs a LUN restore operation that uses the snapshot-restore-file ZAPI:



For Data ONTAP operating in 7-Mode, SnapManager performs the LUN restore operation regardless of the database configuration.

### Steps

- In the **Console Root** tree, connect to the DAG node where the failed database copy resides.

**Note:** If the DAG node is not already added, navigate to the **Add exchange servers to be managed** option.

2. In the **Actions** pane, click **Reseed Wizard**.
3. Coordinate the reseed process using the **Job Control Options** pane:

**Note:** Only failed or suspended database copies are listed. If all of the databases are healthy, a dialog box informs you that **No suspended/Failed databases exist on the selected server**, and the reseed wizard will close. Try to connect to the DAG node where the failed database copy resides again.

If you want to...	Then do this...
<b>Cancel the current backup operation in progress</b>	Select <b>Cancel conflicting backup that is in progress</b> .
<b>Reseed after the current backup operation is complete</b>	Select <b>Wait for running backup to complete</b> .
<b>End the reseed operation</b>	Select <b>Abort reseed if conflicting operation is running</b> .

## Using a Recovery Database

A Recovery Database is a special type of mailbox database that allows you to restore, recover, and mount your database to an Exchange server while the original database is still serving the user. It is a feature available with Exchange Server.

A recovery database is not bound to any server or database, and each Exchange Server can have only one recovery database.

SnapManager automates the mailbox store recovery process. You can add the Exchange mailbox store through management shell cmdlets on Exchange Server, then you can restore the mailbox from the backup copy created and stored in the Recovery Database.

You can restore a database to a Recovery Database by using the `restore-backup` cmdlet with the appropriate parameter in the SnapManager command-line interface.

### Limitations of using a Recovery Database

A Recovery Database enables you to restore, recover, and mount your database to an Exchange server to retrieve mail while the original database is still serving users. These abilities are limited in terms of the number of databases that you can restore, the type of restore operation that you can perform, and the configuration of the restore operation.

- You can use the Recovery Database for Exchange mailbox databases.
- You can use the Recovery Database as a target for restore operations but not for backup operations.
- You can mount only one Recovery Database at any time on an Exchange Server .

## Restoring a mailbox database to the Recovery Database in Exchange Server

You can use the Restore wizard to restore an Exchange database to the Recovery Database.

### About this task

If the backup set is not verified, SnapManager gives you the option to verify the backup set before the restore operation.

You can restore from the following restore sources:

- Backup that was created on the same Exchange server
- Backup that was created on a different Exchange server

After an actual restore process, you can optionally perform a verification operation to verify that your restored database is not corrupted.

### Steps

1. Either select an Exchange server node in the **Scope** pane or connect to the DAG.
2. In the **Actions** pane, click **Restore Wizard**.
3. In the **Which Exchange Server Created the Backups** window, select the restore source.
4. In the **Choose the backup that you want to restore** window, double-click the backup set listed under the database.
5. In the **Select the target for backup restoration** window, select **Restore to the Recovery Database**.
6. In the **Choose the items to be restored** window, select the check boxes corresponding to the production databases that you want to restore.
7. In the **Choose Recovery Database** window, select the destination Exchange server to which you want to restore the backup set.

**Note:** You cannot specify a DAG name as a destination Exchange server.

8. Specify the name of the new Recovery Database to which you want to restore the selected backup set.
9. In the **Type of Restore** window, select a recovery point for your database.
10. Configure the verification settings in the **Verify options** window:
  - a. To check the current logs (for up-to-the-minute restore operations only) and run database verification, in the **Select Test Restore Options** window, select **Check current Logs**.
  - b. If you do not want to perform exhaustive verification of the transaction log sequence and database metadata before the restore process, clear the **Exhaustive verification** check box.
  - c. If the restore server does not have access to the archived backup storage, click **Advanced options**, and then under Archived Backup Access, select **Restore server does not have access to the archived backup storage**.
  - d. To coordinate the backup and restore processes, make the following selections in the Job Control Options pane:

If you want to...	Then do this...
<b>Cancel the current backup operation in progress</b>	Select <b>Cancel conflicting backup that is in progress</b> .
<b>Restore after the current backup operation is complete</b>	Select <b>Wait for running backup to complete</b> .
<b>End the restore operation</b>	Select <b>Abort restore if conflicting operation is running</b> .

11. For an actual restore process, if you want SnapManager to automatically recover and remount your restored databases immediately after the restore process, select **Mount Database Automatically After Restore** in the **Mount Options** window.
12. Verify the restore settings and then click **Finish**.
13. To start the restore process, click **Start Now**.

#### After you finish

After restoring the mailbox database to the Recovery Database, SnapManager will restart the Exchange Information Store and Replication services to make the newly restored mailbox database visible to you.

## When to delete a Recovery Database

Because each Exchange Server can have only one Recovery Database mounted at any given time, after you recover the mailbox from the Recovery Database, you must delete the Recovery Database. You can then recover other mailbox databases.

You can delete the Recovery Database using the SnapManager MMC.

## Restoring mailboxes from unmanaged media

If you archived a SnapManager backup to a Disaster Recovery site, you can still restore from this unmanaged media.

#### About this task

This archived backup might be on tape or another type of cold storage, and no longer tracked by SnapManager. However, you can restore a mailbox or a group of mailboxes and mount this backup back into the Exchange Recovery Database.

Use the following general process as a guideline for restoring a backup from unmanaged media.

#### Steps

1. Create temporary LUNs using SnapDrive.
 

**Note:** For detailed information, refer to your SnapDrive documentation.
2. Restore the database and log LUNs to the temporary LUNs. This might involve restoring from tape.
3. Run the Exchange utility (`eseutil.exe`) with the `/r` option to roll the logs into the database.
 

**Note:** For detailed information about `eseutil.exe`, refer to your Microsoft Exchange documentation or enter `ESEUTIL /?`.
4. Use the Exchange admin tool or Powershell to mount this Store to the Recovery Database.
5. Recover the data.

This involves using Exchange tools or Powershell to recover mailboxes or mailbox items.

6. Dismount the Recovery Database.
7. Delete the temporary LUNs.

## Troubleshooting restore operations

If a restore operation fails with an error message similar to `Error: 0xC00413CA, SnapManager detected the following Exchange writer error. Please retry SnapManager operation. VSS_E_WRITERERROR_RETRYABLE`: (and additional information), you might be able to fix the problem and try to restore again.

### About this task

A `restore.env` log file is created during a restore operation. If you later restore the same data and that log file is still present, the operation fails.

### Steps

1. Check the transaction log LUN for the database for a `restore.env` file: for example, `E01restore.env`.
2. If such a file is present, delete it and retry the restore operation.

## Mailbox restore operation using Single Mailbox Recovery

---

You can use SnapManager for Exchange to launch the Single Mailbox Recovery application to restore your mailbox data. Using Single Mailbox Recovery, you can locate and restore items at any level of granularity directly to an existing mailbox on your Exchange server.

Single Mailbox Recovery works in conjunction with SnapManager for Exchange, allowing you to restore individual Exchange mailbox items that were previously deleted.

For more information about restoring individual Exchange mailboxes using the Single Mailbox Recovery, see the appropriate *Single Mailbox Recovery Administration Guide* for your system.

### Recovering mailbox data

When you need to restore mailbox items to your Exchange server, you can launch the Single Mailbox Recovery application using SnapManager for Exchange to locate and restore these items.

#### Before you begin

You must have the supported version of Single Mailbox Recovery installed on the server that is running the SnapManager for Microsoft Exchange user interface, MMC Snap-in.

Also, if you are using Protection Manager, the source and destination Single Mailbox Recovery servers must be configured for Protection Manager access. See the `sdcli dfm_config` command for help on how to configure SnapDrive to access a Protection Manager server.

**Note:** The Single Mailbox Recovery mount of an archived Snapshot copy currently works on the local node only. If you try to mount an archived Snapshot copy from a remote node, Single Mailbox Recovery mount fails. For example, if you have archived a Snapshot copy on Node 1 and you try to mount this archived copy on Single Mailbox Recovery from the same node, it works properly; however, if you try this mount operation on Node 2, the Single Mailbox Recovery mount fails.

#### Steps

1. Select the Exchange server node in the **Scope** pane, and then select **Restore**.  
SnapManager displays the Exchange server databases in the center pane.
2. Select the name of the backup that you want to restore and click **Run SMBR** in the **Actions** pane.  
**Note:** When SnapManager and Single Mailbox Recovery are installed on the same host, Single Mailbox Recovery usage is tracked through the use of AutoSupport (ASUP) messages.

The Start SMBR dialog box opens.

3. Provide the information required by the Single Mailbox Recovery software such as the database name and (optionally) the transaction log usage. If the **Usage of transaction logs** check box is selected, clear it. (The transaction logs require the software mount the log LUN, but because this is a point-in-time restore, the log LUN has already been deleted, and so the mount fails.)
4. Click **Start SMBR** to launch the Single Mailbox Recovery software and start the mailbox recovery.

#### After you finish

After the Single Mailbox Recovery application restores your mailbox items, you must unmount the Snapshot copies you restored.

## Cleaning up after Single Mailbox Recovery

If you want to recover mailbox data, Single Mailbox Recovery needs to mount the mailbox backup copy that contains the appropriate data. After your mailbox data is recovered, you need to unmount the backup copy. Mounted backup copies use resources and cannot be deleted.

#### About this task

Unmounting the backup copies actually renames the backup copy of the resource group by appending the backup copy name with a timestamp.

#### Steps

1. Select the Exchange server node in the **Scope** pane, and then select **Restore**.  
SnapManager displays the Exchange server the databases in the center pane.
2. Click **SMBR Cleanup** in the **Actions** pane.  
**Note:** When SnapManager and Single Mailbox Recovery are installed on the same host, Single Mailbox Recovery usage is tracked through the use of ASUP messages.
3. Select the LUNs that you want to unmount, and then click **Dismount**.  
To select all the LUNs that were mounted during the restore process, click **Select All**.
4. Close the window when the progress bar indicates that the cleanup process is complete.

## **What SnapManager reports contain**

SnapManager reports list the step-by-step details of every SnapManager operation that you perform, their final statuses, and any error messages that you encounter during the operation.

The SnapManager Report Directory provides folders that group the reports for each of these operation types:

- Backup
- Config
- Debug
- Frequent Recovery Point
- Delete Snapshot
- Restore
- Reseed
- Miscellaneous

**Note:** There might be other folders in the directory if you are running SnapManager in a DAG.

## **SnapManager reports in a DAG**

When the SnapManager MMC snap-in is connecting to the Database Availability Group (DAG), you can view all the reports from all the member servers of the DAG.

The report folders reside on each member server of the DAG. The report folders for the DAG reside on the "owner node" of the DAG. Whenever configuration, backup, or restore operations are performed, the reports generated for the DAG are stored in the report folders of the server which is the owner node of the DAG at the time these operations were performed.

## **Reasons to change the report directory location**

You might need to change the report directory location, either because of limited space or because you want to share the directory between SnapManager and the Exchange server running in a DAG environment.

If you find you have limited space in the current report directory, you can move it to a different location that has more available disk space.

## **Changing the SnapManager report directory**

If you have limited space in your original report directory, you can change the report directory in which to store your reports.

### **About this task**

If you change the directory, you can no longer see the reports that were created before you changed the path. To view any such reports, change the report directory path back to the original path and refresh the view.

**Steps**

1. In the **Actions** pane, click **Report Directory Settings**.
2. Enter or browse to the path name for your new Report Directory.

**Note:** The directory cannot be located on a CIFS share.

3. In the **Actions** pane, click **Refresh**.

## Locating the report directory in a DAG

By default, the SnapManager Report Directory is on the disk on which SnapManager is installed.

**About this task**

- Change the directory from every SnapManager node.
- Use a disk that does not contain Exchange or SnapManager data for the report directory.  
The report directory is restored from its Snapshot copy when you perform a restore operation.

**Steps**

1. In the **Actions** pane, click **Report Directory Settings**.
2. Enter or browse to the new report directory path.
3. In the **Actions** pane, click **Refresh**.

## Viewing SnapManager reports

You can view SnapManager reports from the SnapManager GUI.

**Steps**

1. In the **Scope** pane, click **Reports**.  
SnapManager displays the report folders in the Results pane.
2. In the report folders, double-click the database for which you want to view a report.  
SnapManager displays the report in the Results pane.

**Note:** The reports display in descending order (the latest report displays first).

## Printing SnapManager reports

You can print SnapManager reports either from within the displayed report, or you can open the report in Notepad and print it.

**Steps**

1. In the Scope pane, click **Reports**.
2. Click the directory that contains the report you want to print.

If you want to...	Then...
<b>Print the displayed report directly</b>	<ul style="list-style-type: none"> <li>a. Select the report that you want to print.</li> <li>b. Right-click anywhere within the displayed report in the Result pane, and select <b>Print</b>.</li> </ul>
<b>Print from Notepad</b>	<ul style="list-style-type: none"> <li>a. Right-click the report name in the Result pane and select <b>Open with Notepad</b>.</li> <li>b. From Notepad, select <b>Print</b> in the File menu.</li> </ul>

## Deleting SnapManager reports

You can delete reports from the Report Directory to increase the space available in the directory.

### Steps

1. Click the reports in the Scope pane.
2. Select the report you want in the SnapManager Report Directory.
3. Right-click the report directory or the individual report that you want to delete and select **Delete All** or **Delete**.
4. Click **Yes** in the dialog box that appears.

The report is deleted.

## Control of report auto-archiving

SnapManager can reduce the amount of space reports require by archiving the reports based on their age or number. You can control this process to tune the amount of space used.

The SnapManager Report Directory Setting options window also includes controls for the report auto-archive function. You can set how often the reports directory is checked, and whether reports are archived based on either age or number of reports.

Auto archiving does not start until you set the auto-archiving options.

### Frequency

controls when the reports directory is checked for reports to archive. You can set the frequency to between 0 and 60 minutes, or 0 to 60 hours.

### Older than (days)

when set, specifies after how many days a report is archived. The available range is from 0 to 30 days.

### In excess of

when set, specifies how many reports remain unarchived. If there are more reports than this value, reports are added to the archive, oldest first, until the set number of reports remain.

There is a single archive file, and it is stored in the reports directory. The name of the file is the name of the report directory plus .zip.

# Modifying your database configuration on NetApp storage

---

The SnapManager Configuration wizard enables you to select database verification servers, move Exchange databases and transaction logs to Data ONTAP LUNs, and configure automatic event notification. You must run SnapManager from the system console.

## Moving Exchange data with the Configuration wizard

You might need to move Exchange data from a LUN to a local disk because the data is no longer managed, or you might need to move Exchange databases from a local disk to a LUN. Resource management issues might require a LUN-to-LUN transfer of Exchange data. In these cases, you can use the SnapManager Configuration wizard to move your Exchange data.

### From local disk to LUN

If you want to move Exchange databases, you can use the wizard to unmount the mailbox databases or the Exchange databases , move the Exchange databases and transaction log files to the selected LUN, and remount the databases.

**Note:** SnapManager takes Exchange databases offline during the transfer operation.

The wizard creates a SnapInfo directory that SnapManager uses to store information about the backup sets and the backed up transaction logs.

The wizard disables circular logging for all mailbox databases and databases that are moved to a LUN, to enable up-to-the-minute restoring of those databases and mailbox databases.

The wizard also guides you through several application settings. These settings include enabling notification of SnapManager events using e-mail, Syslog, or the AutoSupport feature.

### From LUN to LUN

Resource management issues might require LUN-to-LUN transfer of Exchange data. An example of a situation that requires a LUN-to-LUN transfer is when you want to consolidate the Exchange data on another storage system.

### From LUN to local disk

You need to move Exchange data from a LUN to a local disk if the data is no longer managed by using SnapManager.

**Note:** Protection Manager integration can be enabled without reinstalling SnapManager for Exchange anytime by rerunning the Configuration wizard. You must first ensure that you have configured SnapDrive for Protection Manager integration. If you have not already done so, you can set up SnapDrive Protection Manager integration via the `sdccli` command, and then restart the SnapDrive service, or you can rerun the SnapDrive setup. For more information on Protection Manager integration, see the SnapDrive help.

Make sure all databases are migrated to LUNs from a local disk. An example of a heterogeneous configuration is Exchange databases residing on LUNs and local storage simultaneously. Exchange databases must reside on NetApp LUNs to enable SnapManager to back up.

On Exchange Server , use the same drive letters or mount points for the Exchange data LUNs on all the nodes of a DAG.

Install SnapManager and SnapDrive for Windows on all member servers of the DAG in Exchange Server. If you plan to use NetApp storage on a couple of nodes and non-NetApp storage for remaining nodes, SnapDrive for Windows is required on all the nodes; however, SnapManager for Exchange is not required. SnapManager can be installed on the node where you plan to run the backup, which is connected to NetApp storage.

- You need to install SnapDrive for Windows on all nodes. Also, make sure there is proper network connectivity to the storage system from all of the nodes.
- You will not be able to use SnapManager to migrate databases (using the Configuration wizard) if SnapManager is not installed on all of the nodes.

You will not be able to use the SnapManager DAG functionality if you do not install SnapManager on all nodes; you use SnapManager only on an Exchange server which has SnapManager installed.

## Moving an individual database to a LUN

The Configuration wizard enables you to move an individual database to a LUN. You use the Select a Database to move to a LUN page to move an individual database to a LUN.

### About this task

The list of available drives changes as you make choices that rule out other choices.

### Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. In the **Disk** list, select the drive letter for the disk to which you want to move the database.
3. Click **<=>**.  
The selected database and the disk gets associated, and the database files are then moved to that disk, when the Configuration wizard moves your data.
4. Step through the wizard, and review the configuration summary in the **Completing the Configuration Wizard** page, and click **Finish**.
5. Click **Start Now**.  
**Note:** If you are moving the location of Exchange databases in this step, it could take some time to complete.
6. Click **OK**.

## Changing the location of a database

The Configuration wizard enables you to change the location of a database that you had associated with a LUN or local disk, if you want to change your configuration plan.

### About this task

The list of available drives changes as you make choices that rule out other choices.

If the Exchange system file path is on the database LUN (placed for a previous version of SnapManager), and you use the Configuration wizard to reconfigure the path of the transaction log, then the Configuration wizard automatically moves the system file path to the same LUN as the file path of the transaction log.

### Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. Under **Database Location Results**, highlight the name of the database that you want to move, and click **Reconfigure**.  
The selected database is returned to the list of unconfigured databases, and you can re-associate it with any available disk.
3. Proceed through the wizard, review the configuration summary in the **Completing the Configuration Wizard** page, and click **Finish**.
4. Click **Start Now**.  
**Note:** If you are moving the location of Exchange databases in this step, it could take some time to complete.
5. Click **OK**.

## Viewing the full path for database files

To migrate a database, you must verify that the database location is correct. The Configuration wizard enables you to view the full path for your database files.

### Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. In the **Configuration** wizard, select a database under **Database Location Results**, and click **Details**.  
The Database Settings window appears for the selected database.
3. Click **Details**.  
The current path and file name for the selected database are displayed, along with the new path and file name if the location of the selected database has changed.
4. Step through the wizard, review the configuration summary in the **Completing the Configuration Wizard** page, and click **Finish**.
5. A dialog box will appear. Click **Start Now**.  
**Note:** If you are moving the location of Exchange databases in this step, it could take some time to complete.
6. Click **OK**.

## Dataset protection policy

You have to specify the protection policy for your datasets using the Configuration wizard. The protection policy assigned to a dataset determines how replications are done. It also determines how long backups are retained on the secondary storage SnapVault target.

You can assign any these protection policies to the dataset:

- Back up
- Remote backups only

- Any customized policy based on the above two policy templates

Both the **Back up** policy and the **Remote backups only** policy have scheduled SnapVault replication that happens in the background along with "on demand" SnapVault replication that happens when you do a backup or verify operation with SnapManager.

The "backup retention" type determines the duration for which the archived backups are retained, and which backups are verified during a "verification only" operation.

## Migration of transport database paths for Exchange Server

The Configuration wizard cannot be used to migrate transport database paths for Exchange Server . The Configuration wizard does not support migration of transport database paths of Exchange transport server role.

When SnapManager runs on Exchange Server , the option to Configure Queue Database is skipped on the Configuration wizard. You cannot access this option when you connect to the Exchange server.

To change the location of the queue database and the queue database transaction logs, you can modify the following keys in the %ExchangeInstallPath%Bin\EdgeTransport.exe.config file:

- QueueDatabasePath <path>
- QueueDatabaseLoggingPath <path>

## Moving transaction logs to a LUN

You can move the transaction logs to a LUN by using the Configuration wizard.

### Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. From the **Database** list in the Configuration wizard, select the database whose transaction logs you want to move.
3. From the **Disk List**, select the drive letter for the LUN or local disk to which you want to move the transaction logs.
4. Click <=>.

The transaction logs for the selected database and LUN gets associated, and the transaction log files are then moved to that LUN when the Configuration wizard moves your data.

5. Step through the wizard, review the configuration summary in the **Completing the Configuration Wizard** page, and click **Finish**.
6. Click **Start Now**.

**Note:** If you are moving the location of Exchange databases in this step, it could take some time to complete.

7. Click **OK**.

## Changing the location of transaction log files

To effectively manage your transaction log files, you can reconfigure their location using the Configuration wizard.

### Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. In the **Configuration** wizard, go to **Transaction Logs Location Results**, select the database for the transaction files that you want to move, and click **Reconfigure**.  
The selected database is returned to the list of unconfigured transaction logs, and you can re-associate it with any available LUN or local disk.
3. Proceed through the wizard, review the configuration summary in the **Completing the Configuration Wizard** page, and click **Finish**.
4. Click **Start Now**.  
**Note:** If you are moving the location of Exchange databases in this step, it could take some time to complete.
5. Click **OK**.

## Viewing the full path for transaction log files

You can view the full path of the transaction log files from the Configuration wizard.

### Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. In the Configuration wizard, select a database in **Transaction Logs Location Results**, and then click **Details**.  
The **Transaction Logs Path** window appears for the selected database. If the location of the transaction log files has changed, the current paths and file names for the transaction logs, for the selected database are displayed, along with the new paths and file names.
3. Step through the wizard, and review the configuration summary in the **Completing the Configuration Wizard** page, and click **Finish**.
4. Click **Start Now**.  
**Note:** If you are moving the location of Exchange databases in this step, it could take some time to complete.
5. Click **OK**.

## Configuration using the SnapManager control file

SnapManager creates a control file that contains configuration information about the Exchange server. You can use the control file either to import or to export the Exchange server configuration.

The configuration data in the control file is represented in XML format. You can edit the file manually using an XML editor.

The configuration data in the control file is grouped into the following sections, so that you can manage it more easily:

- Storage layout settings
- Notification settings
- Verification settings
- Report directory settings
- Backup settings
- Scheduled jobs
- SnapMirror relationship settings

Instead of running the Configuration wizard every time you want to migrate databases and transaction logs to LUNs, you should use the control file in the following scenarios:

- Large configurations
- Mass deployment

## Configuration of SnapManager in a DAG using a control file

You can use a control file instead of the Configuration wizard to configure SnapManager. The control file support for Exchange Server mailbox databases is based at the Database Availability Group (DAG) level, and not at a member server level.

You can export the configuration settings of a DAG to another DAG or import the configuration settings from another DAG. This can be done by using either a SnapManager snap-in or a SnapManager for Exchange PowerShell cmdlet.

Because all database instances in all member servers of a DAG have to be migrated at the same time to keep database replication operating properly, the control file support for Exchange Server mailbox databases is DAG-based, not member server-based. It is not possible to export the configuration of a member server in a DAG to a control file and then import that control file to another member server in the same DAG.

## Importing Exchange Server configuration information using the control file

You can configure an Exchange server by using a control file containing configuration settings exported from another Exchange server. You can import configuration details and specific sections of the Exchange configuration from a control file.

### Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. Select **Use Control File**.
3. To import configuration details for the server, click **Next**.  
The Import or Export Selection window appears.
4. Select **Import**.
5. You can review your configuration settings by selecting **Review current settings in the Configuration wizard**.
6. In **Use control file**, either type the complete path of the control file or use the **Browse** feature to locate and select the file path.

7. Click **Advanced**.
8. In the **Configuration Import/Export Advanced Options** window, specify the configuration settings that need to be imported, or exported.
9. Select **OK**.
10. Click **Next** and then, to load the control file and validate the imported configuration and settings, click **Apply**.
 

If you do not want to apply the new configuration, click **Cancel**.

If you click **Apply**, SnapManager loads the control file, and validates the imported configuration and settings.

## Exporting Exchange Server configuration information using the control file

You can export specific settings of an Exchange server configuration to a control file that you can use to configure other Exchange servers. You can export the current configuration details to a control file, and export a specific section of the current configuration to a control file.

### Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. Select **Use Control File**.
3. To export configuration details for the server, click **Next**.
 

The Import or Export Selection window appears.
4. Select **Export**.
5. In **Use control file**, either type the complete path of the control file or use the **Browse** feature to locate and select the file path.
6. Click **Advanced**.
7. In the **Configuration Import/Export Advanced Options** window, specify the configuration settings that need to be exported.
8. Select **OK**.
9. Click **Next** and then, to load the control file and validate the imported configuration and settings, click **Apply**.
 

If you do not want to apply the new configuration, click **Cancel**.

If you click **Apply**, SnapManager loads the control file, and validates the imported configuration and settings.

## SnapManager control file XML schema

You can set and edit storage layout, notification, verification, report directory, backup, and SnapMirror settings through the control file using an XML schema.

### Storage layout settings XML schema

Use the storage layout settings XML schema to set and edit storage layout settings.

```
<?xml version="1.0" encoding="utf-8"?>
<xsd:schema xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
```

```

attributeFormDefault="unqualified"
elementFormDefault="qualified">
  <xs:element name="SMECONFIG" type="xs:string" />
  <xs:complexType>
    <xs:sequence>
      <xs:element name="HOST_NAME" type="xs:string" />
      <xs:element name="SERVER_NAME" type="xs:string" />
      <xs:element name="STORAGE_LAYOUT" type="xs:string" />
      <xs:complexType>
        <xs:sequence>
          <xs:element name="MTA_SYSTEM_FILES" type="xs:string" />
          <xs:complexType>
            <xs:sequence>
              <xs:element name="MTA_DB_DIR" type="xs:string" />
              <xs:element name="MTA_RUN_DIR" type="xs:string" />
            </xs:sequence>
          </xs:complexType>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
    <xs:element name="SMTP_SYSTEM_FILES_VS" type="xs:string" />
    <xs:complexType>
      <xs:sequence>
        <xs:element name="SMTP_SVC_PATH" type="xs:string" />
        <xs:complexType>
          <xs:sequence>
            <xs:element name="ADD_MSISCSI_DEPENDENCY" type="xs:boolean" />
            <xs:element name="STORAGE_GROUPS" type="xs:string" />
            <xs:complexType>
              <xs:sequence>
                <xs:element name="SG_NAME" type="xs:string" />
                <xs:element name="SG_SYS_PATH" type="xs:string" />
                <xs:element name="SG_LOG_PATH" type="xs:string" />
                <xs:element name="SG_SNAPINFO" type="xs:string" />
                <xs:element name="DATABASES" type="xs:string" />
              </xs:sequence>
            </xs:complexType>
            <xs:sequence>
              <xs:element name="DATABASE" type="xs:string" />
              <xs:complexType>
                <xs:sequence>
                  <xs:element name="DB_NAME" type="xs:string" />
                  <xs:element name="EDB_PATH" type="xs:string" />
                  <xs:element name="STM_PATH" type="xs:string" />
                </xs:sequence>
              </xs:complexType>
            </xs:sequence>
          </xs:complexType>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
</xs:complexType>

```

```

</xs:complexType>
</xs:element>
</xs:sequence>
</xs:complexType>
</xs:element>

```

When you have datasets configured in the host, and the storage group is configured with the dataset policy, the exported control file contains the following information in the storage layout settings:

```

<STORAGE_GROUPS>
  <STORAGE_GROUP>
    <SG_NAME>First Storage Group</SG_NAME>
    <SG_SYS_PATH>Q:\Program Files\Exchsrvr\mdbdata\</SG_SYS_PATH>
    <SG_LOG_PATH>Q:\Program Files\Exchsrvr\mdbdata\</SG_LOG_PATH>
    <SG_SNAPINFO_PATH>Q:\SIF\SME_SnapInfo\</SG_SNAPINFO_PATH>
    <SG_DATASET_POLICY>Backup up</SG_DATASET_POLICY>
    .....
  </STORAGE_GROUP>
</STORAGE_GROUPS>

```

## Notification settings XML schema

Use the notification settings XML schema to specify notification settings.

```

<xs:element name="COMMON_SETTINGS">
<xs:complexType>
  <xs:sequence>
    <xs:element name="NOTIFICATION">
      <xs:complexType>
        <xs:sequence>
          <xs:element name="SEND_EMAIL_NOTIFICATION">
            <xs:complexType>
              <xs:sequence>
                <xs:element name="SMTPSERVER" type="xs:string" />
                <xs:element name="FROM" type="xs:string" />
                <xs:element name="TO" type="xs:string" />
                <xs:element name="SUBJECT" type="xs:string" />
                <xs:element name="NOTIFY_AUTO" type="xs:boolean" />
                <xs:element name="LONG_MSG" type="xs:boolean" />
                <xs:element name="AS_ATTACHMENT" type="xs:boolean" />
                <xs:element name="SEND_ON_FAILURE" type="xs:boolean" />
              </xs:sequence>
            </xs:complexType>
          </xs:element>
          <xs:element name="EMS_ENABLED" type="xs:boolean" />
          <xs:element name="ASUP_ENABLED" type="xs:boolean" />
        </xs:sequence>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
</xs:complexType>

```

```

type="xs:boolean" />                                <xs:element name="ASUP_ON_FAIL"
type="xs:boolean" />                                </xs:sequence>
                                                    </xs:complexType>
</xs:element>

```

## Verification settings XML schema

Use the verification settings XML schema to specify the verification settings.

```

<xs:element name="VERIFICATION">
    <xs:complexType>
        <xs:sequence>
            <xs:element name="VERIFICATION_CLIENT_SETTING" >
                <xs:complexType>
                    <xs:sequence>
                        <xs:element name="VER_SERVER" type="xs:string" />
                    </xs:sequence>
                </xs:complexType>
            </xs:element>
        <xs:element
            name="VERIFICATION_SERVER_SETTING" >
            <xs:complexType>
                <xs:sequence>
                    <xs:element
                        name="ESEUTIL_PATH" type="xs:string" />
                    <xs:element
                        name="AUTO_DRIVELETTER" type="xs:boolean" />
                    <xs:element
                        name="MP_DIR" type="xs:string" />
                    <xs:element
                        name="THROTTLE" type="xs:boolean" />
                    <xs:element
                        name="IO_PAUSE" type="xs:unsignedByte" />
                    </xs:sequence>
                </xs:complexType>
            </xs:element>
        </xs:sequence>
    </xs:complexType>
</xs:element>

```

## Report directory settings XML schema

Use the report directory settings XML schema to specify report directory settings.

```

<xs:element name="REPORT_DIRECTORY" type="xs:string" />
    <xs:element name="BACKUP">
        <xs:complexType>
            <xs:sequence>
                <xs:element name="BACKUP_CLIENT_SETTING" >
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="NAMING_CONVENTION" >
                                <xs:complexType>
                                    <xs:attribute name="GENERIC"
type="xs:boolean" use="required" />
                                    <xs:attribute name="UNIQUE"
type="xs:boolean" use="required" />
                                </xs:complexType>
                            </xs:element>
                        </xs:sequence>
                    </xs:complexType>
                </xs:element>
            </xs:sequence>
        </xs:complexType>
    </xs:element>
    <xs:element
        name="BACKUP_SET_TO_KEEP" type="xs:unsignedByte" />
    <xs:element
        name="BACKUP_SET_TO_KEEP_IN_DAYS" type="xs:unsignedByte" />

```

```

    name="BACKUP_SET_TO_VERIFY" type="xs:unsignedByte" />
    </
  xs:sequence>
  </
  xs:complexType>
  </
  xs:element>

```

## Backup settings XML schema

Use the backup settings XML schema to specify backup settings.

```

<xs:element name="BACKUP_SERVER_SETTING">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="RUN_CMD_PATH" type="xs:string" />
      <xs:element name="RUN_CMD_ARGUMENT" type="xs:string" />
    </xs:sequence>
    </xs:complexType>
  </xs:element>
</xs:sequence>
</xs:complexType>
</xs:element>

```

## SnapMirror relationship settings XML schema

Use the SnapMirror relationship settings XML schema to specify SnapMirror relationship settings.

```

<xs:element name="VERIFICATION_ON_DESTINATION">
  <xs:complexType>
    <xs:element name="SELECTED_DESTINATIONS">
      <xs:complexType>
        <xs:sequence>
          <xs:element maxOccurs="unbounded"
name="SELECTED_DESTINATION">
            <xs:complexType>
              <xs:sequence>
                <xs:element name="SOURCE_FILER"
type="xs:string" />
                <xs:element name="SOURCE_VOLUME"
type="xs:string" />
                <xs:element
name="DESTINATION_FILER" type="xs:string" />
                <xs:element
name="DESTINATION_VOLUME" type="xs:string" />
              </xs:sequence>
            </xs:complexType>
          </xs:element>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
</xs:complexType>
</xs:element>
<xs:element
name="OTHER_AVAILABLE_DESTINATIONS">
  <xs:complexType>
    <xs:sequence>
      <xs:element
maxOccurs="unbounded"
name="AVAILABLE_DESTINATION">
        <xs:complexType>
          <xs:sequence>
            <xs:element name="SOURCE_FILER" type="xs:string" />

```



## Configuring SnapManager application settings

You can configure or change SnapManager application settings at any time after you install SnapManager as long as you run SnapManager from the system console and not from a Terminal Services client.

### Where to access SnapManager application settings

Using the SnapManager user interface, you can easily configure and change SnapManager application settings at any time after you install SnapManager. There are multiple ways with which you can access your application settings.

From the Configuration wizard, you can access settings for only the verification server. To access other verification settings (Override Verification, Throttling, and Access LUN in Snapshot copy), you must open the **Verification Settings** dialog box.

When you open the **Run Command After Operation** dialog box from the Actions pane, you can view or configure only the default settings. However, from within the context of a specific operation, the default settings are presented and then can be modified for this operation only. As an option, the default settings can be updated.

The following table shows the GUI components that you can use to configure SnapManager application settings.

Application settings	Where the setting can be accessed
Add servers to be managed	Actions pane Configuration wizard
Exchange user account	Actions pane
Migrate databases to local disk	Actions pane
Database verification settings <ul style="list-style-type: none"> <li>• Verification server</li> <li>• Override for restore operations</li> <li>• Verification throttling</li> <li>• Mount point</li> </ul>	Actions pane Backup wizard Restore wizard Configuration wizard
Backup settings	Actions pane Backup wizard Backup and Verification
Run Command After Operation default settings	Actions pane Within the context of a backup or database verification operation: Backup wizard, or Backup and Verify window

Application settings	Where the setting can be accessed
Fractional space reservation policy settings <ul style="list-style-type: none"> <li data-bbox="323 340 355 361">•</li> <li data-bbox="360 340 507 361">Current status</li> <li data-bbox="323 388 355 409">•</li> <li data-bbox="360 388 507 409">Policy settings</li> </ul>	Actions pane
Report directory settings	Actions pane
Event notification settings <ul style="list-style-type: none"> <li data-bbox="323 551 355 572">•</li> <li data-bbox="360 551 556 572">e-mail notification</li> <li data-bbox="323 599 355 620">•</li> <li data-bbox="360 599 442 620">Logging</li> </ul>	Configuration wizard Actions pane

## Adding Exchange servers to be managed

You can manage one Exchange server or multiple Exchange servers with SnapManager using the **Add Servers to be Managed** option. You cannot run SnapManager until you successfully add an Exchange server. You can also manage multiple servers.

### Steps

1. In the **Actions** pane, select **Add Servers to be Managed**.
2. Type the name of or browse to the Exchange server that you want to manage.

This setting remains in effect, specifying the default Exchange server, until, or unless you change it.

### Result

Whenever the SnapManager program starts, SnapManager automatically connects to the default Exchange server by using the default security authentication method.

**Note:** If you want to manage a different Exchange server later, use the **Add Servers to be Managed** option to connect and manage an Exchange server.

## Enabling database migration back to local disks

After you have migrated your database for SnapManager configuration, you can enable migration back to local disks.

### About this task

If you enable the **Enable databases and transaction logs to be migrated back to local disk** option, SnapManager disables it the next time you start SnapManager.

### Steps

1. In the **Actions** pane, click **Configuration Wizard Option Settings**.
2. Select the **Enable databases and transaction logs to be migrated back to local disk** check box to enable the migrate-back-to-local-disk feature.
3. Click **OK**.

## Disabling database migration back to local disks

You can disable the migration of your database back to local disks after you have configured SnapManager.

### Steps

1. In the **Actions** pane, click **Configuration Wizard Option Settings**.
2. Clear the **Enable databases and transaction logs to be migrated back to local disk** check box to disable the migrate-back-to-local-disk feature.
3. Click **OK**.

## Considerations for selecting the database verification server

You can manage verification of your backup sets on either your production Exchange server or on a remote verification server, and you can use any of deferred verification, remote verification, and verification throttling to do so.

Verification of databases in a backup set can be done on the production Exchange server (the Windows host system running the Exchange server used to service the users) or on a remote verification system (a different Exchange server with Exchange management tools installed). Running database verification on a production Exchange server is CPU-intensive for the Windows host and also involves a substantial amount of activity on the storage system.

**Note:** When you change the database verification server, this change does not affect any database backup (with verification) or database verification-only jobs that are already scheduled. You must recreate the scheduled backup jobs for the change to apply.

SnapManager offers three methods for managing database verification load that you can use separately or in any combination:

- Deferred verification
- Remote verification
- Verification throttling

Both the deferred verification and remote verification manage database verification load by separating database verification from the backup operation. Verification throttling manages database verification load by slowing down the verification throughput rate.

## Configuring the verification server

After you have configured SnapManager, you can configure your verification server to be the same as your host server or a remote server. Configuring a remote server as the verification server reduces the load on the host server. You can optionally configure the verification server after you have completed SnapManager configuration.

### Before you begin

You must have SnapDrive, SnapManager, and Exchange installed on your verification server. If you configure a remote server as your verification server, the versions of SnapManager and SnapDrive must be the same on both the host and the remote server.

You can also use the following path to configure the verification server: **Actions > Backup Verification Settings > Verification Server**.

### About this task

You can omit the configuration of the verification server during configuration by selecting the **Select the verification server later...** check box in the Configuration wizard.

If Exchange is not installed on the computer that you select, specify the `eseutil.exe` filepath and copy the necessary files to the computer before proceeding.

### Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. In the **Database Verification Server** window, enter or browse to the name of the server that you want to use as the verification server.

## Remote verification prerequisites

Before you use a remote verification server, ensure that you establish the Windows host requirements, LUN requirements, and verification server designation.

Remote verification uses the same mechanisms as local verification, except that the verification occurs on a different host than the one that initiated the backup operation. This is why you need SnapDrive and SnapManager installed on your remote verification server, in addition to FC or iSCSI connectivity to the storage system.

If you use a remote verification server to verify the databases in multiple databases in a single job, you need an additional LUN.

You can designate your verification server from the production Exchange server.

## How remote verification works

SnapManager initiates a backup operation at the primary host, which then contacts the remote verification server. The remote verification server then uses SnapDrive for verification and sends the results back to the primary host.

The basic steps of this process are as follows:

1. The SnapManager backup with verification (or verification only) is initiated on the primary SnapManager host, which is configured to run verifications on the remote verification server.
2. The primary SnapManager host contacts the remote verification server and initiates the verification job.
3. The remote verification server uses SnapDrive to connect to a LUN backed by a Snapshot copy containing the databases to be verified.
4. The remote verification server performs the database verification on the LUN backed by Snapshot copy.
5. When the remote verification server completes the verification, it sends the results back to the primary SnapManager host.

**Note:** You must be careful not to schedule backup operations while a verification operation is in progress. Verification is always performed on a LUN that is backed by a Snapshot copy. If you make a Snapshot copy of the same volume while a LUN backed by Snapshot copy exists, you create a “busy Snapshot copy,” which might cause problems when you attempt to delete some Snapshot copies.

## Viewing or changing the verification server

You can view or change your verification server for verification load management.

### Before you begin

You must be connected to the production Exchange server to view or to change the verification server.

### About this task

Until you specify verification settings, database verification is run from the Exchange server you selected. Verification does not necessarily run on the system from which you opened the **Database Verification Settings** dialog box. A change in the verification server does not affect any database verification jobs that are already scheduled.

### Steps

1. Click **Backup Verification Settings** in the **Actions** pane.

You can also use these paths:

- **Backup wizard > Verification Settings**
- **Restore wizard > Verification Settings**

The **Verification Server** tab is active by default and displays the host name of the current verification server.

2. In the **Verification Server** tab, type or browse to the Exchange server you want to use as the database verification server.

**Note:** If you plan to specify a remote verification server, ensure that the server is set up correctly.

3. Click **OK**.

## Selecting the Snapshot copy access method for database verification

Use the **Access LUN in Snapshot** tab to specify how SnapManager should access database backup Snapshot copies during database integrity verification. Assign either a drive letter or directory path to access the backup Snapshot copy as a mounted LUN.

### Steps

1. Click the **Access LUN in Snapshot** tab.

You can access the **Access LUN in Snapshot** tab from the Verification Settings window of the Configuration wizard, the Backup Verification Settings window, the Backup wizard or the Restore wizard.

2. Assign either a drive letter or directory path to access the backup Snapshot copy as a mounted LUN.

If you want to...	Then do this...
<b>Mount the Snapshot copy on the next available drive letter</b>	Select <b>Automatically assign available drive letter</b> .
<b>Mount the Snapshot copy on a specific NTFS mount point</b>	<p>Do the following:</p> <ol style="list-style-type: none"> <li>Select the <b>Mount in an empty NTFS directory</b> option.</li> <li>Enter or browse to the directory path of an NTFS mount point.</li> </ol> <p><b>Note:</b> This mount point is used if SnapManager is configured to use drive letters but runs out of available drive letters.</p>

- Click **OK**.

## Database verification throttling

SnapManager allows you to throttle the database checksum verification rate with Exchange Server. Database verification throttling enables you to manage your verification load.

**Note:** Database verification is not a support requirement for databases with at least two copies in a DAG. By default, when performing a DAG backup with SnapManager for Exchange, verification is off.

### Database verification throttling options

By default, SnapManager uses `ChkSgFiles`, the Microsoft Exchange integrity verification library, for throttling options. However, this can be changed through the registry to use `Eseutil`, the Microsoft Exchange consistency checker utility.

### How database verification throttling works

`ChkSgFiles` inserts a one-second pause after a given number of I/O operations during the database physical consistency verification.

Optionally, the `Eseutil.exe` can be used (through changing the registry) to insert a one-second pause after a given number of input output (I/O) operations during the database physical consistency verification.

`ChkSgFiles` (or `Eseutil.exe`) reads 512 KB for each database checksum verification I/O operation. Therefore, when configuring the throttling value, the maximum throughput rate for database checksum verification is decreased to the following:

$$512 \text{ KB per I/O} \times x \text{ I/Os per second} = 512 \times x \text{ KBps}$$

You can decrease the maximum throughput rate by decreasing the number of input output operations (x) that elapse between one-second pauses.

Number of I/O operations (x) between one-second pauses	Maximum possible database verification speed	
	Calculation	Maximum speed
100	$512 \text{ KB/IO} \times 100 \text{ IO/sec} = 51,200 \text{ KBps}$	50 MBps
150	$512 \text{ KB/IO} \times 150 \text{ IO/sec} = 76,800 \text{ KBps}$	75 MBps

200	512 KB/IO × 200 IO/sec = 102,400 KBps	100 MBps
250	512 KB/IO × 250 IO/sec = 128,000 KBps	125 MBps

**Note:** Decreasing the ChkSgFiles (or Eseutil.exe) database verification throughput causes the database checksum verification to take longer to complete. This also means that a backup job that is configured with verification also takes longer to complete. However, decreasing verification throughput does not cause a backup job without verification configured to take any longer to complete.

## Calculating the verification throttling sleep interval

You must calculate the appropriate sleep interval value to use for each server. Monitor the read and write performance of the LUNs on which the databases reside, and then calculate the verification throttle setting so that the average and peak physical disk performance is below the maximum values determined by Microsoft.

### About this task

- To monitor the SnapManager progress, watch the Backup Status window that is displayed during a manually launched backup job.
- Monitor the counters during the busiest time of day that a verification or backup with verification job might occur.

### Steps

1. Launch the Windows Performance Monitor utility (also known as “Perfmon”) on the machine that is the Exchange server.
2. Add the Perfmon counters PhysicalDisk\Average Disk sec/Read, and PhysicalDisk\Average Disk sec/Write for each LUN that contains a database.
3. Monitor the pairs of values, Physical Disk\Average Disk sec/Read and PhysicalDisk\Average Disk sec/Write, while a SnapManager verification job or a backup job with verification is in progress.

If monitoring a backup with verification job, be sure to monitor the counters during the verification portion of the job. Multiple databases on the same Exchange server are verified sequentially, so be sure to monitor the counters for all of the database verifications.

4. The throttle setting is global, so select the reading from the database that exhibits the worst performance (highest readings).
5. Compare the measured values to the maximum read and write values recommended by Microsoft.

Physical Disk\Average Disk sec/Read: Average read time should be below 20 ms, Maximum peak read time should be below 50 ms.

Physical Disk\Average Disk sec/Write: The maximum physical disk performance times determined by Microsoft are as follows: Average write time should be below 20 ms, Maximum peak write time should be below 50 ms.

Compare the measured values to the maximum read and write values recommended by Microsoft: less than 20 ms. average time and less than 50 ms. peak time.

6. If the measured values exceed the maximum values recommended by Microsoft, adjust the throttle setting to decrease the verification throughput rate.

## Configuring database verification throttling

You have to configure database verification throttling on all of the Windows host systems that are to perform database verification and running Exchange Server.

### Before you begin

Throttling is accessible only if the verification server is installed with Microsoft Exchange Server.

### Steps

1. From the SnapManager console, click **Backup Verification Settings** in the **Actions** pane.  
You can also use these paths:
  - **Backup Wizard > Backup or Verify Databases and Transaction Logs > Verify databases and transaction Logs > Backup management group > Database Verification Server > Verification Settings**
  - **Restore wizard > Mount options > Verification Settings**
2. Click the **Verification Throttling** tab.
3. Select or clear the **Throttle database checksum verification** check box to enable or disable the throttling feature.
4. If you enabled the throttling feature, enter a positive integer value in the **Pause for 1 second after x I/O operations** box.

The default value is 150. Enter the number of Input Output (I/O) operations to complete before pausing for one second during checking of the database physical consistency.

Using a value in the range of 100 to 250 suits most environments.

5. Click **OK**.

## Throttling entries in the SnapManager backup and verification report

If you enable either Eseutil throttling or ChkSgFiles throttling, the SnapManager Backup and verification report logs messages indicating that the throttling feature is enabled before each database verification operation.

If you use Eseutil throttling, before each database verification, the SnapManager backup and verification report includes the following entry:

`ESEUTIL throttling feature is enabled.`

If you use ChkSgFiles throttling, before each database verification, the SnapManager and verification report includes the following entry:

```
Running Integrity Verification using ChkSgFiles API Throttle (Pause): 1000
ms per X I/O's
```

## Impact of database verification on performance

Database verification can impact the performance of both the Exchange server and the storage system. To overcome the performance impact, you can separate the database verification process and database backup operations.

Verification can degrade Exchange Server response, particularly during peak work hours. There are two options for distributing this load by separating database verification and database backup operations: deferred database verification and remote database verification.

To reduce the load even further, you can use database verification throttling, which can be combined with either or both of the existing options. This feature supports the ability to throttle Eseutil (the Microsoft Exchange consistency checker utility) in database checksum verification mode.

Another option for managing database verification uses the same database checksum algorithm as Eseutil throttling. This feature supports the ability to throttle ChkSgFiles (the Microsoft Exchange integrity verification library) in database checksum verification mode.

## Database verification override during restore operation

You must not select an unverified backup copy as the source of the restore operation. If a verified backup copy is not available when you need to restore and you cannot wait for a verification to complete before restoring, you can override verification and restore from an unverified backup copy.

### Configuring the database verification override option

You can configure SnapManager to override the verification requirement and restore directly from an unverified backup copy.

#### Steps

1. In the **Actions** pane, click **Backup Verification Settings** > **Verification Settings**.

Other ways to open Verification Settings are as follows:

- From the Backup Wizard, go to the **Verify the Databases and Transaction Logs in this Backup** window and click **Verification Settings**.
- From the Restore Wizard, go to the **Verify the Database Integrity in this Backup** window (displayed only if you select to restore from an unverified backup copy) and click **Verification Settings**.

2. Select the **Override Verification** tab.

3. Configure the **Override Database Verification Requirement for Restore** option.

**Note:** If you enable this option, it is reset the next time you start the SnapManager application.

4. A message appears that says "An unverified backup may contain an image of an Exchange database that is physically corrupt. If the backup copy contains physical database corruption, this corruption will persist in the restored database." If you want to enable the **Override Database Verification Requirement for Restore** option, click **OK**; otherwise, click **No**.

5. Click **OK** to close the dialog box and apply your change.

**Note:** If you are using the Restore wizard, by selecting the **Override Database Verification Requirement for Restore** option causes the Verify the database integrity of this backup prior to restoring it option to become accessible so that you can disable it.

If...	Then...
<b>The override option is disabled (the default and recommended setting)</b>	You cannot select an unverified backup copy as the source of the restore operation.
<b>The override option is enabled (not recommended)</b>	If you select an unverified backup copy as the source of the restore operation, you are asked to proceed with the restore operation.
<p><b>Note:</b> This option is reset when you exit SnapManager.</p>	

## Verification override entry in the SnapManager restore report

You must restore only from a verified database to ensure a successful restore operation. Before each database verification, the SnapManager restore report indicates if verification override is enabled. If the feature is enabled, the report includes the warning “Database verification before restore was overridden.”

## Configuring default settings for the Run Command After Operation option

You can configure default values to populate the **Run Command After Operation** dialog box when you open it from the Actions pane, or the Backup and Verify window, or the Backup wizard, to run a command or a script after a backup or verification operation.

### Steps

1. In the **Actions** pane, click **Run Command After Operation**.
2. In the **Specify a computer where...** box, enter or browse to the name of the host on which your program or script resides.
3. In the **Specify the full path...** box, browse to your program or script.
4. Enter the command input string in the **Command Arguments** box.

You can do this using any combination of the following methods:

- To enter text directly into the **Command Arguments** box, click the box and type the desired text.
- To enter a SnapManager variable into the **Command Arguments** box, do the following:
  - a. Click the **Command Arguments** box to position the cursor.
  - b. In the **SnapManager Variables** list, select the variable you want to enter.
  - c. Click **Select**.

**Note:** The `$SnapInfoPath` variable is enclosed within double quotes so that the path name can contain spaces without affecting the script invocation on the Windows command line. If you do not want the double quotes to appear in your command line, remove them from the Command Arguments box.

5. Repeat steps 1 to 4 as required, until the **Command Arguments** box contains the arguments you want to pass to your program or script.
6. Click **OK**.

**Note:** SnapManager verifies that the specified program exists on your system. SnapManager does not run the command until the backup or verification operation is complete.

Whenever the **Run Command After Operation** dialog box is opened from either the Backup and Verify window or the Backup wizard, the boxes are populated with the values you specified as default settings.

## Fractional space reservation

When you create a LUN, Data ONTAP reserves space in the volume containing that LUN so that write operations to that LUN do not fail due to lack of disk space. With fractional reserve, this space is set to less than 100 percent of the total size of the LUNs.

SnapDrive creates and manages LUNs with space reservation enabled. Operations such as creating a Snapshot copy or creating new LUNs can occur only if there is enough available unreserved space. These operations are restricted from using reserved space.

While space reservation is enabled at the LUN level, fractional overwrite reserve amounts are configured at the volume level; that is, fractional space reservation does not control how the total amount of space reserved for overwrites in a volume is applied to individual LUNs in that volume.

The volume has the guarantee option set to `volume` rather than `file`. Fractional reserve is supported by Data ONTAP 7.1 or later. For more detailed information, see the *SAN Administration Guide* for your version of Data ONTAP.

Additional space that is not space-reservation-enabled on the volume is automatically reserved for overwriting blocks that belong to a LUN. By default this additional space is equal to 100 percent of the total size of all space-reserved LUNs in the volume. If space reservation is disabled, write operations to a LUN might fail due to insufficient disk space in the volume and the host application might terminate, report I/O errors, or experience unexpected behavior.

With fractional reserve, the space reserved for overwrites is set to less than 100 percent and the space that is preallocated for space reservation is reduced to that percentage. Fractional reserve is generally used for volumes with LUNs that store data with a low rate of change.

## What can happen with a fractional-space-reserved volume

If a fractional-space-reserved volume runs out of overwrite reserve space, write operations to a LUN fail and the host application might terminate, report I/O errors, or exhibit unexpected behavior. Data ONTAP uses automatic expansion of flexible volumes and automatic deletion of Snapshot copies from flexible volumes to avoid this situation.

When a LUN is fully space reserved, write operations cannot fail due to an out-of-space condition. When the overwrite reserve for a volume is set to less than 100 percent, write operations to the LUNs on that volume might fail when the volume runs low in free disk space.

The automatic expansion of flexible volumes and the automatic deletion of Snapshot copies from flexible volumes monitor the reserved space and take action if the free space becomes scarce. For more detailed information, see the *SAN Administration Guide* for your version of Data ONTAP.

### Automatic expansion of flexible volumes

Data ONTAP automatically expands a nearly full volume into the space preallocated for it in the aggregate. The volume must be a flexible volume with the guarantee option set to **Volume**. You can enable automatic deletion of Snapshot copies and FlexVol expansion features separately or together, with one policy to be applied before the other. When fractional-space-reserved volumes hold LUNs that store Exchange database files, however, you can only use only the automatic FlexVol expansion feature.

### Automatic deletion of Snapshot copies from flexible volumes

Data ONTAP automatically deletes one or more Snapshot copies on a nearly full volume, when you enable the Snapshot copy automatic deletion policy. If the trigger condition is detected, the oldest or newest Snapshot copies are deleted until a configured percentage of the volume is free space. If you do not want to automatically delete Snapshot copies on the volume, you can set the overwrite reserve to 100 percent, by setting the fractional space reserve to 100 percent on the storage system.

This Data ONTAP feature is not designed specifically to support backup and restore operations on Exchange databases. The options for selecting Snapshot copies to be deleted do not have visibility to the automatic backup Snapshot copy deletion criteria configured in SnapManager. You must always retain at least one online backup copy for each database.

## Fractional space reservation policies

Fractional space reservation policies include specific thresholds that determine when SnapManager must delete Exchange backup sets or unmount Exchange databases (or both), because the overwrite reserve utilization for the volume is running low.

If overwrite reserve space runs low for a fractional space-reserved volume, SnapManager prevents the overwrite reserve from becoming fully depleted.

### The default fractional space reservation policy

SnapManager sets a default policy for fractional space reservation. You can use the default values, or change the values that would apply to all storage system volumes.

The default fractional space reservation policy is automatically enabled for any traditional or flexible volume that has an overwrite reserve that is set to less than 100 percent. The volume must also contain LUNs that store Exchange database files, Exchange transaction log files, or SnapManager SnapInfo directories.

#### Default policy

You can use the default policy as-is, allowing the factory default values to be applied to all volumes that contain fractional space-reserved LUNs.

#### Default policy with customized settings

Optionally, you can customize the default policy that is applied to all storage system volume that contains fractional space-reserved LUNs.

#### Volume-specific policies

Optionally, you can override the default policy for any particular volume that contains fractional-space-reserved LUNs, by applying a custom policy.

## Fractional space reservation policy settings

The fractional space reservation policy settings enable you to indicate when SnapManager should begin automatically deleting Snapshot copies and unmounting Exchange databases due to overwrite reserve utilization. You can also specify how many Snapshot copies to retain.

Enabling automatic deletion of Exchange backup Snapshot copies does not necessarily prevent an out-of-space condition on the volume. Therefore, database unmounting is always enabled. If Snapshot copy deletion is enabled, you must configure it to trigger before unmounting the database.

### Deletion of Exchange backup Snapshot copies

SnapManager fractional space reservation policy setting	Factory default value	Configurable values
	Status: enabled	Status: enabled or disabled
Trigger on overwrite reserve utilization	70%	1% through 99%
Number of Snapshot copies to retain	5	1 through 256

## Unmounting of Exchange databases

SnapManager fractional space reservation policy setting	Factory default value	Configurable values
	Status: enabled	Status: enabled
Trigger on overwrite reserve utilization	90%	1% through 99%

## Configuring fractional space reservation policies

You can enable fractional space reservation and also set the value of the fractional space reservation. With fractional reserve, the space reserved for overwrites is set to less than 100 percent of the total size of the space-reserved LUNs in a traditional volume or a flexible volume.

### About this task

- Although automatic deletion of Exchange backup Snapshot copies does not necessarily prevent an out-of-space condition on the volume, it is recommended that the automatic deletion of backups be enabled for every volume that contains fractional-space-reserved LUNs that store Exchange data.
- Data ONTAP includes a separate Snapshot copy autodelete feature. The SnapManager autodelete feature can be used in place of or along with the Data ONTAP autodelete feature.

**Note:** You cannot configure fractional space reservation policies through the Database Availability Group (DAG). You can configure them only when you connect to a member server of the DAG.

### Steps

- In the **Actions** pane, select **Fractional Space Reservation Settings**.
- Click the **Policy Settings** tab.
- Specify which policy you want to view or change.

If you want to access this...	Then do this...
The default policy	In the navigation tree, select <b>Default Policy</b> .
A volume-specific policy	In the navigation tree, select the storage system and then the volume.

- To enable fractional space reservation monitoring, select the **Enable Fractional Space Reservation Monitoring** check box.
- Use **Automatically delete backup sets** to enable or disable automatic deletion of Exchange backup Snapshot copies in fractional-space-reserved LUNs on the volume.

If you want to...	Then...
Enable automatic deletion of Exchange backup Snapshot copies	Select <b>Delete backups that include LUNs which have less than 100% overwrite reservation</b> , and then skip ahead to Step 8.
Disable automatic deletion of Exchange backup Snapshot copies	Clear <b>Delete backups that include LUNs which have less than 100% overwrite reservation</b> , and then proceed to Step 6.

- In the **Trigger point for overwrite reserve utilization** field, type the level of overwrite reserve utilization (in percentage of total reserve) that is to trigger deletion of Exchange backup Snapshot copies.

The value must not be a negative integer that is less than the **Trigger point for overwrite reserve utilization** value in the **Automatic dismount of databases** panel.

7. In the **Number of most recent backup sets to retain** field, type the number of backup sets to be retained if automatic backup set deletion is triggered.

The value must be an integer from 1 through 256 and should be based on the backup creation and verification schedule.

8. Use the **Automatically dismount databases** panel to configure automatic unmounting of Exchange databases in fraction-space-reserved LUNs on the volume.

**Note:** Because automatic deletion of Exchange backup Snapshot copies does not necessarily prevent an out-of-space condition on the volume, SnapManager does not allow you to disable unmounting of databases for any fractional space reservation policy.

9. In the **Trigger point for overwrite reserve utilization** field, type the level of overwrite reserve utilization (in percentage of total reserve) that is to trigger unmounting of Exchange databases.

The value must be an integer from 0 through 99.

**Note:** If Snapshot copy autodelete is enabled, SnapManager requires that this threshold be set to a higher level than the threshold that triggers automatic Snapshot copy deletion. Setting the threshold at a higher value ensures that Snapshot copy autodelete is triggered first.

10. Click **OK**.

## Fractional space reservation policies to manage Exchange data

Fractional space reservation policies enable you to monitor overwrite reserve utilization on fractional space-reserved volumes that contain your Exchange data.

If you store Exchange data on LUNs in a fractional space-reserved volume in a SnapManager environment, you need to avoid an out-of-space condition on the volume such that you have explicit or implicit Exchange-aware control over the deletion of Exchange backup set components.

To address this need, SnapManager provides its own space management tool for monitoring overwrite reserve utilization on the volumes. If overwrite reserve space runs low for a fractional space-reserved volume, SnapManager can take action to prevent the overwrite reserve from becoming fully depleted.

Fractional space reservation policies include specific thresholds that act as trigger points. SnapManager can delete Exchange backup sets or unmount Exchange databases (or both) when the overwrite reserve utilization for the volume reaches the trigger point.

**Note:** If you enable SnapManager e-mail notification, SnapManager sends SMTP e-mail message after an event of SnapManager fractional space reservation policy is complete.

### Automatic unmounting of Exchange databases

Automatic unmounting is triggered if overwrite reserve utilization on the volume reaches the threshold specified by the fractional space reservation policy. SnapManager automatically unmounts the databases and stops the write operations to LUNs in that volume.

The threshold for overwrite reserve is specified by the policy for fractional space reservation. Another component of the fractional space reservation policy is the last-resort action that prevents further consumption of overwrite reserve; hence, automatic unmounting is always enabled.

SnapManager first uses backup set deletion to free some overwrite reserve. If this is not enough, unmounting the affected database prevents further consumption of overwrite reserve. This happens when both components of a fractional space reservation policy are enabled; and the unmounting of databases is triggered at a later level of overwrite reserve utilization than the level that is used to trigger the deletion of Exchange backup Snapshot copies.

**Attention:** If another host or client continues to write data to the affected volume, the overwrite reserve space might still run out and the volume goes offline. For this reason, you must use dedicated volumes for Exchange data.

### Automatic deletion of Exchange backup copies

You can enable automatic deletion of LUN backup copies that store Exchange data. SnapManager checks for the level of overwrite utilization on the volume and triggers automatic deletion of backups if the level of overwrite utilization reaches the threshold.

Automatic deletion of LUN backup copies serves as the Exchange-aware replacement for or adjunct to the feature of Data ONTAP Snapshot copy deletion. SnapManager follows the following sequence of steps:

1. Deletes the oldest Snapshot copies.
2. Retains the specified number of total Snapshot copies on the volume.
3. Retains the most recent backup of any database (if it resides on the volume).
4. Retains any backup copies of databases no longer in existence.

You must select the backup retention level based on your backup copy creation and verification schedule. It is important that at least one verified backup copy remains on the volume if Snapshot copy deletion is triggered. Due to its Exchange-aware features, the automatic deletion of Snapshot copies does not necessarily prevent an out-of-space condition on the volume.

You must set the same number of backup sets to be deleted on database LUNs and transaction log LUNs. If there is a mismatch in this number, SnapManager attempts to delete backup sets based on the fractional reserve policy settings.

#### Example: automatic deletion of backup sets that span multiple volumes

If you have a backup copy that spans multiple volumes, with a different automatic deletion threshold configured on each volume, then for a specific volume, SnapManager deletes Snapshot copies based on the policy for that volume.

In this example, the automatic deletion settings for each volume are configured to take the following actions:

- Volume 1: Delete all but 2 Snapshot copies if 20 percent overwrite reserve utilization is exceeded.
- Volume 2: Delete all but 5 Snapshot copies if 20 percent overwrite reserve utilization is exceeded.
- Volume 3: Delete all but 10 Snapshot copies if 20 percent overwrite reserve utilization is exceeded.

If the 20 percent overwrite reserve utilization threshold for Volume 1 is exceeded, SnapManager deletes all but two Snapshot copies, regardless of the policies for Volumes 2 and 3. If the 20 percent overwrite reserve utilization threshold for Volume 2 is exceeded, SnapManager deletes all but five Snapshot copies, regardless of the policies for Volumes 1 and 3.

## Viewing current fractional space reservation data for a LUN

You can view current fractional space reservation settings to ensure that the policy is in force for each LUN is configured appropriately.

### About this task

Only the **Drive Letter** or **Mount Point** column displays LUN-specific information. All other columns in the **Current Settings** tab display information that applies across the volume that contains the LUN.

The SnapManager fractional space reservation policy includes a separate Exchange-aware automatic deletion feature. The SnapManager automatic deletion feature can be used in place of or along with the Data ONTAP automatic deletion feature; you can also select to disable the SnapManager automatic deletion feature.

### Steps

1. In the **Actions** pane, click **Fractional Space Reservation Settings**.
2. In the **Current Settings** tab, note the space consumption status for each LUN that stores database or SnapInfo directories.

The information displayed in this tab automatically refreshes every 60 seconds.

3. If the **Snapshot Autodelete** column is enabled, investigate the cause of it being enabled and take one of the following preventive actions:
  - Disable the Data ONTAP Automatic deletion of Snapshot copies feature.
  - Ensure that the Data ONTAP Automatic deletion of Snapshot copies feature is configured in such a way that it does not delete Exchange backup set components.

For details about the `snap automatic delete storage system` command, see the `man` page.

4. Click **OK** to close the dialog box.

## Fractional space reservation status data

You can view data about your current space-reservation status in the **Current Settings** tab of the **Fractional Space Reservation Settings** dialog box

If the Storage Snapshot Autodelete option is enabled, the LUN is contained in a FlexVol volume that has overwrite reserve set to less than 100 percent and that also has the Data ONTAP automatic Snapshot copy deletion feature enabled and configured to trigger when the overwrite reserve is nearly full. If Exchange data or SnapManager SnapInfo directories are stored on LUNs contained in a volume with these characteristics, the Data ONTAP Snapshot copy automatic deletion policy might delete Exchange backup set components.

The following columns display SnapManager configuration information:

### Drive Letter or Mountpoint

A SnapManager configuration setting. The drive letter or NTFS mountpoint on which the LUN is mounted.

### Backup Autodelete Trigger (percentage)

A SnapManager fractional space reservation policy setting. The percentage of overwrite reserve utilization that triggers automatic deletion of Exchange backup sets for the volume that contains the LUN.

### Disable Database Trigger (percentage)

A SnapManager fractional space reservation policy setting. The percentage of overwrite reserve utilization that triggers automatic disabling of Exchange databases for the volume that contains the LUN.

The following columns display the fractional overwrite reserve settings and status:

### Fractional Reserve (percentage)

The amount of space reserved for overwrites on the volume that contains this LUN, expressed as a percentage of the total size of all space-reserved LUNs in the volume

### Used Reserve (percentage)

For the volume that contains this LUN, the amount of overwrite reserve in use, expressed in two ways: as a percentage of the total size of all space-reserved LUNs in the volume, and in megabytes

### Available Reserve (MB)

For the volume that contains this LUN, the amount of overwrite reserve available

### Snapshot Autodelete

For the volume that contains this LUN, the state of the Data ONTAP Snapshot copy automatic deletion feature: enabled or disabled

If this LUN stores Exchange data files and is contained in a volume for which the Data ONTAP Snapshot copy automatic deletion feature is enabled, disable this feature on that volume or ensure that it is configured so that it does not delete SnapManager backup set components.

**Note:** The SnapManager fractional space reservation policy triggers are not applicable to fully space-reserved LUNs. If Fractional Overwrite Reserve (percentage) is 100, the LUN is contained in a fully space-reserved volume rather than a fractionally space-reserved volume.

## Event notification options

You can use either the Configuration wizard or the **Auto Notification Settings** dialog box to configure email notifications, syslog event logging, and AutoSupport notifications.

### Operations with email notification support

SnapManager can notify you through email messages (using SMTP) about the success or failure of the following types of events:

- SnapManager backup operation
- Database integrity verification
- SnapManager restore operation
- SnapManager configuration
- SnapManager fractional space reservation policy event execution

You can select *one* of these body messages to include in the body of the email: Send operation results summary or Send verbose operation results.

You must enable the email notifications option, which is disabled by default.

### **SnapManager syslog event logging**

When there is a failure, SnapManager events are posted to the storage system syslog by default. You can disable this option to reduce the load on the network or when you are troubleshooting your system.

### **AutoSupport notification**

If AutoSupport is enabled on both the storage system and SnapManager, technical support receives automatic email notification about any SnapManager events or storage system problems that might occur. This option is enabled by default. You can disable this option to reduce the load on the network or when you are troubleshooting your system.

The AutoSupport daemon monitors the storage system's operations and sends automatic messages to technical support to alert them to potential storage system problems. If necessary, technical support contacts you by email to help resolve a potential system problem. The AutoSupport daemon is enabled by default on the storage system. For more information, see the *System Administration Guide* for your version of Data ONTAP.

### **Limitation of AutoSupport notification to failure events only**

If AutoSupport is enabled, you can limit the SnapManager events that are posted to the storage system syslog and AutoSupport (if allowed for SnapManager) to failure events only. The option is enabled by default.

## **Configuring automatic event notification settings**

You can configure automatic event notification settings for SnapManager. You can enable and configure e-mail notification, advanced event notification settings, advanced e-mail notification settings, storage system syslog settings, and AutoSupport notification settings.

### **Before you begin**

You must have the following ready before you configure automatic event notification settings:

- IP address of the SMTP e-mail server or gateway
- E-mail address of each recipient to whom the notification is to be sent.
- E-mail address of the sender of the notification that you want to use

By default, SMEAutoSender is the name of the notification sender. To specify a sender other than the default, use one of the following formats:

- *SenderAlias<SenderName@SenderDomain>*
- *SenderAlias*
- *SenderName@SenderDomain*

- The text to be appended to the standard subject line, which is included in all notification messages:  
*Backup status at mm\_dd\_yyyy-hh.mm.ss from MachineName*  
By default, the string *SnapManager for Exchange* is appended.

If you select to send the operational results in summary format rather than in verbose format, you can also select the **Include SnapManager Operation Report as an Attachment** option.

### **About this task**

By default, the automatic e-mail notification feature is disabled.

SnapManager relies on and requires an external e-mail host at your site to send e-mail. The e-mail host is a host that runs a mail server that listens on the SMTP port (25).

You can configure the SMTP mail notification settings by selecting either or both of the following two options:

**Only send notification when operation fails**

Specifies that you want e-mail notification sent only when a backup process or a verification process fails (cleared by default).

**Include SnapManager operation report as attachment**

Specifies that you want the status report to be attached to the e-mail notification (cleared by default).

**Steps**

1. In the **Actions** pane (or the **Configure Automatic Event Notification** screen of the Configuration wizard), select **Notification Settings**.

The **Auto Notification Settings** dialog box opens.

2. To enable e-mail notification, select the **Send Email Notification** option.
3. In the relevant text boxes, type the following information.
  - a. In the **SMTP Server** text box, type the host name or the IP address of the SMTP e-mail server or gateway to be used.
  - b. In the **From** text box, type the e-mail address of the sender of the notification.
  - c. In the **To** text box, type the e-mail address of each recipient.

To send to more than one recipient, use a semicolon (;) to separate the addresses.

4. Click **Advanced**.
5. The **Advanced Event Notification Settings** dialog box opens.
6. In the **E-mail Message Content** pane, select the types of body messages to include in the e-mail.
7. If you choose the summary format rather than the verbose format, you can also select the **Include SnapManager Operation Report** as an Attachment option.

8. Click **Apply** to commit your settings.
9. To configure the SMTP mail notification settings, select either or both of the following two options:

- **Only send notification when operation fails** check box
- **Include SnapManager operation report as attachment** check box

10. Click **OK** to apply your settings and close the **Advanced E-mail Notifications Settings** dialog box.

**11. Click **Send a Test Email**.**

SnapManager sends a test e-mail notification that uses the settings you specified and displays what that e-mail message looks like.

12. If you want to post SnapManager events to the storage system syslog, select **Log SnapManager Events to Storage System Syslog**.

12. If you want to enable automatic notification of syslog entries to technical support, and if SnapManager is configured to log events to the storage system syslog, select **Send AutoSupport Notification**.
13. If you want to limit SnapManager event logging to failure events, select **On failure only**.
14. Click **OK** (or **Next**, if you are using the Configuration wizard).

## Advanced administration

---

For very large implementations, you need to be aware of SnapManager's configuration limits. You also should be aware of requirements for advanced configurations.

### Maximum configurations supported by SnapManager

Microsoft Exchange Server 2013 has data configuration requirements with respect to the number of databases per server and the number of Recovery databases.

#### Recovery Database

A Recovery Database enables you to restore, recover, and mount your database to a different Exchange server. Each Exchange server can have only one Recovery Database mounted at any time.

#### Member servers in a DAG

A Database Availability Group (DAG) can have a maximum of 16 Exchange member servers.

#### Maximum supported databases

The maximum number of databases (including the Recovery Database) that can be associated with each Exchange server are as follows:

- Standard Edition: 5
- Enterprise Edition:
  - Exchange 2013 CU1: 50 databases per node
  - Exchange 2013 CU2: 100 databases per node

### Service account requirements for archiving backup sets with SnapVault (7-Mode environments only)

To archive backup sets with SnapVault on Data ONTAP systems operating in 7-Mode, the SnapManager service account should be the same account you used to configure SnapDrive access to the DataFabric Manager server. If you cannot use the same account, the SnapManager service account needs specific permissions on the server.

On the DataFabric Manager server, you can assign permissions to the SnapManager service account using one of the following methods:

If you want to...	Then...
Assign specific permissions	<p>Assign the SnapManager service account a role on the DataFabric Manager server with the following capabilities:</p> <ul style="list-style-type: none"> <li>• DFM.DataBase.Read Global</li> <li>• DFM.DataSet.Write Global</li> <li>• DFM.Policy.Read Global</li> <li>• DFM.BackupManager.Backup Global</li> <li>• DFM.BackupManager.Read Global</li> <li>• DFM.BackupManager.Restore Global</li> </ul> <p>You can use the <code>dfm role create</code>, <code>dfm role add</code>, and <code>dfm user add</code> commands to create the role, add the capabilities, and create the user.</p>
Assign full-control permissions	<p>Assign the SnapManager service account full-control rights on the DataFabric Manager server as shown in the following examples:</p> <ul style="list-style-type: none"> <li>• On Windows:</li> </ul> <pre>dfm user add -r GlobalFullControl MyDomain\snapuser</pre> <ul style="list-style-type: none"> <li>• On UNIX:</li> </ul> <pre>dfm user add -r GlobalFullControl MyDomain\\snapuser</pre>

## Concurrent backup verification

You can run multiple backup verification jobs concurrently on the verification server when the verification requests originate on different servers from the server where the backup is created.

**Note:** SnapManager executes multiple verification requests from the same host serially.

### SnapManager workflow for concurrent backup verification

The workflow for concurrent backup verification is as follows:

1. After you create a new backup copy at the source backup server, SnapManager immediately sends it for verification to the remote verification server.  
The remote server uses the same Job ID as the source server.
2. The backup job sent to the remote server is displayed as running on the source server.  
The remote server shows the active job as running or queued, depending on its position in the queue.
3. After SnapManager verifies the first backup set at the remote server, the job runs until all the backup sets are created and verified.

**Note:** There is a maximum of four backup jobs from different servers that SnapManager can verify simultaneously. SnapManager places the subsequent jobs in the queue.

The job can be a full backup job that includes backup verification, a deferred backup verification job, or a backup verification job that is initiated as part of a restore job.

You can create a new full backup when the deferred integrity verification job is running.

## Concurrent backup verification during a restore operation

You can restore a backup copy when the verification job is running. When you submit a restore job with verification, SnapManager performs one of the following tasks:

- Aborts the restore job if any backup or deferred integrity verification jobs are running
- Cancels all the running jobs and then starts the restore operation

**Note:** The status of the database of all the canceled verification jobs remains unverified.

- Waits for all the currently running jobs to finish and then starts the restore operation

When you select to have the restore operation cancel all the currently running jobs, the restore operation performs the following tasks:

1. Stops processing new jobs that are in the queue
2. Disables all SnapManager scheduled tasks in the Windows scheduler
3. If SnapManager is creating a full backup, cancels the full backup operation
4. If SnapManager is running a Frequent Recovery Point operation, cancels the Frequent Recovery Point operation
5. If one or more verification jobs are running, cancels all running jobs
6. Waits for the full backup operation to stop
7. Runs the restore operation

## Concurrent backup verification and Frequent Recovery Point backup

Concurrent backup verification involves Frequent Recovery Point backup running in parallel with full backup.

- If SnapManager is waiting for VSS Snapshot copy to start, then Frequent Recovery Point backup is running.
- If SnapManager is running backup verification for the new VSS Snapshot copy backup, then Frequent Recovery Point backup starts.
- If SnapManager is creating a new VSS Snapshot copy backup, then Frequent Recovery Point backup fails to start.

You can select restore options for such backups from among the Job Control Options in the Restore window.

## Managing integrity verification jobs

You can view, move, and cancel queued and running integrity verification jobs. You can perform all the tasks from the Current Job Status pane. By default, job management is enabled in SnapManager.

### Steps

1. In the Scope pane, select an Exchange server.

SnapManager displays the queued and running jobs for that server in the Current Job Status pane.

2. Select the job you want to manage.
3. Position your mouse over the job.

A floating menu appears.

4. Manage the jobs as described in the following table.

If you want to...	Then do this...
<b>Retrieve job information</b>	Select the job. SnapManager displays the job information in the Results pane.
<b>Move a job up within a queue</b>	Select <b>Move Up</b> in the floating menu.
<b>Move a job down within a queue</b>	Select <b>Move Down</b> in the floating menu.
<b>Cancel a queued or running job</b>	Select <b>Cancel Job</b> in the floating menu.

## Upgrading SnapManager

---

When a new version of SnapManager becomes available, you can upgrade your existing installation using an interactive wizard or a command.

### Before you begin

- Your host and storage systems must meet the minimum requirements for the version of SnapManager to which you are upgrading.
- You should have backed up your Exchange databases using SnapManager.
- SnapManager operations must have been stopped.

### About this task

You do not need to stop Exchange services while you upgrade SnapManager. Exchange can continue to run while you upgrade.

After installing or upgrading to this release, you must specify SnapMirror and SnapVault destination volumes in the Database Verification Settings dialog box before you can schedule backup set verification on those volumes. In the Database Verification Settings dialog box, click **Verification Server > Verification on Destination Volumes**, then choose the destination volumes.

## Upgrading SnapManager interactively

You can use the SnapManager installation wizard to interactively upgrade SnapManager.

### About this task

When you run the SnapManager installation utility as part of an upgrade, you are asked whether you want to retain your current installation directory. Because changing your installation directory might cause existing scripts that rely on the SnapManager installation directory path to fail, you should typically retain the current directory.

### Steps

1. Download the software from the NetApp Support Site.  
*[NetApp Downloads: Software](#)*
2. Exit SnapManager, if you have not already done so.
3. Double-click the downloaded .exe file.
4. Complete the pages in the SnapManager installation wizard to upgrade SnapManager.

## Upgrading SnapManager from a command line

You can quickly run the SnapManager upgrade program unattended, in silent mode, from the Windows command line.

### Steps

1. Download the product installer from the NetApp Support Site.  
*[NetApp Downloads: Software](#)*

2. Exit SnapManager, if you have not already done so.
3. From a Windows command prompt of the target host, change to the location where you downloaded the product installer.
4. Choose the type of silent upgrade you want to upgrade SnapManager:

- Silent upgrade with user credentials:

```
installer.exe /s /v"/qn REINSTALLMODE=vomus REINSTALL=ALL
SVCUSERNAME=Domain\UserName SVCUSERPASSWORD=Password
SVCCONFIRMUSERPASSWORD=Password [/L*V DirPath\LogFileName]"
```

- Silent upgrade without user credentials:

```
installer.exe /s /v"/qn REINSTALLMODE=vomus REINSTALL=ALL [/L*V
DirPath\LogFileName]"
```

- *installer* is the name of the .exe file.
- *Domain\UserName* is the user account that Windows uses to run SnapManager.

This SnapManager service account must have specific permissions on the Windows host and the Exchange Server. For more information, see *Installation and Setup Guide*:

[SnapManager 7.2 for Microsoft Exchange Server Installation and Setup Guide for Data ONTAP operating in 7-Mode](#)

[SnapManager 7.2 for Microsoft Exchange Server Installation and Setup Guide For Clustered Data ONTAP](#)

- *Password* is the password for the specified user account.
- *DirPath\LogFileName* is the location and name of a log file, which is useful for troubleshooting.

The asterisk (\*) specifies that all installation information (such as status messages, non-fatal warnings, and error messages) should be logged.

### Example

The following examples display how to upgrade SnapManager:

```
"SME7.1_x64.exe" /s /v"/qn REINSTALLMODE=vomus
REINSTALL=ALL SVCUSERNAME=MKTG2\Administrator
SVCUSERPASSWORD=password SVCCONFIRMUSERPASSWORD=password /L*V C:
\SME_upgrade.log"
```

```
"SME7.1_x64.exe" /s /v"/qn REINSTALLMODE=vomus
REINSTALL=ALL /L*V C:\SME_upgrade_without_user.log"
```

# Repairing, reinstalling, and uninstalling SnapManager

---

You can repair, reinstall, or uninstall SnapManager as needed.

## Repairing SnapManager

In many cases, you can correct stability issues with SnapManager by repairing the software. Repairing the software fixes missing or corrupt files, shortcuts, and registry entries.

### Steps

1. In **Control Panel**, navigate to your installed programs.
2. Select **SnapManager for Microsoft Exchange Server**.
3. Click **Repair**.

### Result

Windows configures SnapManager for Microsoft Exchange Server.

## Reinstalling SnapManager

You might reinstall SnapManager if you want to install an older version of it or if you ran a repair operation that did not resolve a stability issue within the software.

### About this task

You do not need to stop Exchange Server before or during the SnapManager software reinstallation process.

### Steps

1. Record the locations of existing SnapInfo directories.
2. Uninstall SnapManager.
3. Reinstall SnapManager.
4. Configure SnapManager with the same SnapInfo directory locations that were used by SnapManager before the reinstallation.

If you configure SnapManager with *different* SnapInfo directory locations than were used previously, then SnapManager no longer has records of any backups taken before you reinstalled it.

## Moving Exchange data to a local disk

Before uninstalling SnapManager, move your Exchange data store from the storage system to a local disk to ensure there is no data loss in case you encounter an issue during the uninstallation operation.

### Before you begin

Ensure there is enough space on your local disk before moving your database back to it.

### Steps

1. Use Exchange Management Console or Exchange Management Shell to move your databases, transaction logs, and system files from the LUNs onto a local disk.
2. Confirm that the data was moved correctly and that your Exchange server is functioning normally.
3. Disconnect or delete your LUNs from your SnapManager host.

For information about disconnecting or deleting LUNs, see the SnapDrive documentation.

## Uninstalling SnapManager in interactive mode

You can interactively uninstall SnapManager and all its components by using the Windows Add or Remove Programs utility. You can also remove the SnapManager report directory. Unless it is specified for a particular upgrade path or for a particular troubleshooting situation, you do not need to uninstall SnapManager before reinstalling it or upgrading to a newer version.

### About this task

Ensure that you uninstall SnapManager from all the nodes of the Database Availability Group (DAG).

**Attention:** Do not attempt to remove a currently installed version of SnapManager using an interactive method other than the approach described here; using another approach might cause system problems that result from unknown remaining files.

You do not need to stop the Exchange service or remove the Exchange databases from the LUNs before you uninstall SnapManager.

**Note:** You might need to reboot after you uninstall. The installer will inform you if this is the case.

### Steps

1. Close SnapManager.
2. In the **Control Panel**, select **Add or Remove Programs**, and then select the entry for SnapManager.
3. Click the following buttons, depending to what you want to remove:

If you want to...	Then...
<b>Remove only the SnapManager software and leave the report directory</b>	Click <b>Remove</b> .

If you want to...	Then...
Remove both the SnapManager software and the report directory	<ol style="list-style-type: none"> <li>Click Change.</li> <li>Click Remove Reports.</li> <li>Click Remove.</li> </ol>

4. Click Yes.

## Uninstalling SnapManager using the CLI

You can uninstall SnapManager by running the software installation utility from the command-line interface. You can also uninstall the SnapManager software under the control of a script for an unattended uninstallation. Unless it is specified for a particular upgrade path or for a particular troubleshooting situation, you do not need to uninstall SnapManager before reinstalling it or upgrading to a newer version.

### About this task

Ensure that you uninstall SnapManager from all the nodes in the DAG.

**Attention:** Do not attempt to remove a currently installed version of SnapManager using an unattended method other than the one described here; using another method might cause system problems that result from unknown remaining files.

**Note:** You might need to reboot after you uninstall. The installer will inform you if this is the case.

### Steps

1. From a Windows command prompt of the target host, change to the location where you downloaded the product installer.
2. Enter the following command at the command prompt:

```
installer.exe /v"REMOVE=ALL[REMOVEREPORTFOLDER=1][/L*V DirPath
\LogFileName] /qb"
```

Variable	Description
<i>installer</i>	The name of the .exe file; either <i>setup.exe</i> or <i>SME7.1.exe</i> , depending on the SnapManager installation media being used
<i>DirPath\LogFileName</i>	<p>Optional</p> <p>If you specify this option, detailed information about the installation is written to the specified log file. You can use this information to investigate details about how a particular instance of SnapManager is installed.</p> <p>The asterisk (*) is a wildcard character that specifies that all the installation information (such as status messages, nonfatal warnings, and error messages) should be logged.</p> <p><i>DirPath</i> is the fully qualified name of the directory in which the installation logs are created or overwritten.</p> <p><i>LogFileName</i> is the name of the file to which the installation information is written.</p> <p>For example:</p> <pre>/L*V C:\SME_Uninstall.log</pre>

**Example**

```
"SME7.1.exe" /v"REMOVE=ALL REMOVEREPORTFOLDER=1 /qb"
```

## **SnapManager backup archiving**

---

You can use SnapManager to create offline archives of Snapshot copies containing SnapManager backup sets. Archiving data enables you to create a complete, self-consistent replica of your data, should you have to recover it. You can use several methods to archive your data.

### **Why organizations archive data**

The main reason for archiving data is disaster recovery. Archiving helps you to recover data damaged or accidentally deleted due to human error, due to accidental deletion, hardware failure, or natural calamity. Space constraints, historical analysis, and litigation often require the older data to be archived.

### **Guidelines for archiving SnapManager backups**

Before you start archiving your data, you should consider the guidelines for choosing the type of data to be archived, naming, and using the appropriate protocols.

- Archive only verified backups.  
If you are not sure whether a backup copy is verified, use the SnapManager Restore window to check.
- Archive a complete backup set.
- Archive the most recent backup copy.
- Archive all databases together.  
Do not archive individual databases unless you know which Snapshot copies contain the appropriate databases and transaction logs for a specific time.
- Keep in mind that single backup version is used during backup archiving when multiple databases are selected to create a backup. This means that a backup version in OnCommand Unified Manager Core Package includes backups from multiple databases. When deleting a remote backup, all backups in that backup version will be deleted.
- If you use the unique naming convention, look for the Snapshot copy with the most recent date and time.
- If you did not use the unique name option when you created the backup, look for the most recent Snapshot copy in the storage system's LUN drive volume named `/exchsnap_servername_recent` or `/exchsnap_servername_recent_backupmgmtgroup`. This is for backward compatibility with earlier versions of SnapManager, which did not include the **Run Command After Operation** feature.
- Do not use the CIFS or NFS protocols to archive LUNs.  
Use the storage system's `dump` command or an NDMP backup application to archive LUNs.

**Note:** If the system is busy, the network is slow, or the load is primarily on the DataFabric Manager server or the storage system, there is a time lag between the creation of a backup and the appearance of the archive in the Restore view.

- Consider the following factors:
  - The archive method you use
  - Service Level Agreements for disaster recovery

- The number of SnapManager backups performed each day
- Exchange client activity schedules
- Backup verification time

## Methods of archiving SnapManager backups

You can use different methods for selecting the components of your backup to archive.

- Use Network Data Management Protocol (NDMP) or the storage system's `dump` command to archive LUNs directly from the storage system to the archive medium.
- Mount the LUNs in a SnapManager backup Snapshot copy and share it; then use an industry-standard backup utility to copy the LUNs contents to the archive medium.
- Create an Exchange copy-type backup directly on the archive medium using an Exchange-aware backup application.

## Archives created with NDMP or the dump command

You can use Network Data Management Protocol (NDMP) or the Data ONTAP `dump` command to archive each of the LUNs that contain data for the backup set that you want to archive directly from the storage system to the archive medium, without involving Exchange or the Exchange server at all.

NDMP and the `dump` command are the most efficient ways to archive LUN drive files. LUN Snapshot copies are made, copied to the archive medium, and deleted.

For more information about backing up storage system data to tape, see the *Data Protection Guide* for your version of Data ONTAP.

When using NDMP or the `dump` command to archive your SnapManager backups, specify the database LUN by using its absolute path name, in which the three variables represent the following text strings:

```
/vol/volume_name/.snapshot/snapshot_name/LUN_name
```

### **volume\_name**

Name of the volume containing the data to be archived

### **snapshot\_name**

Name of the Snapshot copy containing the LUNs to be archived

### **LUN\_name**

Name of the LUN containing the data to be archived

For example, the absolute path name of the LUN `exch1db.1un` in the Snapshot copy `exchsnap__SRVR3_01-20-2006_12.05.58_Daily` on the volume `ExchVoln` is represented as follows:

```
/vol/ExchVol/.snapshot/exchsnap__SRVR3_01-20-2006_12.05.58__Daily/exch1db.1un
```

## Evaluation of the NDMP and dump command method of archiving

The NDMP and the `dump` command are the most efficient methods of creating archives of the LUN drive files, you archive more data than you need.

### Advantages

- Because the NDMP and the `dump` command methods do not rely on mounting a Snapshot copy, you do not risk creating busy Snapshot copies.
- Because the NDMP and the `dump` command methods archive the entire raw LUN, restoring involves replacing the LUNs.
- If your archive procedure does not send the data over the network, the NDMP and the `dump` command methods can be significantly faster than other methods.

### Disadvantages

- Because you are archiving raw LUNs, the entire LUN containing the Exchange data is archived, so you archive more data than you need. If archiving extra data is undesirable, you can use an industry-standard backup utility to back up the corresponding SnapInfo directory. Coordinate this so that the two pieces of the archive are kept together for later retrieval.
- If you archive the SnapInfo directory separately, you must ensure that you get the SnapInfo directory backed up directly from the Exchange server and the Exchange data extracted from the LUN backed by Snapshot copy from different locations into the same archive.

## Example using Run Command After Operation feature to archive SnapManager backups

After creating SnapManager backups, you can run a script to archive your backups to tape for long-term data security. You can accomplish this either by using the Run Command After Operation feature or by hard-coding the Snapshot copy names into the script.

Assume that you want to run a script in the following environment:

- Run the script on the computer running Exchange and SnapManager.
- The name of the Exchange server is SRVR3.
- The name of the storage system is `storagesystem1`.
- The name of the LUN containing the Exchange databases is `exch1db.lun`.
- The name of the LUN containing the Exchange transaction logs and the SnapInfo directory is `exch1logs.lun`.
- The name of the volume that contains the Exchange database LUN is `ExchDBvol`.
- The name of the volume that contains the Exchange transaction log LUN is `ExchLogvol`.
- The NetApp Data ONTAP PowerShell Toolkit is installed on the Exchange server.

To run a script using the SnapManager Run Command After Operation feature, use the following value with the `-RunCommand` parameter:

```
C:\scripts\archive.bat $ExchSnapshot $InfoSnapshot
```

The values passed into the script provide values similar to the following, respectively:

```
exchsnap__SRVR3_01-20-2012_12.05.58__Daily
```

```
exchloginfo__SRVR3_01-20-2012_12.05.58__Daily
```

The archive.bat script calls a Windows PowerShell script that passes the values into the PowerShell script:

```
powershell.exe C:\Scripts\archive.ps1 %1 %2
```

If you do not use the Run Command After Operation feature and you do not use the unique naming convention for your backups, you can hard code the Snapshot copy names into the script.

For example, the dump command for the database might look like this PowerShell command, for an Exchange server named SRVR3:

```
invoke -nassh -command "dump Of nrst0a" /vol/ExchDBVol/.snapshot/
exchsnap__SRVR3__recent__Daily/exch1db.lun
```

In the following script, the archive.ps1 script performs the archive of the SnapManager backup to tape via ndmp:

```
#Backup up the LUN in the snapshot containing the Exchange database.
#For this example, a level 0 backup is performed
#to the tape storage device named nrst0a.
import-module dataontap

$ExchSnapshot = $args[0]
$InfoSnapshot = $args[1]


$ExchDBPath = "/vol/ExchDBVol/.snapshot/" + $ExchSnapshot + "/"
exch1db.lun"
$ExchLogPath = "/vol/ExchLogVol/.snapshot/" + $ExchSnapshot + "/"
exch1db.lun"
$ndmpDBcmd = "dump Of nrst0a" + $ExchDBPath
$ndmpLogcmd = "dump Of nrst0a" + $ExchLogPath

Connect-NaController -Name storageSystem1 -Transient
invoke-nassh -command $ndmpDBcmd
invoke-nassh -command $ndmpLogcmd
```

## Archives created using an industry-standard backup utility

If you want to use an industry-standard backup utility to archive your SnapManager backups, you can mount the LUNs backed up by the Snapshot copy that you want to archive, and then use the backup utility to copy the archive data to your archive medium.

In this case, the NTFS data is backed up, rather than the raw LUNs, as shown in the following diagram:



**Note:** You do not need to mount the LUN on the Exchange server; you can use another computer for archiving.

Your archive must include the following components:

- The SnapInfo directory backed up directly from the Exchange server
- The Exchange data extracted from the LUN backed up by a Snapshot copy

## Evaluation of an industry-standard backup utility method of archiving

Using an industry-standard backup utility to archive your SnapManager backups enables you to select exactly which data you archive; however, you must be careful to avoid scheduling any backups while the archiving is performed.

### Advantages

- Because you are archiving NTFS data rather than raw LUNs, you can archive exactly the data that you need, and no more.
- The procedures and tools used for this method are familiar and available to you.

### Disadvantages

- Because this method relies on mounting a Snapshot copy, you must be careful to avoid scheduling any backups during the archive process, because creating a Snapshot copy of a mounted Snapshot copy results in a Snapshot copy that cannot be deleted.
- You must make sure that you get the SnapInfo directory backed up directly from the Exchange server and the Exchange data extracted from the LUN backed by a Snapshot copy from different locations into the same archive.

## Example: Using an industry-standard backup utility to archive SnapManager backups

You can use a script to mount a Snapshot copy of the LUNs that contain the Exchange databases; then back up the databases using an industry-standard backup utility, unmount the LUNs, and back up the SnapInfo directory.

Assume that you use a script in the following environment:

- The script is run on the computer running Exchange and SnapManager.
- The drive letter for the Exchange database is S.
- The Snapshot copy is mounted as drive V.

The drive letter used for the LUN mount must be available when you run the script.

The following script mounts a Snapshot copy of the LUNs that contain the Exchange databases, and it backs them up using an industry-standard backup utility. It then unmounts the LUNs and backs up the SnapInfo directory. The /N and /D options are used to show the name of the tape and the description options.

```

REM Mount a LUN backed by the Snapshot copy of the Exchange
REM database as drive letter V:\.
sdcli snap mount -k s -s %1 -d v
REM Use Windows Backup to back up the database files. The path
REM to the databases in your environment might be different.
ntbackup backup "V:\Program Files\Exchsrvr\mdbdata" /N %1 /D %2
REM Dismount the Snapshot copy mounted as drive letter V:\.
REM The below example is shown with an optional parameter (-f)
REM to forcefully disconnect the drive letter.
sdcli snap unmount -d v -f
REM Use Windows Backup to back up the snapinfo directory.
REM This backup appends the media so as not to
REM overwrite the database backup. The path to the
REM snapinfo directory is passed as the third
REM parameter on the command line used to launch
REM this script.
ntbackup backup %3 /T%1 /A

```

To run this script using the SnapManager Run Command After Operation feature, use the following command:

```
C:\SnapManager\Scripts\scriptname.txt $ExchSnapshot $InfoSnapshot
```

The command-line parameters %1 and %2 might provide values similar to the following, respectively:

- exchsnap\_\_SRVR3\_01-20-2012\_12.05.58\_\_Daily
- eloginfo\_\_SRVR3\_01-20-2012\_12.05.58\_\_Daily

If you prefer not to use the Run Command After Operation feature and you are not using the unique naming convention for your backups, you can hard-code (manually enter) the Snapshot copy names into the script.

For example, the `dump` command for the database might look like this, for an Exchange server named SRVR3:

```
rsh storageSystem1 dump 0f nrst0a
/vol/Exch/.snapshot/exchsnap__SRVR3__recent__Daily/exch1db.lun
```

## How to use the Exchange Backup Agent to archive Exchange backups

You can use Exchange backup copies to archive your SnapManager backup copies. You use Windows Backup (which, in turn, uses the Exchange Backup Agent) to create a backup copies of your Exchange databases. In this case, the Exchange data itself is captured and archived.

Because you do not use SnapManager to create the backup copy, you also do not use SnapManager to perform the restore process, so the SnapInfo directory does not need to be archived.

**Note:** Use a copy or differential backup for this method. Performing any other type of backup operation causes your existing SnapManager backup copies to become unusable for an up-to-the-minute restore operation.

### Evaluation of the Exchange Backup Agent method of archiving

When you use Exchange Backup Agent to archive Exchange backup copies, you do not need to archive the SnapInfo directory using the Exchange Backup Agent method, and Exchange data itself is captured and archived.

#### Advantage

The advantage of this approach is that the procedures and tools used for this method are familiar and available to you.

#### Disadvantage

The disadvantage of this approach is that no SnapManager backup operation can be initiated while this method is in use.

### Example: Using Exchange Backup Agent to archive Exchange backup copies

You can run a script to archive your SnapManager backup copies to tape using Exchange Backup Agent. The script launches Windows Backup and tells it to perform a copy-type backup operation through the Exchange Backup Agent.

#### Example script

Before you can use this script, you need to use the Windows Backup GUI to create a .BKS file that contains the databases you want to back up. This script assumes that it is run on the computer running Exchange and SnapManager.

```
REM Launch Windows Backup
ntbackup backup "@c:\MyConfig.BKS" /M copy
```

**Note:** Although this method does not use SnapManager backup copies, it can be run using the Run Command After Operation feature.

## If you use a centralized backup model

When you use a centralized backup model, you can initiate a backup job by running a command on a centralized backup server, or by running a command on the Exchange server through the third-party agent that came with your backup software.

Use your centralized backup software documentation to determine which command to run and which server to run it on. Then, you can use the Run Command After Operation feature to trigger the backup.

If you plan to run the command on the Exchange server, then you need take no extra steps.

If you plan to run the command on a centralized backup server, install SnapManager on it. Configure the SnapManager application on the Exchange server to run the command on a remote machine (the centralized backup server).

**Note:** To run the command remotely, you need to install SnapManager on the remote server even if it does not have Exchange server installed. Although you need to install SnapManager, you do not need to completely configure it.

You can run a command remotely in many ways, but this document only covers doing it through SnapManager.

## Automatic backup archiving using the Run Command After Operation feature

SnapManager enables you to automatically run your own program or script after a backup or database verification operation, but only after the operation finishes successfully. This feature, called Run Command After Operation, is typically used to automatically archive a backup copy.

Make sure that you archive all components of a backup set together. If any one of the databases is offline and you initiate a backup process using the Run Command After Operation, the backup operation fails for that particular job.

### Command arguments supported by the Run Command After Operation feature

The Run Command After Operation feature supports some predefined variables that can pass operation-specific information to your program or script.

#### \$ExchSnapshot

Expands to the name of an Exchange database Snapshot copy, as shown in the following examples:

```
exchsnap__winsnrvr3__01-31-2012_15.03.09
exchsnap__winsnrvr3__recent
```

The number of database Snapshot copies in a SnapManager backup set depends on the number of volumes used to store the databases included in the backup operation.

#### \$InfoSnapshot

Expands to the name of a SnapInfo directory Snapshot copy, as shown in the following examples:

```
eloginfo__winsnrvr2__01-31-2012_15.03.09
eloginfo__winsnrvr2__recent
```

**\$SnapInfoName**

Expands to the name of the SnapInfo directory, as shown in the following examples:

WINSRVR3\_01-23-2012\_16.21.07\_Daily

WINSRVR3\_recent

If you use this variable, you must also provide the correct path to the directory.

**\$SnapInfoPath**

Expands to the name of the SnapInfo subdirectory, as shown in the following examples:

I:\SME\_SnapInfo\EXCH\_WINSRVR3\SG\_database\_name  
\WINSRVR5\_01-06-2012\_11.52.36

The \$SnapInfoPath variable is automatically enclosed within double quotation marks so that the actual path name can contain spaces without affecting the script invocation on the Windows command line.

If you do not want the double quotes to appear in your command line, remove them from the Command Arguments field in the Run Command After Operation dialog box.

## Specifying the command to be run by the Run Command After Operation feature

While specifying a backup or verification operation, you can use Run Command After Operation to specify the details of the command that you want to run after the operation finishes. Using default values for the command arguments frees you from entering the command information each time you initiate a backup or verification operation.

### Before you begin

- You must install SnapManager on a computer that you specify, but you do not need to configure it.
- If your script is stored on a network share, use the UNC path rather than the network drive letter to specify the script's location.

Add the network location to Internet Explorer's list of trusted sites; otherwise, your command might fail.

### About this task

If more than one Snapshot copy of the same type exists for your command, include that Snapshot copy's variable name once for every Snapshot copy. If you do not, you cannot create a complete archive. For example, if you enter myscript \$exchSnap \$exchSnap, the command that is generated is myscript snap1 snap2.

### Steps

1. Use the Backup wizard or the Backup and Verify window to specify a full database backup, a transaction log only backup, or a database verification.
2. Select **Yes, Run a command after this operation** in the Backup wizard or **Run Command After Operation** in the Backup and Verify window.
3. Under **Specify a computer where...**, select the host on which your program or script resides.
4. Under **Specify the full path...**, select your program or script.
5. Type the command input string in **Command Arguments**.

You can directly enter text or select the variable you want to enter from the SnapManager Variables list.

6. Repeat Steps 4 and 5 as needed until the Command Arguments box contains all of the arguments that you want to pass to your program or script.
7. Click **OK** and return to the backup or verification setup process.

After the backup or verification finishes successfully, the command that you specified, with the Snapshot copy and SnapInfo information for this backup operation inserted, runs automatically.

#### **Example: script and the required parameters**

Suppose that you want to run the following script:

```
C:\SnapManager Scripts\scriptname.bat
```

You need the following parameters:

- Name of the database Snapshot copy
- Name of the specific SnapInfo directory
- NTFS path to the SnapInfo directory

## **Enabling the launch of SnapManager scripts from a UNC path**

If SnapManager is installed on a Windows Server system and you want to launch a script from a Uniform Naming Convention (UNC) path, you need to add the host where the script is located to the Internet Explorer list of trusted sites.

#### **About this task**

- If SnapManager is installed on a Windows Server system and you attempt to launch a script from a UNC path, SnapManager might hang. When this happens, the log file shows no indication of the failure and the last line of the log text “The specified command has been launched successfully.” does not appear.
- The supported Windows Server ships with Internet Explorer's Enhanced Internet Explorer Security Configuration enabled. This setting is restrictive and prevents batch files located on a network share from running. The network location needs to be added to Internet Explorer's list of trusted sites.
- The system you log on to can be the same system that is running SnapManager, or it can be a different server. Log on using the same user account that SnapManager is configured to use.
- The security settings are specific to each user account on each machine. You need to repeat this procedure for on all machines that are used to run scripts and for all users who might launch the script on each machine. You will also need to repeat these steps if you start using a different user account to run SnapManager.

#### **Steps**

1. Log in to the Windows Server system that will be running the script.
2. Launch Internet Explorer.
3. In the menu bar, select **Tools > Internet Options**.

4. In the **Internet Options** dialog box, select **Security > Local intranet > Sites**.
5. In **Add this Web site to the zone**, enter the host name of the machine from which the script will be launched and click **Add**.
6. Click **OK**.
7. To verify your changes, browse to the network on which the script resides and launch the script. The script should now run normally without security prompts. The script might fail to complete properly because the SnapManager variables are not passed to the script in this test.

## **SnapManager command-line reference**

SnapManager 6.1 for Microsoft Exchange provides SnapManager command-line functionality, enabling you to create scripts to run SnapManager without using the graphical user interface (GUI).

### **Guidelines for using the SnapManager for Exchange PowerShell command-line interface**

Before you start using the SnapManager for Exchange PowerShell command-line interface, you must be aware of parameter and option presentation requirements.

To run the Microsoft Exchange cmdlets, use the Microsoft Exchange Management Shell.

Observe the following guidelines when using the SnapManager command-line interface:

- All parameters and options are case-insensitive; for example, using the option `-Daily` gives the same results as using `-daily`.
- Some of the options must be invoked in a particular order.  
For best results, use the order given in the syntax for all options.
- When a parameter value string contains spaces, be sure to enclose it in single quote; for example, use `'database 1'` rather than `database 1`.

### **Launching SnapManager for Exchange PowerShell**

You can launch the SnapManager for Exchange PowerShell from the Windows Start menu to use the command-line interface to perform various SnapManager operations.

#### **Step**

1. Open the SnapManager for Exchange PowerShell interface.

Depending on your Windows operating system, perform one of the following:

For Win2008:

- Select **Start > Programs > NetApp > SnapManager for Exchange PowerShell**.

For Win2012:

- Click **Start > SnapManager for Exchange PowerShell**.

The SnapManager for Exchange PowerShell command-line interface is displayed.

#### **After you finish**

On the Windows Server system with **User Access Control** enabled, use the menu option **Run As Administrator** to prevent any access errors related to the Windows Registry, or any other Windows server resources.

## new-backup

The new-backup command enables you to initiate a backup or verification job, with all of the options available through the SnapManager GUI.

```
new-backup

-ClusterAware True|False

-VerifyOnDestVolumes src_storage_system_list:src_vol
    dest_storage_system:dest_vol
-Verify
-Server server_name

-Database 'database1', 'database2',...
-ManagementGroup Standard|Weekly|Daily
-ActiveDatabaseOnly True|False
-PassiveDatabaseOnly True|False
-BackupTargetServer server_name
-ActivationPreference ActivationPreferenceNum
-UpdateMirror True|False
-VerDestVolume True|False
-NoUTMRestore True|False
-NoTruncateLogs False
-Throttle throttle_val
-VerificationServer server_name
-UseMountPoint True|False

-MountPointDir mountpoint_dir
-RetainBackups no_of_days_to_retain_backup
-RetainDays no_of_days_delete_backup
-Command True|False
-RunCommand win_path_and_script_name
-GenericNaming True|False

-RecoveryPoint win_path_and_script_name
-ReportProgress True|False
-ArchiveBackup True|False

-ArchivedBackupRetention Hourly|Monthly|Daily|Weekly|Unlimited
-RetainUtmBackups no_of_log_backups_to_retain
-RemoteAdditionalCopyBackup [Boolean]
-AdditionalCopyBackupDAGNode [String]
-RetainRemoteAdditionalCopyBackup [Integer]
-RetainRemoteAdditionalCopyBackupDays [Float]
-ArchiveRemoteAdditionalCopyBackup [SwitchParameter]
common parameters
```

### Description

The new-backup command enables you to back up Exchange databases.

This command also supports the following common parameters:

- -Debug (-db)
- -ErrorAction (-ea)
- -ErrorVariable (-ev)
- -OutBuffer (-ob)
- -OutVariable (-ov)

- `-Verbose (-vb)`
- `-Confirm`

To learn more about common parameters, see [Help about\\_ubiquitous\\_parameters](#).

## Parameters

### **-ClusterAware** `True|False`

Short name: `c1`

Assumes significance only when scheduling jobs in cluster configurations, by facilitating scheduling the same job in multiple cluster nodes to improve fault tolerance.

In the case of DAG, if a job is scheduled with `-ClusterAware`, the job runs only if the host in which it is scheduled is the active node of the DAG.

### **-verifyOnDestVolumes** `src_storage_system_list:src_vol:dest_storage_system:dest_vol`

Short name: `vermirror`

Overrides the existing SnapMirror relationships.

### **-verify**

`Verify` is a switch parameter that can be specified when the command is run. If the parameter is specified, the Windows PowerShell runtime resolves its value as true. If the parameter is not specified, which is typically the default, the parameter value is resolved as false.

Short name: `ver`

Verifies the backed up SnapManager databases and transaction logs.

### **-Server** `Exchange_server_name`

Short name: `svr`

Specifies the target Exchange server name. You can also specify the DAG name.

In a cluster configuration, you must specify `-Server` explicitly in all of the cmdlets to perform all operations. If `-Server` is not specified explicitly in a stand-alone server, SnapManager uses the local machine as the default to run the following cmdlets: `new-backup`, `verify-backup`, `restore-backup`, `get-backup`, and `delete-backup`.

### **-Database** `'database1', 'database2',...`

Short name: `dbs`

Lists the databases in the following format:

`-Database 'DBName1', 'DBName2'`

If you do not specify `Database`, the cmdlet backs up all the databases on the server specified with the `-server` parameter.

You can back up all mailbox databases in the DAG without specifying each member server of the DAG by using the following command:

`new-backup -Server DAG1`

If a database (for example, DB1) is replicated on three member servers in the DAG, you can back up this database by using the following command:

`new-backup -Server DAG1 -Database 'DB1'`

This backs up all three copies of the database on all three member servers.

**Note:** If the server name specified is a DAG name, and if the `-ActiveDatabaseOnly` or the `-PassiveDatabaseOnly` option is not

specified, to select a database for backup you have to specify both the database and the server on which the database resides. For example, specify the databases by using the following command:

**-Database 'db1\server1, db1\server2'**

**-ActiveDatabaseOnly True|False**

Short name: activedb

Backs up the active databases. This is a switch parameter.

**Note:** If neither the **-ActiveDatabaseOnly** nor the **-PassiveDatabaseOnly** option is specified, all passive and active databases in the DAG are included in the backup.

**-PassiveDatabaseOnly True|False**

Short name: passivedb

Backs up the passive databases. This is a switch parameter.

**Note:** If neither the **-ActiveDatabaseOnly** nor the **-PassiveDatabaseOnly** option is specified, all passive and active databases in the DAG are included in the backup.

**-BackupTargetServer *server name***

Short name: bkupsrv

Backs up the databases on the specified server. If you do not specify **BackupTargetServer**, mailbox databases on all member servers in the DAG are backed up. For example, the command `new-backup -server DAG1` backs up all databases on all member servers in the DAG.

**-ActivationPreference *ActivationPreferenceNum***

Short name: actpref

Back up the databases with the specified ActivationPreference number.

**Note:** You can use the following cmdlet from the Microsoft Exchange 2010 and 2013 PowerShell to get the list of ActivationPreference numbers on the member servers for a database: `Get-MailboxDatabase -Identity databasename | fl`

**-ManagementGroup Standard|Weekly|Daily**

Short name: mgmt

Specifies the frequency of a backup or verify operation that is scheduled to be performed on a daily, weekly, or standard basis.

**-UpdateMirror True|False**

Short name: updmir

Starts a SnapMirror synchronization after the backup operation.

SnapMirror updates the specified volume to reflect incremental updates to a source volume. If a SnapManager volume is enabled for SnapMirror use, the SnapMirror destination is updated from the source volume.

**-VerDestVolume True|False**

Short name: verdest

Verifies the SnapMirror destination volume.

**-NoUTMRestore True|False**

Short name: noutm

Denies the retention of up-to-the-minute restore ability.

The logs are also deleted for any backup copies that you delete as part of this backup operation. This does not retain the up-to-the-minute restore ability for older backup copies that remain after the delete phase of a backup operation.

**-NoTruncateLogs *False***

Short name: notrunc

Denies the backup of truncated transaction logs. This option can be used to conserve space on the LUN containing the backed up Exchange transaction logs.

**-Throttle *throttle\_val***

Short name: throt

*throttle\_var* is an integer value that defines the throttling value to be used during the verification operation.

**-VerificationServer *server\_name***

Short name: versvr

Overrides the preconfigured SnapManager verification settings. It denotes the host to be used as the verification server for the verification phase of a backup operation. This is a switch parameter.

**-UseMountPoint *True|False***

Short name: mp

Mounts the Snapshot copy to a NTFS directory. During a SnapManager verification operation, the Snapshot copies are mounted in a default NTFS directory for database verification. This option is effective when there are no drives available to mount the Snapshot copies during database verification. The value of this parameter overrides the preconfigured SnapManager verification settings.

**-MountPointDir *mountpoint\_dir***

Short name: mpdir

Specifies which mount-point directory a Snapshot copy is to be mounted to, during database verification.

**-RetainBackups *no\_of\_days\_to\_retain\_backup***

Short name: rtbackups

Specifies the number of backup copies to retain after the delete phase of a SnapManager backup operation.

**-RetainDays *no\_of\_days\_to\_delete\_backup***

Short name: rtdays

Specifies the number of days after which a backup copy is to be deleted. Specifies that the deletion of backup copies is to be based on an "older than number of days" policy. Use this option in conjunction with RetainBackups. Absence of this option denotes a deletion policy based on retained backup count.

**-Command *True|False***

Short name: cmd

Indicates that RunCommand is to be used after the current operation. This is a switch parameter.

**-RunCommand *win\_path\_and\_script\_name***

Short name: runcmd

Specifies the complete path name of and runs the specified command after the SnapManager backup or verification operation is complete.

**Note:** You must specify this command explicitly; the preconfigured command does not run after the backup or verification operation.

**-GenericNaming True|False**

Short name: gen

Specifies the generic naming convention to be used for the SnapManager backup sets.

**-RecoveryPoint *win\_path\_and\_script\_name***

Indicates a Frequent Recovery Point backup.

**Note:** If you specify the parameter new-backup, all the other parameters except -Server, StorageGroup-UpdateMirror, -ClusterAwareLcrBackupCopyRemoteCCRNode are ignored.

**-ReportProgress True|False**

Short name: repprog

Displays the operation status and progress information in the PowerShell output. If you do not use this switch parameter, the progress information is logged only to the report file, and not to the PowerShell output.

**Note:** Avoid using this switch for scheduled backup and verification jobs.

**-ArchiveBackup True|False**

Short name: arch

Creates an archive of the backup copy created on the primary node. The primary node is the node on which the backup operation is initiated. Include this parameter if you have datasets configured on your primary node and want to archive the backup copy to the SnapVault secondary storage system.

**-ArchivedBackupRetention Hourly|Monthly|Daily|Weekly|Unlimited**

Short name: archret

Determines the retention time for the archives that were created using the ArchiveBackup parameter. The retention can be hourly, monthly, daily, weekly, or unlimited.

**Note:** Do not use ArchivedBackupRetention without using the parameter ArchiveBackup. If you use the parameter ArchiveBackup only, the daily retention type is used by default.

**-RetainUtmBackups *no\_of\_log\_backups\_to\_retain***

Defines the number of log backups to retain.

**-RemoteAdditionalCopyBackup [Boolean]**

Short name: remCopyBk

Specifies any backups that also have a remote additional backup copy.

Default value: true = DAG

If this parameter is specified, if databases are replicated to other member servers, a remote additional backup (copy based without log truncation) of the selected databases of the remote nodes is created. This applies only to DAG connections.

**-AdditionalCopyBackupDAGNode [String]**

Short name: remCopyBkNode

Specifies DAG nodes allowed to have remote additional (copy-based) backups.

All DAG member nodes are allowed.

**-RetainRemoteAdditionalCopyBackup [Integer]**

Short name: `rtRemCopyBk`

Specifies the number of remote additional (copy-based) backups to keep on remote nodes.

There is no default value. This is used when the backup completes to determine which previous backups to retain.

**-RetainRemoteAdditionalCopyBackupDays [Float]**

Short name: `rtRemCopyBkDays`

Specifies the number of days to keep backups when doing a remote additional (copy-based) backup on a remote node.

Using this parameter indicates that deletion is based on an “older than number of days” policy. If this parameter is not used, the deletion policy is based on the number of retained additional backups.

**-ArchiveRemoteAdditionalCopyBackup [SwitchParameter]**

Short name: `archRemCopyBk`

Creates an archive of the remote additional (copy-based) backup created on the remote cluster member nodes.

Default value: false

Include this parameter if you have datasets configured on those nodes and want to archive the remote additional copy-based backup to another secondary SnapVault storage system.

**Example: Creating a new backup copy of two databases in standard backup management group, and deleting older backup copies**

This command creates a new backup copy of two databases on EXCHSRVR in the Standard backup management group (the default), using the unique naming convention (the backup copy is named using the date-time stamp). RunCommand is performed, and all older backup copies are deleted except for the eight most recent.

```
new-backup -Server 'EXCHSRVR' -ManagementGroup 'Standard' -  
NoTruncateLogs $False -RetainBackups 8 -Command - RunCommand 'C:  
\WINDOWS\system32\svrpt1.bat' -CommandServer 'SNAPMGR-48' -Database  
'DBName1', 'DBName2' -Verify -VerificationServer 'Snapmgr-48' -  
Throttle 150 -UseMountPoint -MountPointDir 'C:\ProgramFiles\NetApp  
\SnapManager for Exchange\SnapMgrMountPoint'
```

**Example: Creating a new backup of two databases in standard backup management group, and retaining older backup copies**

This command creates a new backup of two databases on EXCHSRVR3 in the Standard backup management group (the default), using the unique naming convention (the backup are named using the date-time stamp). No verification or other command is performed after this operation, and no older backup copies are deleted.

```
new-backup -Server 'EXCHSRVR3' -Database 'DBName1', 'DBName2'
```

### Example: Creating a new backup in Daily backup management group, verifying backup copies, and deleting older backup copies

This command creates a new backup in the Daily backup management group. The backup is verified on SRVR7, and all older Daily backup copies are deleted except for the three most recent. No SnapMirror replication is initiated after the backup, even if the volume is a SnapMirror source volume.

```
new-backup -Server 'EXCHSRVR3' -Database 'DBName1', 'DBName2' -  
VerificationServer -SRVR7 -retainbackups 3 -ManagementGroup daily
```

### Example: Creating a new backup and verifying destination volume

This command creates a backup and performs verification on the destination volumes:

```
new-backup -Server SNAPMGR-55 -ManagementGroup Standard -  
NoTruncateLogs $False -Database 'DBName1' -Verify -VerificationServer  
SNAPMGR-55 -VerDestVolume -UpdateMirror $False
```

### Example: Creating a frequent recovery backup

This command creates a frequent recovery backup from the specified database in "SNAPMGR-50":

```
new-backup -Server SNAPMGR-50 -Database 'DBName1', 'DBName2' -  
UpdateMirror -RecoveryPoint
```

### Example: Creating archive of primary backup at the secondary database

This command creates an archive of the primary backup at the secondary database location:

```
new-backup -Server exchange1 -Database 'DBName1', 'DBName2' -  
GenericNaming -ManagementGroup Standard -NoTruncateLogs $False -  
RetainBackups 8 -Verify -VerificationServer exchange2-Throttle 200 -  
UseMountPoint - MountPointDir 'C:\Program Files\NetApp\SnapManager for  
Exchange\SnapMgrMountPoint' $False -ArchiveBackup -  
ArchivedBackupRetention Monthly
```

### Example: Creating a backup of the specified databases

This command creates a backup of the specified databases on the server "SNAPMGR-DAG1":

```
new-backup -Server SNAPMGR-DAG1 -Database 'Mailbox Database  
1248233294\SNAPMGR-06-VM2', 'Mailbox Database 1333555666\SNAPMGR-06',  
'Mailbox Database 1456789\SNAPMGR-06-VM2'
```

### Example: Backing up databases on the specified server in the DAG

This command creates a backup of the specified databases on the server "SNAPMGR06":

```
new-backup -Server 'SNAPMGR06-DAG1' -ClusterAware -ManagementGroup  
'Standard' -NoTruncateLogs $False -RetainBackups 3 -dbs 'Mailbox  
Database 0294565900', 'Mailbox Database 0793619176', 'DB1' -Verify -  
UseMountPoint -BackupTargetServer snapmgr-06
```

**Example: Backing up the databases with the specified ActivationPreference number on the specified server in the DAG**

This command creates a backup of the databases with the ActivationPreference number 2 on the server "SNAPMGR06":

```
new-backup -Server 'SNAPMGR06-DAG1' -ClusterAware -ManagementGroup 'Standard' -NoTruncateLogs $False -RetainBackups 3 -Verify -UseMountPoint -ActivationPreference 2 -BackupTargetServer snapmgr-06
```

**Example: Creating a backup of active DAG databases and performing remote additional backups**

This command creates a backup of the active DAG databases and also performs a remote additional (copy-based) backup of the databases that are replicated to member servers DAG1-NODE1 and DAG1-NODE2. If datasets are configured on DAG1-NODE1 or DAG1-NODE2 member servers, then SnapVault creates additional backups on the secondary storage system.

```
new-backup -Server 'SM-DAG1' -ClusterAware -ActiveDatabaseOnly -RemoteAdditionalCopyBackup $true -ArchiveRemoteAdditionalCopyBackup -AdditionalCopyBackupDAGNode DAG1-NODE1, DAG1-NODE2 -RetainRemoteAdditionalCopyBackup 4 -BackupTypeFullBackup
```

## verify-backup

The `verify-backup` command enables you to verify the backup sets using the SnapManager for Exchange PowerShell command-line interface.

```
verify-backup
  -ClusterAware True|False
  -VerifyOnDestVolumes src_storage_system_list:src_vol:
    dest_storage_system:dest_vol
  -Server Exchange_server_name

  -ManagementGroup Standard|Weekly|Daily
  -Database 'database1', 'database1',...
  -ActiveDatabaseOnly True|False
  -PassiveDatabaseOnly True|False
  -BackupTargetServer server_name
  -ActivationPreference ActivationPreferenceNum
  -UpdateMirror True|False
  -VerDestVolume True|False
  -Throttle throttle_val
  -VerificationServer server_name
  -UseMountPoint True|False
  -MountPointDir mountpoint_dir
  -VerifyBackups no_of_backups_to_verify
  -ReportProgress True|False
  -ArchiveBackup True|False
  -ArchivedBackupRetention Hourly|Monthly|Daily|Weekly|Unlimited
  -VerifyArchiveBackup True|False
  common parameters
```

## Description

This command enables you to verify the SnapManager backup sets with all of the options available through the SnapManager GUI.

This command also supports the following common parameters:

- `-Debug (-db)`
- `-ErrorAction (-ea)`
- `-ErrorVariable (-ev)`
- `-OutBuffer (-ob)`
- `-OutVariable (-ov)`
- `-Verbose (-vb)`
- `-Confirm`

To learn more about common parameters, see help about\_ubiquitous\_parameters.

## Parameters

### **-ClusterAware** `True|False`

Short name: `cl`

Assumes significance only when scheduling jobs in cluster configurations, by facilitating scheduling the same job in multiple cluster nodes to improve fault tolerance.

In the case of DAG, if a job is scheduled with `-ClusterAware`, the job runs only if the host in which it is scheduled is the active node of the DAG.

### **-VerifyOnDestVolumes** `src_storage_system_list:src_vol:dest_storage_system:dest_vol`

Short name: `vermirror`

Overrides the existing SnapMirror relationships.

### **-Server** `Exchange_server_name`

Short name: `svr`

Specifies the target Exchange server name.

In a cluster configuration, you need to specify `-Server` explicitly in all of the cmdlets to perform all operations. If `-Server` is not specified explicitly in a stand-alone server, SnapManager uses the local machine as the default to run the following cmdlets: `new-backup`, `verify-backup`, `restore-backup`, `get-backup`, and `delete-backup`.

### **-Database** `'database1', 'database2', ...`

Short name: `dbs`

Lists the databases in the following format:

```
-Database 'DatabaseName1', 'DatabaseName2'
```

If you do not specify `Database`, the cmdlet verifies backups of all the databases on the server specified with the `Server` parameter.

You can verify backups of all mailbox databases in the DAG without specifying each member server of the DAG in the following way:

```
verify-backup -server DAG1
```

**Note:** If the server name specified is a DAG name, and if the `-ActiveDatabaseOnly` or the `-PassiveDatabaseOnly` option is not specified, to select a database for backup you have to specify both the database and the server on which the database resides. For example, specify the databases in the following way:

```
-Database 'db1\server1', 'db1\server2'
```

**-ActiveDatabaseOnly True|False**

Short name: `activedb`

Verifies the active databases. This is a switch parameter.

**Note:** If neither the `-ActiveDatabaseOnly` nor the `-PassiveDatabaseOnly` option is specified, backups of all active and passive databases in the DAG are included in the verify-backup operation.

**-PassiveDatabaseOnly True|False**

Short name: `passivedb`

Verifies the passive databases. This is a switch parameter.

**Note:** If neither the `-ActiveDatabaseOnly` nor the `-PassiveDatabaseOnly` option is specified, backups of all active and passive databases in the DAG are included in the verify-backup operation.

**-BackupTargetServer *server name***

Short name: `bkupsvr`

Verifies the databases on the specified server.

If you do not specify `BackupTargetServer`, backups of mailbox databases on all member servers in the DAG are verified. For example, the command `verify-backup -server DAG1` will verify the backups of databases on all member servers in the DAG.

**-ActivationPreference *ActivationPreferenceNum***

Short name: `actpref`

Verifies the backups of databases with the specified `ActivationPreference` number on each member server.

**Note:** You can use the following cmdlet from the Microsoft Exchange 2013 PowerShell to get the list of `ActivationPreference` numbers on the member servers for a database:

```
Get-MailboxDatabase -Identity databasename | fl
```

**-ManagementGroup Standard|Weekly|Daily**

Short name: `mgmt`

Specifies the backup or verify operation that is scheduled to be performed on a daily, weekly, or standard basis.

**-UpdateMirror True|False**

Short name: `updmir`

Starts a SnapMirror synchronization after the backup operation.

SnapMirror updates the specified volume to reflect incremental updates to a source volume. If a SnapManager volume is enabled for SnapMirror use, the SnapMirror destination is updated from the source volume.

**-VerDestVolume True|False**

Short name: verdest

Verifies the SnapMirror destination volume.

**-Throttle *throttle\_val***

Short name: throt

Defines the throttling value to be used during the verification operation.

**-VerificationServer *server\_name***

Short name: versvr

Overrides the pre-configured SnapManager verification settings. It denotes the host to be used as the verification server for the verification phase of a backup operation. This is a switch parameter.

**-UseMountPoint True|False**

Short name: mp

Mounts the Snapshot copy to a NTFS directory. During a SnapManager verification operation the Snapshot copies are mounted in a default NTFS directory for database verification. This option is effective when there are no drives available to mount the Snapshot copies during database verification. The value of this parameter overrides the pre-configured SnapManager verification settings.

**-MountPointDir *mountpoint\_dir***

Short name: mpdir

Specifies which mount point directory a Snapshot copy is to be mounted to, during database verification.

**-VerifyBackups *no\_of\_backups\_to\_verify***

Short form: verbkups

Verifies the unverified SnapManager backup sets. The default value is one, but you can specify the number of backup copies that you want to verify.

**-ReportProgress True|False**

Short form: reppro

Gets the operation status and progress information in the PowerShell output. If this switch is not used, the progress information is logged on only to the report file and not to the PowerShell output.

**Note:** Avoid using this switch for scheduled backup and verification jobs.

**-ArchiveBackup True|False**

Short name: arch

Creates an archive of the backup copy created on the primary node. Primary node is the node where backup operation is initiated. Include this parameter if you have datasets configured on your primary node and want to archive the backup copy to the SnapVault secondary storage system.

**-ArchivedBackupRetention Hourly|Monthly|Daily|Weekly|Unlimited**

Short name: archret

Determines the retention time for the archives that were created using the ArchiveBackup parameter. The retention can be hourly, monthly, daily, weekly, or unlimited.

**Note:** Do not use ArchivedBackupRetention without using ArchiveBackup. If you use the parameter ArchiveBackup only, the daily retention type is used by default.

**-VerifyArchiveBackup True|False**

Short name: verarch

Specifies that the backup copy that has to be verified is an archived backup copy. If you do not specify VerifyArchiveBackup, and one local and one archived backup copy exist with the same name, SnapManager verifies the local backup copy.

**Example: Verifying a backup copy**

This command verifies the backup copies in a specified database.

```
verify-backup -Server SNAPMGR-48 -ManagementGroup Standard -Database
'DBName1' -VerifyBackups 1 -VerificationServer Snapmgr-48 -Throttle
150 -UseMountPoint -MountPointDir 'C:\Program Files\NetApp
\SnapManager for Exchange\SnapMgrMountPoint'
```

**Example: Verifying backup copies using another verification server**

This command verifies the backup copies using SNAPMGR-48 as the verification server.

```
verify-backup -Server SNAPMGR-55 -ManagementGroup Standard -Database
'DBName1','DBName2' -VerifyBackups 5 -VerificationServer SNAPMGR-48 -
UseMountPoint -MountPointDir 'C:\Program Files\NetApp\SnapManager for
Exchange\SnapMgrMountPoint'
```

**Example: Verifying the backups of active databases**

This command verifies the backups of active databases.

```
verify-backup -Server 'SNAPMGR06-DAG1' -ManagementGroup 'Standard' -
dbs 'Mailbox Database 0294565900','Mailbox Database 0793619176','DB1' -
-VerifyBackups 1 -VerificationServer 'SNAPMGR-06' -UseMountPoint -
MountPointDir 'C:\Program Files\NetApp\SnapManager for Exchange
\SnapMgrMountPoint' -ActiveDatabaseOnly
```

## delete-backup

The delete-backup command enables you to delete backup sets.

```
delete-backup

-Server server_name
-backup backup_name
-NoUTMRestore True|False

-Database 'database1', 'database2',...
-RemoteBackup True|False
-ArchiveBackup True|False
    common parameters
```

## Description

Use this command to delete backup sets.

This command also supports the following common parameters:

- `-Debug (-db)`
- `-ErrorAction (-ea)`
- `-ErrorVariable (-ev)`
- `-OutBuffer (-ob)`
- `-OutVariable (-ov)`
- `-Verbose (-vb)`
- `-Confirm`

To learn more about common parameters, see help about\_ubiquitous\_parameters.

## Parameter

### **`-Server Exchange_server_name`**

Short form: `svr`

Use this parameter to specify the name of the backup server.

**Note:** In a DAG configuration, specify `Server` explicitly in all of the cmdlets to perform an operation. If you do not specify `Server` explicitly in a stand-alone server, SnapManager uses the local machine, as default, to run the following cmdlets: `new-backup`, `verify-backup`, `restore-backup`, `get-backup`, and `delete-backup`.

**Note:** For Exchange Server , you cannot specify the DAG name in this parameter, because it is not supported.

### **`-backup backup_name`**

Specifies the name of the backup set.

### **`-NoUTMRestore True|False`**

Short name: `noutm`

Denies the retention of up-to-the-minute restore ability.

The logs are also deleted for any backups that you delete as part of this backup operation. This does not retain the up-to-the-minute restore ability for older backups that remain after the delete phase of a backup operation.

### **`-Database 'database1', 'database2',...`**

Short name: `dbs`

Specifies the name of the database to be deleted.

This is a mandatory parameter if the cmdlet is deleting a database of Exchange Server 2013.

### **`-RemoteBackup True|False`**

Deletes the archived backup copies.

If `RemoteBackup` is not specified, the local backup copies get deleted.

### **Example: Deleting a specific backup set of a specified database**

This command deletes the backup set `exchsnap__SNAPMGR-55_11-10-2006_13.39.16`.

```
delete-backup -backup exchsnap_SNAPMGR-55_11-10-2006_13.39.16
```

**Example: Deleting a specific backup set that belongs to a specified database without up-to-the-minute restore ability**

This command deletes the backup set exchsnap\_KRISHNA-SVR18\_11-01-2012\_18.12.22 without up-to-the-minute restore ability.

```
delete-backup -Database 'DBName1' -backup exchsnap_KRISHNA-SVR18_11-01-2012_18.12.22 -NoutmRestore -verbose -confirm
```

## get-backup

The `get-backup` command enables you to retrieve backup sets that meet specified criteria.

```
get-backup

-Server Exchange_server_name
-ManagementGroup Standard|Weekly|Daily

-Database 'database1', 'database2',...
-Backup name_of_the_backup
-RecoveryPoint True|False
-Details True|False
common parameters
```

### Description

This cmdlet enables you to retrieve backup sets that meet the input criteria specified in the PowerShell command-line interface.

This command also supports the following common parameters:

- `-Debug (-db)`
- `-ErrorAction (-ea)`
- `-ErrorVariable (-ev)`
- `-OutBuffer (-ob)`
- `-OutVariable (-ov)`
- `-Verbose (-vb)`
- `-Confirm`

To learn more about common parameters, see help `about_ubiquitous_parameters`.

### Parameter

**-Server *Exchange\_server\_name***

Short name: `svr`

Specifies the target Exchange server name.

In a cluster configuration, you need to specify `-Server` explicitly in all of the cmdlets to perform all operations. If `-Server` is not specified explicitly in a stand-alone server, SnapManager uses the local machine as the default to run the

following cmdlets: new-backup, verify-backup, restore-backup, get-backup, and delete-backup.

**-ManagementGroup Standard|Weekly|Daily**

Short name: mgmt

Specifies the backup copy or verify operation that is scheduled to be performed on a daily, weekly, or standard basis.

**-Database 'database1', 'database2', ...**

Short name: dbs

Lists the databases in the following format:

`-Database 'DBName1', 'DBName2'`

If you do not specify Database, the cmdlet retrieves all databases.

**-Backup *name\_of\_the\_backup***

Shows the specified full backup details.

If you do not specify Backup, get-backup displays all of the full backup copies.

**-RecoveryPoint True|False**

Shows the recovery point.

**-Details True|False**

Shows the details of the full backup copy and the recovery point.

**Example: Showing the backup copies of a management group**

This command shows the backup copies of the standard management group in the database DBName1.

```
get-backup -Server SNAPMGR-48 -ManagementGroup Standard -Database 'DBName1'
```

**Example: Showing all full backup copies**

This command shows all of the backup copies of the database DBName1.

```
get-backup -Server SNAPMGR-48 -Database 'DBName1'
```

**Example: Showing all of the full backup copies with recovery points**

This command shows all of the backup copies of the database DBName1 with recovery points.

```
get-backup -Server SNAPMGR-48 -Database 'DBName1' -RecoveryPoint
```

**Example: Showing a specific full backup copy with recovery points**

This command shows a specific backup copy of the backup set exchsnap\_snapmgr-50\_03-01-2007\_08.00.00 in the database DBName1 with recovery points.

```
get-backup -Server SNAPMGR-48 -Database 'DBName1' -Backup exchsnap_snapmgr-50_03-01-2007_08.00.00 -RecoveryPoint
```

**Example: Showing a specific full backup copy with complete details**

This command shows the backup set exchsnap\_snapmgr-50\_03-01-2007\_08.00.00 with complete details.

```
get-backup -Server SNAPMGR-48 -Database 'DBName1' -Backup
exchsnap_snapmgr-50_03-01-2007_08.00.00 - Details
```

#### Example: Showing both full backup copy and recovery points in detail

This command shows the backup set exchsnap\_snapmgr-50\_03-01-2007\_08.00.00 with the recovery points in detail.

```
get-backup -Server SNAPMGR-48 -Database 'DBName1' -Backup
exchsnap_snapmgr-50_03-01-2007_08.00.00 - RecoveryPoint -Details
```

## restore-backup

The `restore-backup` command enables you to restore the databases.

```
restore-backup

-Backup name_of_the_backup
-RestoreLastBackup restore_last_backup
-VerifyOnDestVolumes src_storage_system_list:src_vol:
  dest_storage_system:dest_vol
-Verify
-VerifyMetadata True|False
-ExhaustiveVerification True|False
-Server Exchange_server_name
-DestinationServer dest_server_name
-AutoMount True|False
-TestRestore True|False

-Database 'database1', 'database2',...
-Rehomemailbox True|False
-BkUpServer backup_server_name
-SnapInfoDirectory snapinfo_dir_path
-SnapVaultSecondary svm_name:volume
-PointInTime True|False
-VerDestVolume True|False
-VerificationServer verf_server_name

-OverrideVer True|False
-CheckLog True|False
-Destination database_name
-DestinationServer dest_server_name
-RecoveryPointTime recvry_pt_time_stamp
-CancelBackup True|False
-WaitForBackupComplete True|False
-RecoveryPoint True|False
-RestoreArchivedBackup True|False
-NoAccessToRemoteBackup True|False
-VerifyArchiveBackup True|False
  common parameters
```

### Description

This command enables you to restore backup sets with all of the options available through the GUI.

This command also supports the following common parameters:

- `-Debug (-db)`
- `-ErrorAction (-ea)`
- `-ErrorVariable (-ev)`

- `-OutBuffer (-ob)`
- `-OutVariable (-ov)`
- `-Verbose (-vb)`
- `-Confirm`

To learn more about common parameters, see `help about_ubiquitous_parameters`.

## Parameters

### **-Backup *name\_of\_the\_backup***

Short form: `bkup`

The name of the backup set that you want to restore.

### **-RestoreLastBackup *restore\_last\_backup***

Short form: `rstlast`

Restores backup copies without specifying the name.

If you try to use `Backup` and `RestoreLastBackup` together, SnapManager ignores `RestoreLastBackup` and uses `Backup` during the restore operation.

A typical usage example of the `RestoreLastBackup` parameter is as follows:

```
restore-backup -restorelastbackup = 1 -backup = "backup name"
```

If the value of `RestoreLastBackup` is 1, SnapManager ignores this parameter and uses the `Backup` during the restore operation.

**Note:** The default value of this parameter is 0, which means that SnapManager restores the latest backup. If the value is 1, SnapManager restores the second-to-latest backup.

### **-VerifyOnDestVolumes *src\_storage\_system\_list:src\_vol:dest\_storage\_system:dest\_vol***

Short form: `vermirror`

Overrides the existing SnapMirror relationships.

### **-Verify**

`Verify` is a switch parameter that can be specified when the command is run. If the parameter is specified, the Windows PowerShell runtime resolves its value as true. If the parameter is not specified, which is typically the default, the parameter value is resolved as false.

Short name: `ver`

Verifies the backed up SnapManager databases and transaction logs. `-Verify` is a switch parameter.

### **-VerifyMetadata `True|False`**

Short form: `vermetadata`

Verifies the metadata and transaction logs.

### **-ExhaustiveVerification `True|False`**

Short form: `exhver`

Performs exhaustive database verification.

### **-Server *Exchange\_server\_name***

Short name: `svr`

Specifies the target Exchange server name.

In a DAG configuration, you need to specify `-Server` explicitly in all of the cmdlets to perform all operations. If `-Server` is not specified explicitly in a stand-alone server, SnapManager uses the local machine as the default to run the following cmdlets: `new-backup`, `verify-backup`, `restore-backup`, `get-backup`, and `delete-backup`.

**-DestinationServer *dest\_server\_name***

The name of the target server where the Recovery Database is to be created.

`DestinationServer` is specified to restore to a Recovery Database during the restore operation.

**-AutoMount *True|False***

Short form: `mt`

Mounts the databases automatically after the restore operation.

**-TestRestore *True|False***

Short form: `test`

Performs a test restore operation.

Default value: `False`

**-Database '*database1*', '*database2*', ...**

Short name: `dbs`

Lists the databases in the following format:

`-Database 'DBName1', 'DBName2'`

If you do not specify `Database`, the cmdlet restores all databases.

**-Rehomemailbox *True|False***

Updates the user accounts associated with mailboxes in restored databases to point to the mailbox server with the new name.

This is an optional parameter with the `restore-backup` cmdlet of another server.

**-BkUpServer *backup\_server\_name***

Short form: `bksvr`

Specifies the name of the server on which the backup copy was created. Use this parameter only with `RestoreFromServer` where the backup copy was originally created.

**-SnapInfoDirectory *snapinfo\_dir\_path***

Short form: `sifdir`

Specifies the SnapInfo directory path for the archived backup set during the restore operation. Use this parameter only with `RestoreFromServer`.

**-SnapVaultSecondary *svm\_name:volume***

Short form: `vaultsec`

Specifies the backup vault from which you want to restore a database. This parameter applies to clustered Data ONTAP only.

Although the restore operation can take only one Snapshot copy from a secondary destination, if the `lun_DB` and `lun_LOG` are separated on different volumes on the destination, then you can provide this type of information in the following format:

`-SnapVaultSecondary svm1:volume1, svm2:volume2, svm3:volume3`

**-PointInTime <True|False>**

Short form: `pit`

Performs a point-in-time restore operation.

**-VerDestVolume True|False**

Short form: verdest

Verifies the SnapMirror destination volume.

**-VerificationServer *verf\_server\_name***

Short form: versvr

Overrides the preconfigured SnapManager verification settings. It specifies the host to be used as the verification server for the verification phase of a backup operation.

**-OverrideVer True|False**

Short form: ovr

Overrides the verification of the databases.

**-CheckLog True|False**

Short form: chklog

Specifies the transaction logs to be restored.

**-Destination *database\_name***

Specifies where the backup copy need to be restored.

You can restore to the same database or to the Recovery Database. The default value is **tosamesg**. To restore the backup copy to the Recovery Database, enter **torsg**.

**-DestinationServer *dest\_server\_name***

Short form: dstsvr

Specifies the name of the destination Exchange server.

**-RecoveryPointTime *recvry\_pt\_time\_stamp***

Specifies the recovery point timestamp.

The timestamp for each recovery point can be seen from the output of the **get-backup** cmdlet. If the specified timestamp does not match any of the recovery points shown in the backup copies, the **restore-backup** cmdlet returns an error message showing the available recovery points before and after the timestamp.

**Note:** -RecoveryPointTime option overrides **PointInTime** if you specify both.

**-CancelBackup True|False**

Pauses all the active scheduled backup jobs on the current Exchange Server, or on all nodes in the DAG environment, and cancels the current backup copy before performing the restore operation.

When the restore operation completes, SnapManager enables the paused scheduled backup jobs only. All the other inactive jobs do not change.

**-WaitForBackupComplete True|False**

Pauses all the active scheduled backup jobs on the current Exchange Server, or on all nodes in the DAG environment, and waits for the current backup operation to complete before performing the restore operation. When the restore operation completes, SnapManager enables only the paused scheduled backup jobs. All the other inactive jobs do not change.

**-RecoveryPoint True|False**

Specifies if the backup set is a Frequent Recovery Point backup.

**-RestoreArchivedBackup True|False**

Short form: `rstarchbkup`

Restores database from an archived backup.

**-NoAccessToRemoteBackup True|False**

Short form: `noaccessarchivebkup`

Specifies that there is no direct access to the secondary storage system.

SnapManager uses the proxy server to access the secondary storage system.

**-VerifyArchiveBackup True|False**

Short name: `verarch`

Specifies that the backup copy that has to be verified is an archived backup copy. If you do not specify `VerifyArchiveBackup`, and one local and one archived backup copy exist with the same name, SnapManager verifies the local backup copy.

**Example: Restoring database**

This command restores `exchnap_SNAPMGR-55_11-10-2012_13.36.24` to the specified database.

```
restore-backup -server SNAPMGR-48 -Database "DBName1" -backup
exchnap_SNAPMGR-55_11-10-2012_13.36.24
```

**Example: Restoring from an archive**

This command restores `exchnap_SNAPMGR-54_11-10-2006_14.47.18` that was created on the archived server SNAPMGR-54:

```
restore-backup -server SNAPMGR-48 -Database "DBName1" -BkUpServer
SNAPMGR-54 -backup exchnap_SNAPMGR-54_11-10-2012_14.47.18 -
SnapInfoDirectory 'K:\SME_Snap\InfoEXCH_SNAPMGR-48A\SG_WZ00\12-
04-2012_14.47.18'
```

**Example: Restoring backup sets created on different Exchange Server**

This command restores `exchnap_SNAPMGR-54_11-10-2012_14.47.18` that was created on the server SNAPMGR-54:

```
restore-backup -server 'SNAPMGR-48' -Database "DBName1" -BkUpServer
SNAPMGR-54 -backup exchnap_SNAPMGR-54_11-10-2012_14.47.18 -
SnapInfoDirectory 'K:\SME_SnapInfo'
```

**Example: Restoring a specified recovery point time backup**

This command restores `exchnap_snapmgr-50_03-01-2012_08.00.00` at the recovery point time `03-01-2012_08:55:00`:

```
restore-backup -Server snapmgr-50 -Database "DBName1" -backup
exchnap_snapmgr-50_03-01-2012_08.00.00 -RecoveryPointTime :
03-01-2012_08:55:00
```

**Example: Restoring a database from the SnapVault location**

This command restores `exchnap_snapmgr-54_11-10-2006_14.47.18` SnapVault backup from location

```

sn_vserver_dev:Vol_Exch2k13_pbkup_db1_vault,sn_vserver_dev:Vol_Exch2k13_pbkup_db1_
log_sif_vault

restore-backup -Database "MailboxDatabase01" -backup
exchsnap__snapmgr-54_11-10-2006_14.47.18 -SnapVaultSecondary
sn_vserver_dev:Vol_Exch2k13_pbkup_db1_vault,sn_vserver_dev:Vol_Exch2k1
3_pbkup_db1_log_sif_vault

```

## reseed-database

The `reseed-database` command enables you to reseed a passive database copy that is in a failed state and restore the copy (using the latest local Snapshot copy; an archived copy cannot be used) to a healthy state.

```

reseed-database

-Server Exchange_server_or_DAG_name
-Database 'database_name'
-ReseedNode server_name
-CancelBackup True|False
-WaitForBackupComplete True|False
common parameters

```

### Description

The `reseed-database` command enables you to reseed a passive database copy that is in a failed state (Failed, FailedandSuspended, or Suspended) to restore it to the latest local Snapshot copy and place the passive database copy in a healthy state.

This command also supports the following common parameters:

- `-Debug (-db)`
- `-ErrorAction (-ea)`
- `-ErrorVariable (-ev)`
- `-OutBuffer (-ob)`
- `-OutVariable (-ov)`
- `-Verbose (-vb)`
- `-Confirm`

To learn more about common parameters, see [Help about\\_ubiquitous\\_parameters](#).

### Parameters

**-Server** *Exchange\_server\_or\_DAG\_name*

Short name: `svr`

Specifies either the target Exchange server name or the DAG name.

In a cluster configuration, you must specify `-Server` explicitly in all the cmdlets to perform all operations. If `-Server` is not specified explicitly in a stand-alone server, SnapManager uses the local machine as the default to run the following cmdlets: `new-backup`, `verify-backup`, `restore-backup`, `get-backup`, `delete-backup`, and `reseed-database`.

**Note:** If you specify a DAG name, you should also provide a server name for the `-ReseedNode` parameter.

**-Database 'database\_name'**

Short name: `dbs`

This parameter is required and provides the database name in the following format:

`-Database 'DBName1'`

If you do not specify a `Database`, you will be prompted to provide a value. If you do not, the cmdlet displays an error. For example:

```
reseed-database : Cannot bind argument to parameter
'Database' because it is an empty screen.
```

**-ReseedNode 'serverI'**

Short name: `rsn`

Provides the server name in the following format:

`-ReseedNode 'server1'`

You should not specify the `-ReseedNode` parameter without a DAG name for the `-Server` parameter.

**-CancelBackup True|False**

Short name: `cancelbkup`

Pauses all the active scheduled backup jobs on the current Exchange Server, or on all nodes in the DAG environment, and cancels the current backup copy before performing the reseed operation. When the reseed operation is complete, SnapManager enables the paused scheduled backup jobs only. All the other inactive jobs do not change.

**Note:** If this parameter is set to `False` or not specified, the reseed operation is aborted if any backup operation is in progress.

**-WaitForBackupComplete True|False**

Short name: `waitbkup`

Pauses all the active scheduled backup jobs on the current Exchange Server or on all nodes in the DAG environment, and waits for the current backup operation to finish before performing the reseed operation. When the reseed operation completes, this parameter enables only the paused scheduled backup jobs. All the other inactive jobs do not change.

**Note:** If this parameter is set to `False` or not specified, the reseed operation is aborted if any backup operation is in progress.

**Example: Reseeding a passive database copy**

The following command reseeds the specified copy of the database “Mailbox Database 1248233294” on the server “SNAPMGR-VM1” (on the DAG “SNAPMGR-DAG1”):

```
reseed-database -Server SNAPMGR-DAG1 -Database 'Mailbox Database
1248233294' -ReseedNode 'SNAPMGR-06-VM2'
```

You can also use the following command to perform the same operation:

```
reseed-database -Server 'SNAPMGR-06-VM2' -Database 'Mailbox Database
1248233294'
```

**Example: Reseeding a passive database copy that is replicated on three member servers in the DAG**

If a database (for example, DB1) is replicated on three member servers in the DAG, you can reseed this database by using the following examples:

- Submit the job at the DAG level:

```
reseed-database -Server DAG1 -Database 'DB1' -ReseedNode 'Server1'
```

- Submit the job at the server level:

```
reseed-database -Server Server1 -Database 'DB1'
```

## Get-JobStatus

The `Get-JobStatus` command enables you to view the status of the queued, running, and finished jobs.

```
Get-JobStatus
  -Server Exchange_server_name
  -ShowChildJobs True|False
  common parameters
```

### Description

Specify the server name to view a particular job status. This command also supports the following common parameters:

- `-Debug (-db)`
- `-ErrorAction (-ea)`
- `-ErrorVariable (-ev)`
- `-OutBuffer (-ob)`
- `-OutVariable (-ov)`
- `-Verbose (-vb)`
- `-Confirm`

To learn more about common parameters, see `help about_ubiquitous_parameters`.

### Parameters

**-Server *Exchange\_server\_name***

Short name: `svr`

Specifies the name of the Exchange server for which you monitor the job status.

This is an optional parameter. If you do not specify this parameter, the name of the server that runs this cmdlet becomes the default host name.

**-ShowChildJobs *True|False***

Short name: `cj`

Displays all the child jobs of the running and the finished jobs.

This is an optional parameter.

**Example: Displaying all jobs**

This command displays all the jobs that are handled by the Exchange server Exchange1.

```
Get-JobStatus -Server Exchange1
```

**Example: Displaying child jobs**

This command displays all the child jobs of the running jobs, the finished jobs, and the parent-level jobs that are managed by the Exchange server Exchange1.

```
Get-JobStatus -Server Exchange1 -ShowChildJobs
```

## Change-JobPriority

If a job is queued, Change-JobPriority enables you to move a SnapManager job to a different priority in the queue.

```
Change-JobPriority
  -Server Exchange_server_name
  -JobID numeric_job_id
  -Priority position_of_the_job
  -SourceBackupServer name_of_server_thatCreates_the_backup
```

### Description

Change-JobPriority enables you to move a job into a different priority in the queue. You can view the current queue with the Get-JobStatus command.

This command also supports the following common parameters:

- -Debug (-db)
- -ErrorAction (-ea)
- -ErrorVariable (-ev)
- -OutBuffer (-ob)
- -OutVariable (-ov)
- -Verbose (-vb)
- -Confirm

To learn more about common parameters, see help about\_ubiquitous\_parameters.

### Parameters

**-Server Exchange\_server\_name**

Short name: svr

Specifies the target Exchange server name.

In a DAG configuration, you need to specify Server explicitly in all of the cmdlets to perform all operations. If Server is not specified explicitly in a standalone server, SnapManager uses the local machine as the default to run the following cmdlets: new-backup, verify-backup, restore-backup, get-backup, and delete-backup.

**-JobID *numeric\_job\_id***Short name: `id`

Used to identify a particular job that is being handled by the SnapManager server. This is a required parameter if you do not specify `AllJobs`.

**-Priority *position\_of\_the\_job***Short name: `p`

Specifies the position to which you want to move it. This is a required parameter.

**-SourceBackupServer *name\_of\_server\_thatCreates\_the\_backup***Short name: `bksvr`

Specifies the name of the Exchange server that creates the backup. This is an optional parameter. If you do not specify this parameter, the name of the source backup server specified by the parameter `Server` becomes the default name for the server.

**Example: Changing job priority**

This command changes the priority of the deferred integrity verification job that is queued in the remote verification server `VerificationServer1` (with Job ID 123) to priority 1.

```
Change-JobPriority -Server VerificationServer1 -SourceBackupServer
Exchange1 -JobID 123 Priority 1
```

## Cancel-Job

The `Cancel-Job` command enables you to cancel the jobs that are in queued or in running state.

```
Cancel-Job
  -Server Exchange_server_name
  -JobID numerical_job_id
  -AllJobs True|False
  common parameters
```

**Description**

If the job is in the queue, SnapManager removes the job from the queue. If the job is running, this cmdlet cancels the running job. This command also supports the following common parameters:

- `-Debug (-db)`
- `-ErrorAction (-ea)`
- `-ErrorVariable (-ev)`
- `-OutBuffer (-ob)`
- `-OutVariable (-ov)`
- `-Verbose (-vb)`
- `-Confirm`

To learn more about common parameters, see `help about_ubiquitous_parameters`.

**Parameters****-Server *Exchange\_server\_name***

Short name: svr

Specifies the name of the Exchange server.

If you do not specify this parameter, the default host is the name of the server that runs this cmdlet. This is an optional parameter.

**-JobID *numerical\_job\_id***

Short name: id

Identifies a particular job that is handled by the SnapManager server.

This is a required parameter if you did not specify **AllJobs**.

**-AllJobs *True|False***

Short name: all

If this parameter is set to true, all the jobs including those in running and in queued states are cancelled.

This is a required parameter if you did not specify **JobID**.

**Example: Cancelling a job**

This command cancels the job running in the Exchange server Exchange1 with Job ID 123.

```
Cancel-Job -Server Exchange1 -JobID 123
```

**Example: Cancelling a job managed by a remote integrity verification server**

This command cancels all the queued and running jobs managed by the remote integrity verification server.

```
Cancel-Job -Server VerificationServer1 -AllJobs
```

## Export-config

This cmdlet enables you to export the SnapManager configuration control file.

```
Export-config
  -Server Exchange_server_name
  -ControlFilePath name_of_control_file_and_path
  -Section comma_separated_list_of_section_names
  common_parameters
```

**Description**

The **Export-config** cmdlet enables you to export the SnapManager control file that contains the configuration information which you can later use to configure SnapManager on other systems by using the **import-config** command.

**Parameters****-Server *Exchange\_server\_name***

Short name: svr

Specifies the name of the Exchange server whose configuration you want to export as an XML control file.

**-ControlFilePath *name\_of\_control\_file\_and\_path***

Short name: config

Specifies the output XML file name and path.

**-Section *comma\_separated\_list\_of\_section\_names***

Short name: sect

Specifies the list of section names separated by comma to export.

You can export the following sections of the control file:

storage, notification, verification, report, backup, scheduledjob, snapmirrorvolume.

If you do not specify Section, SnapManager assumes that all sections should be exported.

#### Example: Exporting a control file

This command exports the specified configuration to the SMEConfig\_12\_18\_2007\_01.12.57.xml control file.

```
export-config -Server Exchange1 -ControlFilePath "C:\Program Files
\NetApp\SnapManager for Exchange\SMEConfig_12_18_2007_01.12.57.xml"
```

## Import-config

The Import-config cmdlet enables you to import the SnapManager configuration control file.

```
Import-config

-Server Exchange_server_name
-ControlFilePath name_of_control-file_and_path
-Section comma_separated_list_of_section_names_to_import
-AllowLocal True|False
-ValidateAndApply True|False
-Username username
-Password password
-ClusterAware True|False
```

#### Description

This command enables you to import the SnapManager control file that contains the server configuration information. You can import either a section of the control file, or the complete control file.

This command also supports the following common parameters:

- -Debug (-db)
- -ErrorAction (-ea)
- -ErrorVariable (-ev)
- -OutBuffer (-ob)
- -OutVariable (-ov)
- -Verbose (-vb)
- -Confirm

To learn more about common parameters, see help about\_ubiquitous\_parameters.

**Parameters****-Server *Exchange\_server\_name***

Short name: svr

Specifies the name of the Exchange server to which you want to import the control file.

**-ControlFilePath*name\_of\_control-file\_and\_path***

Short name: config

Specifies the location of the control file to import.

Specify -ControlFilePath with the control file name. If the control file is not in the current directory, the full file name path must be given.

**-Section *comma-separated\_list\_of\_section\_names\_to\_import***

Short name: sect

Specifies the comma-separated list of names of sections to import.

You can import the following sections from the control file:

storage, notification, verification, report, backup, scheduledjob, snapmirrorvolume.

If you do not specify particular sections, SnapManager imports all sections in the control file.

**-AllowLocal *True|False***

Short name: tolocal

Migrates the databases to the local disk.

**-ValidateAndApply *True|False***

Short name: apply

This is an optional parameter. By default, the value is false, that indicates to perform only validation. If the value is set to true, this parameter indicates to perform validation and apply if the validation is successful.

**-Username *username***

Short name: usr

Verifies the user name before creating a scheduled job.

**-Password *password***

Short name: pwd

Verifies the user credentials before creating a scheduled job.

**-ClusterAware *True|False***

Short form: cl

Assumes significance only when scheduling jobs in cluster configurations, by facilitating scheduling the same job in multiple cluster nodes to improve fault tolerance.

In the case of DAG, if a job is scheduled with -ClusterAware, the job runs only if the host in which it is scheduled is the active node of the DAG.

**Example: Importing sections of the control file**

This command imports the specified sections from the control file `sme_config.xml` to Exchange server Exchange1.

```
import-config -Server Exchange1 -ControlFilePath "C:\Program Files
\NetApp\SnapManager for Exchange\sme_config.xml" -Section
```

```
storage,notification,verification,report,backup,scheduledjob,snapmirro  
rvolume -ValidateAndApply
```

## **Copyright**

---

Copyright © 2019 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

## Trademark

---

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

*<http://www.netapp.com/us/legal/netapptmlist.aspx>*

## **How to send comments about documentation and receive update notifications**

---

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

*[doccomments@netapp.com](mailto:doccomments@netapp.com)*

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account [@NetAppDoc](#).

You can also contact us in the following ways:

- NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

# Index

## A

Activation Preference number  
    displayed by Database Filter display [26](#)  
Activation Preference Number  
    mailbox database [27](#)  
    original DAG database [27](#)  
active databases  
    displayed by Database Filter [26](#)  
age-based Snapshot copy deletion [28](#)  
application  
    settings [83](#)  
application settings  
    accessing [83](#)  
    restriction for configuring  
        restriction for changing [83](#)  
archives  
    creating with an industry-standard backup utility [116](#)  
archiving  
    automatic [120](#)  
    Exchange Backup Agent  
        example [119](#)  
    guidelines [113](#)  
    industry-standard backup utility  
        evaluation [117](#)  
        example [118](#)  
    methods [114](#)  
    NDMP and dump  
        evaluation [115](#)  
archiving SnapManager backups  
    using the Run Command After Operation feature  
        [115](#)  
Auto Notification  
    event settings [99](#)  
automatic deletion  
    of Snapshot copies  
        orphaned SnapInfo folders [30](#)  
    requirements for [30](#)  
automatic event notification  
    configuring [100](#)  
AutoSupport  
    configuring [100](#)  
AutoSupport notification  
    for events [99](#)  
    limiting to failure events [99](#)

## B

backing up databases  
    for the first time [14](#)  
    using a schedule [16](#)  
    using the Backup wizard [19](#)  
backup  
    using remote additional copy backup [22](#)  
Backup and Verify window  
    using to create backup copies [20](#)  
backup archives  
    creating with an industry-standard backup utility [116](#)  
    creating with Exchange Backup Agent [119](#)

backup copies  
    archiving [119](#)  
    problem deleting [35](#)  
backup copy  
    restoring to another mailbox server [60](#)  
    unmounting [67](#)  
backup grouping  
    gapless DAG example [23](#)  
backup jobs  
    planning for databases in DAGs [17](#)  
    using the Backup and Verify window [20](#)  
    using the new-backup command [125](#)  
backup management group  
    changing [45](#)  
backup model  
    centralized [120](#)  
backup naming  
    conventions [10](#)  
    options for [10](#)  
backup naming conventions  
    for Snapshot copies [10](#)  
backup operations  
    using Frequent Recovery Points [43](#)  
backup options fro transaction logs [39](#)  
backup set  
    changing backup management group [45](#)  
backup set management  
    on a production Exchange server [85](#)  
    on a remote verification server [85](#)  
backup sets  
    verifying the initial set [15](#)  
backup settings  
    XML schema [81](#)  
backup types  
    overview of [11](#)  
backup verification  
    concurrent jobs [104](#)  
Backup wizard  
    setting up and using [19](#)  
backups  
    guidelines for archiving [113](#)  
benefits and features  
    overview of [8](#)  
best practices [11](#)  
busy Snapshot copy  
    problems caused by [38](#)

## C

Cancel-Job command [149](#)  
CCR-enabled database  
    restoring using SnapManager [49](#)  
centralized backup model [120](#)  
Change-JobPriority command [148](#)  
ChkSgFiles [88](#)  
ChkSgFiles integrity verification library [90](#)  
ChkSgFiles utility  
    for verifying database files [35](#)  
cleaning up after

mailbox data [67](#)

cmdlets

- managing remote additional backups [24](#)
- Microsoft Exchange
  - usage guidelines [124](#)

command

- specifying [121](#)

command-line interface

- SnapManager for Exchange PowerShell usage
  - guidelines [124](#)
  - SnapManager,usage guidelines [124](#)

command-line reference

- SnapManager commands [124](#)

commands

- Cancel-Job [149](#)
- Change-JobPriority [148](#)
- delete-backup [136](#)
- Export-config [150](#)
- get-backup [138](#)
- Get-JobStatus [147](#)
- Import-config [151](#)
- new-backup [125](#)
- reseed-database [145](#)
- restore-backup [140](#)
- verify-backup [132](#)

comments

- how to send feedback about documentation [156](#)

concurrent backup verification

- during a restore operation [104](#)
- with Frequent Recovery Point backup [104](#)
- workflow [104](#)

configuration

- of Exchange data [71](#)
- of gapless DAG backups [23](#)

configuration requirements

- Exchange Server [103](#)

Configuration wizard

- changing transaction log path [72](#)
- to move Exchange databases [71](#)
- using to change the location of log files [75](#)

control file

- exporting [76](#)
- XML schema [77](#)

controlling report auto-archiving [70](#)

controlling space used for reports [70](#)

**D**

DAG

- Activation Preference Number [27](#)
- backups using remote additional copy backup [22](#)
- gapless backup [22](#)
- member servers configuration limitation [103](#)
- remote additional copy backup feature [22](#)

DAG backup jobs

- planning [17](#)

DAG backups

- configuring using remote additional copy [23](#)

DAG member servers

- restoring database copies to [49](#)

database

- automatic unmounting [96](#)

changing the location using the Configuration wizard [72](#)

restoring

- Restore window [56](#)
- Restore Wizard [55](#)

restoring to a Frequent Recovery Point [57](#)

testing a restore,for up-to-the-minute restore operations [57](#)

verification

- throttling [88](#)

Database Availability Group

- See DAG

Database Availability Group (DAG)

- restore operations within [59](#)

Database Availability Groups

- planning backups for more than nine nodes [17](#)
- planning backups for nine nodes [17](#)

database backup

- using the Backup wizard [19](#)

database backup jobs

- planning for in DAGs [17](#)

database corruption

- identifying with utility programs [35](#)

database files

- viewing the full path to [73](#)

Database Filter

- using to display active databases [26](#)
- using to display databases to back up in a DAG [26](#)
- using to display passive databases [26](#)

database integrity

- verifying with utility programs [35](#)

database integrity checks

- in backup sets [35](#)
- utilities used to provide [35](#)

database migration

- disabling [85](#)
- enabling [84](#)

database path

- viewing [73](#)

database restore operation

- to a Frequent Recovery Point [57](#)

database restore operations

- using SnapManager [49](#)

database verification

- when verification can be performed [38](#)

database verification server

- considerations for selecting [85](#)

database verification workload

- deferred [36](#)
- load management [36](#)
- remote [36](#)
- remote verification on a SnapVault destination [36](#)
- verification throttling [36](#)

databases

- backing up for the first time [14](#)
- backing up with a schedule [16](#)
- maximum supported [103](#)
- overview of backing up [10](#)
- strategy for backing up [11](#)
- verifying the initial backup set [15](#)

databases per Exchange Server

- maximum supported [103](#)

databases per storage group

maximum supported [103](#)  
 dataset protection policy [73](#)  
 Deferred database verification [37](#)  
 delete-backup command  
   purpose [136](#)  
 disaster recovery  
   role of recovery database in [62](#)  
 documentation  
   how to receive automatic notification of changes to [156](#)  
   how to send feedback about [156](#)  
 drive letters [36](#)  
 dump command  
   using to archive LUNs [114](#)

## E

e-mail notification  
   configuring [100](#)  
 email notification support  
   for events [99](#)  
 error 0xC00413CA  
   during a restore operation [65](#)  
 errors  
   busy Snapshot copy [35](#)  
   during restore operations [65](#)  
 Eseutil  
   command-line tool to find physical-level corruption [35](#)  
 Eseutil consistency checker [90](#)  
 Eseutil utility [88](#)  
 event notification  
   Auto Notification Settings [99](#)  
   options for [99](#)  
 events  
   configuring notification [100](#)  
 Exchange Backup Agent  
   using to archive Exchange backup copies [119](#)  
 Exchange backup copies  
   archiving [119](#)  
 Exchange backups  
   automatic deletion [97](#)  
 Exchange data  
   configuration of [71](#)  
   moving back to local disk [110](#)  
 Exchange databases  
   backing up using the new-backup command [125](#)  
   moving [71](#)  
 Exchange mailbox  
   when to recover using SnapManager [49](#)  
 Exchange server  
   configuring the control file [76, 77](#)  
   default server [84](#)  
   restoring backups to another Exchange server [60](#)  
 Exchange Server  
   configuration requirements [103](#)  
 Exchange Server 2013 [76](#)  
 Exchange Server databases  
   restore operations for [59](#)  
 Exhaustive transaction log sequence checks [51](#)  
 explicit deletion  
   of backup or Snapshot LUN copies [31](#)  
   of multiple backups

of individual backups [31](#)  
 Export-config command [150](#)

## F

failed passive database copy  
   reseed using reseed-database command [145](#)  
   reseeding using Reseed Wizard [61](#)  
 features and benefits  
   overview of [8](#)  
 feedback  
   how to send comments about documentation [156](#)  
 flexible volumes  
   automatic expansion [93](#)  
 FlexVol volumes  
   automatic expansion [93](#)  
 fractional reserve  
   to prevent write operations from failing [93](#)  
 fractional space reservation  
   configuring policy [94](#)  
   monitoring [98](#)  
   policies  
    configuring [95](#)  
   status data [98](#)  
   to prevent failure of write operations [93](#)  
   viewing current data [98](#)  
 fractional space reservation policies  
   managing Exchange data [96](#)  
 fractional space reservation policy  
   default [94](#)  
   settings [94](#)  
 Fractional space reservation policy settings [94](#)  
 fractional space reserve [93](#)  
 Frequent Recovery Point  
   backup [44](#)  
   described [43](#)  
   description of [43](#)  
   operations [44](#)  
   reports [44](#)  
   restoration [44](#)  
   restoring data to a [43](#)  
   restoring databases to a [57](#)  
   when verification occurs [44](#)

## G

gapless backup  
   DAG [22](#)  
 gapless DAG backups  
   configuring using remote additional copy [23](#)  
   grouping [23](#)  
 get-backup command  
   purpose [138](#)  
 Get-JobStatus command [147](#)

## I

Import-config command [151](#)  
 individual backup  
   explicit deletion [32](#)  
 individual database  
   move to LUN [72](#)

- information
  - how to send feedback about improving documentation [156](#)
- integrity verification
  - destination SnapMirror volume
    - methods of performing [37](#)
- interactive upgrade
  - installing SnapManager [107](#)

**J**

- JobID parameter [148](#)

**L**

- LCR-enabled database
  - restoring using SnapManager [49](#)
- log database signature [51](#)
- log files
  - changing locations [75](#)
- LUN Clone Split Restore method
  - versus typical best practice
    - for only affected LUNs [51](#)
- LUN clone split status
  - verifying [52](#)
- LUN requirements [36](#)
- LUNs
  - write operations stopped [93](#)

**M**

- mailbox data
  - cleaning up after [67](#)
  - restoring with Single Mailbox Recovery [66](#)
    - restoring with Single Mailbox Recovery [66](#)
- mailboxes
  - restoring from archive [55](#)
- management groups
  - overview of [11](#)
- Microsoft Exchange cmdlets
  - usage guidelines [124](#)
- migration
  - of Exchange data [71](#)
- mount options [55](#)
- mount points [36](#)
- moving Exchange databases
  - local disk to LUN [71](#)
  - LUN-to-LUN
    - LUN-to-local-disk [71](#)
- multiple management groups
  - managing retention [29](#)

**N**

- naming
  - backup naming conventions [10](#)
- naming conventions
  - for Snapshot backup copies [10](#)
  - overview of [11](#)
- NDMP
  - using to archive LUNs [114](#)
- new-backup command
  - for specifying Remote Additional Copy Backup [24](#)
    - purpose [125](#)
  - Non-exhaustive transaction log sequence checks [51](#)
  - notification
    - XML schema [79](#)
  - NTFS data
    - back up example [116](#)

**P**

- passive database copies
  - displayed by Database Filter [26](#)
- point-in-time restore [50](#)
- policies
  - fractional space reservation [94](#)
- PowerShell command-line interface
  - SnapManager for Exchange usage guidelines [124](#)
- Priority parameter [148](#)
- problem launching scripts [122](#)
- product overview
  - features and benefits [8](#)
- production Exchange server [38](#)

**Q**

- quantity-based Snapshot copy deletion [28](#)

**R**

- Recovery Database
  - capabilities of [62](#)
  - configuration restrictions [103](#)
  - delete [64](#)
  - limitations of [62](#)
  - restoring databases to [49](#)
  - use of [62](#)
- Recovery Point Objective
  - meeting [17](#)
- Recovery Storage Group
  - configuration limits [103](#)
- registry keys
  - used for UTM retention [39](#)
- reinstalling
  - SnapManager [109](#)
- remote additional backup
  - control.xml example [25](#)
- remote additional backups
  - controlling retention of [25](#)
  - defining servers [25](#)
  - deleting [25](#)
  - managing [24](#)
  - new-backup parameters for specifying [24](#)
  - number retained [25](#)
- remote additional copy
  - using to configure gapless DAG backups [23](#)
- remote backup retention
  - modifying [23](#)
- remote database verification [37](#)
- remote verification
  - on SnapVault destination [37](#)
  - prerequisites [86](#)
  - working [86](#)

remote verification server [38](#)  
 repairing  
     SnapManager [109](#)  
 report auto-archiving  
     controlling [70](#)  
 report directory  
     changing [68](#)  
     locating in a DAG [69](#)  
     locating in a Windows cluster [69](#)  
     XML schema [80](#)  
 reports  
     DAG [68](#)  
     deleting [70](#)  
     on operations performed  
         on operation final status  
         on error messages [68](#)  
     printing [69](#)  
     viewing [69](#)  
 reseed  
     restore failed passive database copy  
         using reseed-database command [145](#)  
     restoring failed passive database copy  
         using Reseed Wizard [61](#)  
 Reseed Wizard  
     reseeded a failed passive database copy [61](#)  
 reseed-database command  
     purpose of [145](#)  
 restore  
     decisions [53, 54](#)  
     destination [54](#)  
     mailbox data [66](#)  
     point-in-time [50](#)  
     Snapshot copies [50](#)  
     type [54](#)  
     up-to-the-minute [50](#)  
     verification options [54](#)  
 restore data  
     to a Frequent Recovery Point [43](#)  
 restore database  
     from a backup copy [54](#)  
     types of restore operations using SnapManager [49](#)  
 restore error 0xC00413CA  
     troubleshooting [65](#)  
 restore operation  
     guidelines [53](#)  
     in a DAG environment  
         for a passive database copy [52](#)  
     passive database  
         restore operation for [52](#)  
     Restore window method [56](#)  
     Restore wizard method [55](#)  
     restoring mailbox data using Single Mailbox Recovery [66](#)  
     testing of up-to-the-minute restore operations [57](#)  
     to a Frequent Recovery Point [57](#)  
     using a Recovery Database [62](#)  
 restore operations  
     failures [65](#)  
     troubleshooting [65](#)  
     types of  
         guidelines for choosing [53](#)  
     within a Database Availability Group  
         from backup database copies  
             from backup database sets [59](#)  
 restore report  
     verification override entry [92](#)  
 restore Snapshot copies  
     explicit deletion [34](#)  
 restore storage group  
     from a backup copy [54](#)  
 restore time  
     decreasing [50](#)  
 restore-backup command  
     purpose [140](#)  
 restoring from unmanaged media [64](#)  
 retention management  
     for multiple management groups [29](#)  
 rstrsnap\_files  
     deleting [50](#)  
 Run Command After Operation feature  
     command arguments [120](#)  
     predefined variables [120](#)  
     using to archive SnapManager backups [115](#)  
 Run Command After Operation option  
     configuring [92](#)

**S**

scripts  
     launching from UNC paths [122](#)  
     to run SnapManager [124](#)  
 Server parameter [148](#)  
 service accounts  
     requirements for 7-Mode SnapVault [103](#)  
 SIF folders [30](#)  
     *See also* SnapInfo folders  
 silent mode  
     upgrading SnapManager unattended in, from the CLI [107](#)  
 Single Mailbox Recovery  
     role of in restoring mailbox data [66](#)  
     using to restore mailbox items [66](#)  
 SnapInfo folders  
     orphaned folder deletion [30](#)  
 SnapInfo Snapshot copies  
     explicit deletion [32](#)  
     SnapInfo Snapshot copies [32](#)  
 SnapManager  
     command-line reference [124](#)  
 SnapManager Configuration wizard  
     for moving Exchange data [71](#)  
 SnapManager control file  
     import or export [75](#)  
 SnapManager for Exchange  
     upgrading  
         prerequisites for upgrading [107](#)  
 SnapManager for Exchange PowerShell  
     launching [124](#)  
 SnapManager for Exchange PowerShell command-line interface  
     usage guidelines [124](#)  
 SnapManager for Microsoft Exchange Server  
     backup overview [10](#)  
     backup strategy [11](#)  
     features and benefits [8](#)  
     reinstalling [109](#)

repairing [109](#)  
 SnapManager restore  
     capabilities [49](#)  
 SnapManager scripts  
     to run SnapManager without using the GUI [124](#)  
 SnapManager snap-in [68](#), [76](#)  
 SnapManager upgrade  
     interactive mode [107](#)  
     interactively installing SnapManager [107](#)  
 SnapMirror relationship  
     XML schema [81](#)  
 Snapshot copies  
     automatic deletion [93](#)  
     automatic deletion of  
         requirements for deletion [30](#)  
     managing [30](#)  
     naming conventions [10](#)  
 Snapshot copies per volume [27](#)  
 Snapshot copy  
     access method [87](#)  
 Snapshot copy busy [35](#)  
 Snapshot copy names  
     hard coding [115](#)  
 SnapVault  
     service account requirements for 7-Mode [103](#)  
 SnapVault destination  
     remote verification [37](#)  
 software uninstallation  
     uninstallation script  
         using to uninstall software [111](#)  
     uninstalling unattended  
         using the CLI [111](#)  
 SourceBackupServer parameter [148](#)  
 space reservation  
     fractional reserves [93](#)  
 space used by reports  
     controlling [70](#)  
 storage groups  
     maximum per cluster [103](#)  
     maximum supported [103](#)  
 storage layout  
     XML schema [77](#)  
 suggestions  
     how to send feedback about documentation [156](#)  
 syslog entries  
     for events [99](#)  
 syslog event logging  
     configuring [100](#)

## T

test restore  
     verifying current logs and verifying the database [57](#)  
 throttling report  
     entries [90](#)  
 transaction log  
     changing log path [72](#)  
 transaction log backup  
     used by Frequent Recovery Point restore [43](#)  
 transaction log backups  
     using a schedule [17](#)  
 transaction log files  
     managing with the Configuration wizard [75](#)

Transaction log sequence verification [51](#)  
 transaction logs  
     backing up with a schedule [17](#)  
     backup options [39](#)  
     changing locations [75](#)  
     managing with the Configuration wizard [75](#)  
     moving to LUN [74](#)  
     truncated [43](#)  
     viewing full path [75](#)  
 transport database [74](#)  
 troubleshooting  
     busy Snapshot copy [35](#)  
     integrity verification [47](#)  
     problem deleting backup copies [35](#)  
     reinstalling SnapManager [109](#)  
     repairing SnapManager [109](#)  
     restore operations [65](#)  
 Twitter  
     how to receive automatic notification of  
         documentation changes [156](#)

## U

UNC path [122](#)  
 uninstallation  
     interactive [110](#)  
 unmanaged media  
     restoring from [64](#)  
 unmounting  
     back up copy [67](#)  
 up-to-the-minute restore [50](#)  
 up-to-the-minute restore ability [31](#)  
 up-to-the-minute retention  
     policies [39](#)  
 upgrades  
     performing in unattended mode from the CLI [107](#)  
 upgrades for SnapManager  
     prerequisites for upgrading [107](#)  
 UTM retention  
     policies [39](#)  
     registry keys for [39](#)  
     scenario 1 [41](#)  
     scenario 2 [41](#)  
     scenario 3 [42](#)  
     scenario 4 [42](#)  
     triggering [39](#)  
 UTM-related switches  
     valid input for [39](#)

## V

verification  
     accessing Snapshot copies [87](#)  
     override [91](#)  
 performance  
     impact [90](#)  
 sleep interval  
     calculating [89](#)  
 throttling  
     configuring [90](#)  
     options [88](#)  
     sleep interval [89](#)

- working [88](#)
- viewing or changing the server [87](#)
- XML schema [80](#)
- verification override
  - report entry [92](#)
- verification server
  - configuring [85](#)
- verification status reporting [38](#)
- Verification throttling [37](#)
- verify-backup command
  - purpose [132](#)
  - verifying multiple backup sets [36](#)
  - Verifying multiple backup sets [36](#)
  - volumes
    - fractional reserve [93](#)
- W**
  - workload management
    - methods of managing workloads [36](#)