



SnapManager[®] 7.2.2 for Microsoft[®] SQL Server[®]

Installation and Setup Guide

For Clustered Data ONTAP[®]

June 2018 | 215-12168_B0
doccomments@netapp.com

 **NetApp[®]**


Contents

Product overview	4
Deployment workflow	6
Preparing for deployment	7
Storage layout requirements	7
SnapManager dedicated servers	10
SnapManager licensing	10
Supported configurations	11
Supported storage types	12
Windows host requirements	13
Service account requirements	15
Installing SnapManager	16
Installing SnapManager interactively	16
Installing SnapManager from the command line	17
Migrating databases to NetApp storage	19
Connecting SnapManager to SQL Server instances	19
Migrating databases and configuring SnapManager for SQL Server instances	20
Preparing storage systems for SnapMirror and SnapVault replication	25
Understanding the differences between SnapMirror and SnapVault	25
Preparing storage systems for SnapMirror replication	26
Preparing storage systems for SnapVault replication	27
Backing up and verifying your databases	30
SnapManager backup overview	30
Defining a backup strategy	30
Backing up your databases for the first time	33
Verifying the initial backup set	34
Scheduling recurring backups	35
Scheduling recurring transaction log backups	36
Scheduling recurring backup set verifications	37
Where to go next	39
Copyright information	40
Trademark information	41
How to send comments about documentation and receive update notifications	42
Index	43

Product overview

SnapManager for Microsoft SQL Server is a host-side component of the NetApp integrated storage solution for SQL Server, offering application-aware primary Snapshot copies of SQL 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:



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 SQL Server AlwaysOn 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 SQL 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 SQL Server database backups, configure policy-based backup retention, set up point-in-time and up-to-the-minute restore operations and proactively monitor your SQL Server environment with periodic email alerts. PowerShell cmdlets are available for easy scripting of backup and restore operations.

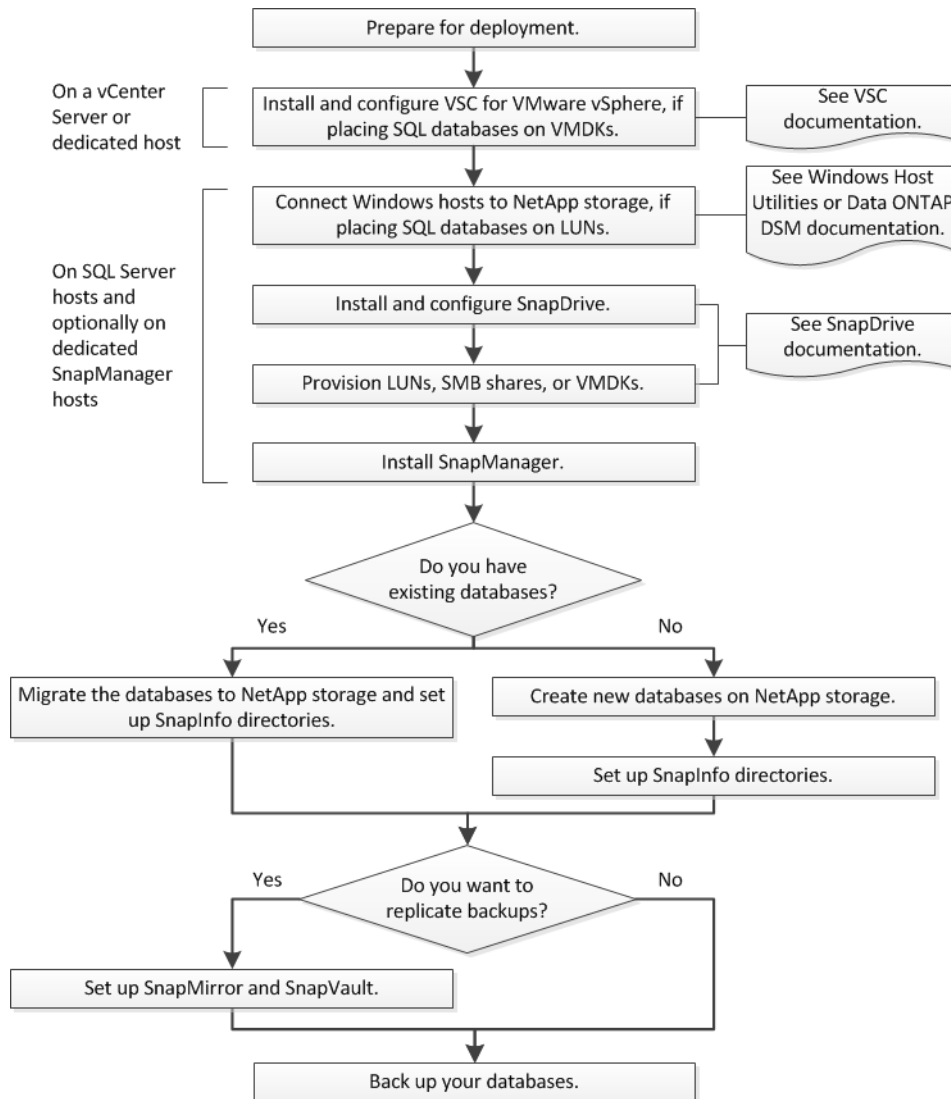
In addition to these major features, SnapManager offers the following:

- Integrated FlexClone software enables you to create space-efficient point-in-time copies of production databases for testing or data extraction (FlexClone license required)
- Simplified migration of existing databases to NetApp storage with an easy-to-use Configuration wizard
- Nondisruptive, automated backup verification

- Fast reseeding of databases in an AlwaysOn cluster
- Federated database backup of multiple SQL Server instances and databases
- Support for backup of LUNs, SMB shares, and VMDKs
- Support for physical and virtualized infrastructures
- Support for iSCSI, Fibre Channel, FCoE, SMB 3.0 (CIFS), RDM, and VMDK over NFS and VMFS

Deployment workflow

Before you can create backups with SnapManager, you need to install the SnapDrive for Windows and SnapManager software, and provision NetApp storage. You can then migrate your databases to the storage system or create new databases in the system.



Preparing for deployment

Before you deploy SnapManager, you need to determine your storage layout, choose a SnapManager configuration, verify that you have the required licenses, and make sure that your Windows hosts meet the minimum requirements.

Steps

1. Plan how to lay out your databases on NetApp storage.
2. Decide whether you are going to use a SnapManager dedicated server for administration or verification.
3. Verify that you have the required licenses.
4. Verify SnapManager support for your configuration and storage type.
5. Verify that your Windows hosts meet SnapManager requirements.
6. Set minimum permissions for SQL Server and SnapManager service accounts.

Related references

[Storage layout requirements](#) on page 7

[SnapManager dedicated servers](#) on page 10

[SnapManager licensing](#) on page 10

[Supported configurations](#) on page 11

[Supported storage types](#) on page 12

[Windows host requirements](#) on page 13

[Service account requirements](#) on page 15

Storage layout requirements

A well-designed storage layout ensures that SnapManager can properly back up your databases and that you can meet your recovery objectives.

The following sections define the storage layout requirements and restrictions for LUNs, VMDKs, and SMB shares. You should take several considerations into account when defining your storage layout, including the size of the database, its rate of change, and the frequency with which you perform backups. For more information about these considerations, see [NetApp Technical Report 4353: Best Practices Guide for Microsoft SQL Server and SnapManager 7.1 for SQL Server with Clustered Data ONTAP](#).

LUN and VMDK requirements

You need dedicated LUNs or VMDKs for the following:

1. Master, model, and msdb system databases
2. Tempdb
3. User database files (.mdf and .ndf)
4. User database transaction log files (.ldf)
5. The SnapInfo directory

The SnapInfo directory stores information about backed up files. You create one or more SnapInfo directories when you migrate databases to NetApp storage.

Snapshot copies are volume-wide, so each LUN should be in a dedicated volume and each VMDK should be in a dedicated datastore and volume.

A LUN or VMDK can contain user database files (.mdf and .ndf) for a single database or for multiple databases. Which model you choose, or whether you choose a combination of the models, will depend primarily on the number of databases you are backing up. While a single database per LUN or VMDK offers the simplest mapping scheme, it will almost always be impractical when you are working with hundreds of databases.

You also need to consider backup and restore performance:

- A single database per LUN or VMDK can result in faster restore times but slower backup times because you are backing up more volumes.
- Multiple databases per LUN or VMDK can result in longer restore times if you are not restoring all the databases in the volume.


This is because SnapManager restores individual databases by mounting the Snapshot copy and then performing a file-level copy.

Follow the same layout model for transaction log files (.ldf). No matter which model you choose, you should place transaction logs and database files on different LUNs or VMDKs because transaction logs are usually backed up more frequently than database files.


For more information about considerations for transaction log files, see [NetApp Technical Report 4353: Best Practices Guide for Microsoft SQL Server and SnapManager 7.1 for SQL Server with Clustered Data ONTAP](#).

LUN and VMDK sample layouts


The following graphic shows how you might lay out storage with a single database per LUN:



The following graphic shows how you might lay out storage with multiple databases per LUN:



The following graphic shows multiple databases per VMDK:



LUN and VMDK restrictions

- You cannot store database files on the same LUN or VMDK as the SnapInfo directory.
- You cannot store database files on a LUN that hosts NTFS volume mount points.
- You cannot store database items of any type on a LUN or VMDK that hosts the SQL Server.
- You cannot store database files, transaction logs, or the SnapInfo directory on a SAN boot LUN or a LUN containing any other directories or files (including system paging files).
- If a database contains multiple filegroups or secondary (.ndf) files, you can spread the database's files across two or more LUNs or VMDKs, provided those LUNs or VMDKs do not contain files for any other databases.

Note: In exceptional cases, you can override these restrictions by enabling the unrestricted database layout option. For more information, see the [SnapManager 7.2 for Microsoft SQL Server Administration Guide](#).

- For configurations where database files, transaction logs, or SnapInfo directories reside on VMDKs, those VMDKs cannot share the same datastore (VMFS or NFS) as VMDKs that store system partitions for any virtual machines.
- For configurations where database files, transaction logs, or SnapInfo directories reside on VMDKs, NetApp's Virtual Storage Console for VMware vSphere must not back up the datastores on which those VMDKs reside.

SMB share requirements

There are no requirements or restrictions for databases on SMB shares. You can place all SQL databases and the SnapInfo directory on the same SMB share or on multiple SMB shares.

SnapManager dedicated servers

Ordinarily, you install SnapManager on each Windows host running SQL Server software. From this *base configuration*, you can administer SnapManager locally or remotely. Depending on your needs, you might also want to install SnapManager on a dedicated administration or verification server.

- An *administration server* lets you manage SnapManager remotely from a host of your choosing. You might want to avoid using a primary SQL Server host for SnapManager administration, or it might simply be more convenient to use a dedicated server.
- A *verification server* lets you offload backup set verification from a primary SQL Server host. SnapManager's optional backup set verification feature uses the Microsoft SQL Server Database Consistency Checker (DBCC) to verify the page-level integrity of databases. Because verification is a CPU-intensive operation that can degrade SQL Server performance, it is a best practice to run the utility on a dedicated server. The verification server must have iSCSI or FC connectivity with the storage system.

Tip: You can configure SnapManager to perform verification during or after backup. You can also configure it to verify the mirror or vault copy on the target storage system rather than the primary copy on the source system.

You can administer SnapManager from a verification server if it is more convenient than configuring a separate administration server.

SnapManager licensing

A SnapManager license and several storage system licenses are required to enable SnapManager operations. The SnapManager license is available in two licensing models: *per-server licensing*, in which the SnapManager license resides on each SQL Server host, and *per-storage system licensing*, in which the SnapManager license resides on the storage system.

SnapManager license requirements are as follows:

License	Description	Where required
SnapManager per-server	A host-side license for a specific SQL Server host. Licenses are required only for SQL Server production hosts on which SnapManager is installed and for the optional verification server. No SnapManager license is required for the storage system or for the optional administration server.	On the SnapManager host. A SnapManager suite license is not required on source and destination storage systems when using per-server licensing.
SnapManager per-storage system (SnapManager suite)	A storage-side license that supports any number of SQL Server hosts. Required only if you are not using a per-server license on the SnapManager host. Note: Trial licenses are available for per-storage system licensing only.	On source and destination storage systems.

License	Description	Where required
SnapRestore	A required license that enables SnapManager to restore and verify backup sets. Restores include file level restores.	On source storage systems. Required on SnapMirror destination systems to perform remote verification. Required on SnapVault destination systems to perform remote verification and to restore from a backup.
FlexClone	An optional license for cloning databases.	On source storage systems. Required on SnapVault destination systems when creating clones from a backup.
SnapMirror	An optional license for mirroring backup sets to a destination storage system.	On source and destination storage systems.
SnapVault	An optional license for archiving backup sets to a destination storage system.	On source and destination storage systems.
Protocols	The following licenses are required: <ul style="list-style-type: none"> • For LUNs, the iSCSI or FC license • For SMB shares, the CIFS license • For NFS-type VMDKs, the NFS license • For VMFS-type VMDKs, the iSCSI or FC license 	On source storage systems. Required on SnapMirror destination systems to serve data if a source volume is unavailable.

Supported configurations

You can use the NetApp Interoperability Matrix to verify SnapManager support for your configuration before you install or upgrade SnapManager.

The following table shows the currently supported software configurations:

Windows Server	SQL Server	SnapDrive for Windows
2016 (Standard and Datacenter)	<ul style="list-style-type: none"> • 2016 (Standard and Enterprise) • 2014 SP1 (Standard and Enterprise) • 2014 (Standard and Enterprise) • 2012 SP3 (Standard and Enterprise) • 2012 SP2 (Standard and Enterprise) • 2012 SP1 (Standard and Enterprise) • 2012 (Standard and Enterprise) 	Bundled

Windows Server	SQL Server	SnapDrive for Windows
2012 R2 (Standard and Datacenter)	<ul style="list-style-type: none"> • 2014 (Standard and Enterprise) • 2012 SP2 (Standard and Enterprise) • 2012 SP1 (Standard and Enterprise) • 2012 (Standard and Enterprise) • 2008 R2 SP3 (Enterprise and Datacenter) • 2008 R2 SP2 (Enterprise and Datacenter) • 2008 SP3 (Standard, Enterprise, and Datacenter) 	Bundled
2012 (Standard and Datacenter)	<ul style="list-style-type: none"> • 2014 (Standard and Enterprise) • 2012 SP2 (Standard and Enterprise) • 2012 SP1 (Standard and Enterprise) • 2012 (Standard and Enterprise) • 2008 R2 SP3 (Enterprise and Datacenter) • 2008 R2 SP2 (Enterprise and Datacenter) • 2008 SP3 (Standard, Enterprise, and Datacenter) • 2005 SP4 (Standard, Enterprise, and Datacenter) 	Bundled
2008 R2 SP1 (Standard, Enterprise, and Datacenter)	<ul style="list-style-type: none"> • 2014 (Standard and Enterprise) • 2012 SP2 (Standard and Enterprise) • 2012 SP1 (Standard and Enterprise) • 2012 (Standard and Enterprise) • 2008 R2 SP3 (Enterprise and Datacenter) • 2008 R2 SP2 (Enterprise and Datacenter) • 2008 SP3 (Standard, Enterprise, and Datacenter) • 2005 SP4 (Standard, Enterprise, and Datacenter) 	Bundled

Related information

[NetApp Interoperability Matrix Tool](#)

Supported storage types

SnapManager supports a wide range of storage types on both physical and virtual machines. Verify support for your storage type before you install or upgrade SnapManager.

Machine	Storage type
Physical server	<ul style="list-style-type: none"> • FC-connected LUNs • iSCSI-connected LUNs • SMB3 shares residing on a Storage Virtual Machine (SVM) (clustered Data ONTAP 8.2 and later)

Machine	Storage type
VMware VM	<ul style="list-style-type: none"> • RDM LUNs connected via FC HBA • RDM LUNs connected via iSCSI HBA • iSCSI LUNs connected directly to the guest system by the iSCSI initiator • VMDKs on VMFS or NFS datastores • A guest system connected to SMB3 shares residing on a Storage Virtual Machine (SVM) (clustered Data ONTAP 8.2 and later)
Hyper-V VM	<ul style="list-style-type: none"> • Passthrough LUNs connected via FC HBA • Virtual FC (vFC) LUNs connected via virtual Fibre Switch • Passthrough LUNs connected via iSCSI HBA • iSCSI LUNs connected directly to the guest system by the iSCSI initiator • A guest system connected to SMB3 shares residing on a Storage Virtual Machine (SVM) (clustered Data ONTAP 8.2 and later)

Windows host requirements

Windows hosts must meet the requirements for the base SnapManager configuration and for each of the optional dedicated SnapManager servers.

Important notes

- You can install SnapManager on any combination of physical machines or virtual machines. If you install SnapManager on virtual machines, the verification server must also be installed on a virtual machine.
- You must use the same version of SnapManager on all of the hosts.
- When you schedule database backups, you can select the individual databases in an AlwaysOn Availability Group cluster or you can select the cluster itself, which backs up all of the nodes in the cluster.
If you want to create backups at the node level, you must install SnapManager and SnapDrive for Windows only on the nodes that you plan to back up. If you want to create backups at the Availability Group level, you must install SnapManager and SnapDrive for Windows on all of the nodes.
- Except in VMDK configurations, you can connect the remote verification server to the storage system by using a different protocol from the protocol that you used to connect to the base configuration.

Host requirements per server type

The following table lists the host requirements for the base SnapManager configuration and for the optional administration and verification servers:

Requirement	Base	Administration	Verification	Notes
SQL Server	Yes	No	Yes	The version of SQL Server that is running on the verification server must be the same or higher than the version that is running on the base configuration.
.NET Framework 4.0, 4.5.x, 4.6 or 4.7	Yes	Yes	Yes	
Windows Failover Cluster Automation Server	Yes	Yes	Yes	Installing SnapManager on a Windows Server 2012 Failover Cluster node.
Windows PowerShell 3.0	Yes	Yes	Yes	
SnapDrive for Windows	Yes	Yes	Yes	Upgrading or installing a SnapDrive for Windows version that is supported by your version of Data ONTAP.


Port and connection requirements for Windows Firewall implementations

If you have enabled Windows Firewall on your hosts, TCP port 808 must be available (both inbound and outbound) for SnapManager communications, including communications with the optional verification servers and administration servers.

Additional requirements for VMDK support

For VMDK support, NetApp Virtual Storage Console (VSC) 4.2.2 for VMware vSphere must be installed on a Windows host in your network. You can install VSC for VMware vSphere on the vCenter Server host or on a dedicated host.

The following illustration shows VSC running in a SnapManager deployment:



Related references

[SnapManager dedicated servers](#) on page 10

Service account requirements

A service account is a user account created explicitly to provide a security context for services running on Windows Server. You must specify a service account when you install or update a service. Before working with SnapManager, you must ensure that both the SQL Server and SnapManager service accounts have the required permissions on the Windows host.

SQL Server service account requirements

SnapManager initiates some SQL Server operations that require access to the file system and Windows registry. For that reason, the SQL Server service account should be a domain user account with permission to write to the file system and registry.

SnapManager service account requirements

The SnapManager service account must have administrator privileges on the SQL Server host. Requirements for special cases are described in the following table:

To...	The SnapManager service account...
Use SnapInfo directory on a CIFS share	Must be a domain user account.
Use Windows Authentication to connect with SQL Servers	Must have the sysadmin role in each SQL Server instance installed on the host.
Run SnapManager from a Group Managed Service Account (gMSA)	Must have rights on the SQL Server host.
Use SnapManager with Microsoft Message Queuing in workgroup mode	Must be a local user account. The service account cannot be a domain user account.

Installing SnapManager

Ordinarily, you install SnapManager on each Windows host running SQL Server software. Depending on your needs, you might want to install SnapManager on a dedicated administration or verification server. You can use an interactive wizard or the command line to install the product.

Before you begin

- You should have backed up your SQL databases.
- SnapDrive for Windows must be installed.

About this task

When you schedule database backups, you can select the individual databases in an AlwaysOn Availability Group cluster or you can select the cluster itself, which backs up all nodes in the cluster. If you want to back up at the node level, you need to install SnapManager and SnapDrive for Windows only on the nodes you plan to back up. If you want to back up at the Availability Group level, you must install SnapManager and SnapDrive for Windows on all nodes.

During installation, SnapManager configures the SQL Browser service to start automatically and installs the following Windows components:

- SQL 2005 backward compatibility components
- SQL Server 2012, 2014, or 2016 CLR types
- SQL Server 2012 R2, 2014, or 2016 Management Objects
- Microsoft Visual C++ 2012 Redistributable Package (x64)

Related tasks

[Installing SnapManager interactively](#) on page 16

[Installing SnapManager from the command line](#) on page 17


Installing SnapManager interactively

You can use the SnapManager installation wizard to interactively install SnapManager on a Windows host.

Steps

1. Download the SnapManager for SQL Server software from the NetApp Support Site.
[NetApp Downloads: Software](#)
2. Double-click the downloaded .exe file.
3. Complete the steps in the SnapManager installation wizard to install SnapManager.

Most of the fields in the wizard are self-explanatory. The following table describes fields for which you might need guidance:

Field	Description
Account	<p>The user account that Windows uses to run SnapManager. This <i>SnapManager service account</i> must have specific permissions on the Windows host and the SQL Server. For details, see Service account requirements on page 15.</p> <p>Specify the account name by using one of the following formats:</p> <ul style="list-style-type: none"> <i>DomainName\UserName</i> <i>UserName@DomainName</i> <p>Example:</p> 
License Type	<p>The license type that you purchased, which is either per-server licensing or per-storage system licensing. For details, see SnapManager licensing on page 10. Leave the License Key field blank if you are using per-server licensing and would prefer to specify the license key in the SnapManager console.</p>
Password	<p>The password for the specified account. Leave the password blank if you entered a group Managed Service Account in the Account field.</p>

Installing SnapManager from the command line

You can run the SnapManager installation program unattended, in silent mode, from the Windows command line.

Steps

1. Download the SnapManager for SQL Server installer from the NetApp Support Site.
[NetApp Downloads: Software](#)
2. From a Windows command prompt on the local host, change to the directory where you downloaded the product installer.
3. Enter the following command at the command prompt:

```
installer.exe /s /v"/qn SILENT_MODE=1 [USERNAME=UserName]
[COMPANYNAME=CompanyName] [ISX_SERIALNUM=LicenseKey]
[INSTALLDIR=InstallDirectory] SVCUSERNAME=Domain\UserName
SVCUSERPASSWORD=Password SVCCONFIRMUSERPASSWORD=Password [/L*v DirPath
\LogFileName]"
```

Enter the following for each variable:

Variable	Description
<i>installer</i>	The name of the .exe file
<i>UserName</i>	The name of the product administrator. If not specified, SnapManager retrieves the default value from the Windows registry.
<i>CompanyName</i>	The name of your company. If not specified, SnapManager retrieves the default value from the Windows registry.

Variable	Description
<i>LicenseKey</i>	The per-server license key. Leave this field blank if you are using per-storage system licensing, or if you would prefer to specify the per-server license key in the SnapManager console. For details, see SnapManager licensing on page 10.
<i>InstallDirectory</i>	An alternate installation directory. If not specified, SnapManager uses the default directory: C:\Program Files\NetApp\SnapManager for SQL Server
<i>Domain\UserName</i>	The user account that Windows uses to run SnapManager. This <i>SnapManager service account</i> must have specific permissions on the Windows host and the SQL Server. For details, see Service account requirements on page 15.
<i>Password</i>	The password for the specified user account. Leave the password blank if you entered a group Managed Service Account in the Account field.
<i>DirPath\LogFileName</i>	The location and name of an installation log file, which is useful for troubleshooting. The asterisk (*) specifies that all installation information (such as status messages, nonfatal warnings, and error messages) should be logged.

Example

```
"SMSQL7.2_x64.exe" /s /v"/qn SILENT_MODE=1 ISX_SERIALNUM=123
SVCUSERNAME=MKTG2\Administrator SVCUSERPASSWORD=examplepwd!
SVCCONFIRMUSERPASSWORD=examplepwd! /L*V C:\SMSQL_Install.log"
```

Migrating databases to NetApp storage

After you have provisioned NetApp storage with SnapDrive for Windows, you can migrate your databases to the storage system or create new databases in the system. In either case, you use the Configuration wizard to create the SnapInfo directory that SnapManager uses to store information about backed-up files.

Before you begin

- The SQL database property AutoClose must be set to **FALSE**.
- For an AlwaysOn Availability Group, the readable secondary value must be set to **Yes** for all replicas used with SnapManager.
- Database names must not include any of the following characters: \ / : * ? " < > | [,]
- Database names must not end with spaces unless the following DWORD (32-bit) registry was set to 1 and the SnapManager Service was restarted:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Network Appliance\SnapManager for SQL Server
\Server\HaveDBWithTrailingSpaces
```

Related tasks

[Connecting SnapManager to SQL Server instances](#) on page 19

[Migrating databases and configuring SnapManager for SQL Server instances](#) on page 20

Connecting SnapManager to SQL Server instances


Before you can migrate your databases, you need to connect SnapManager to your SQL Server instances.

Steps

1. From the Windows **Start** menu, click **SnapManager for SQL Server**.

The SnapManager console appears.

2. In the **Console Root** tree, click **SnapManager for SQL Server**.



The Add SQL Instance to be managed dialog box appears.

3. To choose an SQL instance, select one from the drop-down list, type the name in the text box, or click **Browse**.

Note the following special cases:

If you have...	Do this...
A server that does not have a default SQL Server instance	Specify one of the named instances (Server\Instance) instead of the server name. SnapManager adds all of the instances on the server when you specify one of the named instances.

If you have...	Do this...
An Availability Group	Specify each Availability Group Replica that you plan to back up.

- Under **Login Details**, choose the authentication method that you want SnapManager to use to connect to the SQL Server:

Use Windows authentication

SnapManager connects to the SQL Server using the Windows account under which SnapManager runs (the SnapManager service account). This is the most common method.

Use SQL Server authentication

SnapManager connects to the SQL Server using an account defined on the SQL Server. That account must have sysadmin server role privileges on the SQL Server instance.

SnapManager uses the selected authentication method for all SQL Servers and SQL Server instances in a cluster, so each server must be set up for that authentication method.

- Click **Add**.

The Add SQL Instance to be managed dialog box closes and SnapManager launches the Configuration wizard.

After you finish

Use the Configuration wizard to configure SnapManager for the SQL Server instance and to migrate the SQL Server databases to NetApp storage.

Migrating databases and configuring SnapManager for SQL Server instances

Before you can back up your databases using SnapManager, you need to run the SnapManager Configuration wizard for each SQL Server instance. You use the Configuration wizard to migrate databases to NetApp storage and to configure SnapManager for the SQL Server instances.

Before you begin

You must have shut down any applications that are accessing the databases that you want to migrate. Database migration fails if an application accesses a database.

About this task

Attention: SnapManager takes databases offline when it migrates them. SnapManager takes the *entire* SQL Server offline when it migrates the system databases.


You use the SnapManager Configuration wizard to select a verification server, migrate data files and log files to NetApp storage, map those files to their respective SnapInfo directories, and configure automatic event notification.

When migrating databases, SnapManager ensures that the files are placed in locations that meet SnapManager configuration requirements. If you use a separate tool to migrate databases, run the Configuration wizard to ensure that the files are in correct locations. Incorrectly located files can impair SnapManager operations.

Even if you create new databases directly on NetApp storage, you need to run the Configuration wizard to create a mapping between those databases and the SnapInfo directory.

Steps

- If the **Configuration** wizard is not open, click **Configuration Wizard** in the **Actions** menu.



2. On the **Start** page, click **Next**.


This page includes an option to use a *control file*. A control file contains configuration details about an SQL Server instance. You might use this option at another time to export and then import a configuration. For information about using control files, see the [SnapManager 7.2 for Microsoft SQL Server Administration Guide](#).

3. On the **Verification Settings** page, define how SnapManager should verify backup copies:


For this field...	Do this...
Verification Server	For optimal performance, choose a remote verification server, which offloads work from the SQL production server.
SQL Server Connection	<p>Choose an authentication method for SnapManager to connect to the SQL Server on the verification server during backup verification:</p> <p>Use Windows authentication</p> <p>SnapManager connects to the SQL Server using the Windows account under which SnapManager runs (the SnapManager service account). This is the most common method.</p> <p>Use SQL Server authentication</p> <p>SnapManager connects to the SQL Server using an account defined on the SQL Server. The account must have sysadmin server role privileges on the SQL Server instance.</p>
Access a mounted LUN in snapshot	<p>Keep the default option for mounting Snapshot copies to an empty NTFS directory.</p> <p>SnapManager mounts Snapshot copies to the verification server when it verifies backup copies. Using an empty NTFS directory is typically better than assigning drive letters because the verification server can run out of drive letters if there are more backup copies than available drive letters.</p> <p>For a Windows Failover Cluster, the mount point directory must be a shared disk.</p>

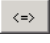
4. On the **Database Selection** page, assign your database files and transaction logs to LUNs, SMB shares, or VMDKs and then view the results of your selections:

- a. In the **Database Selection** pane, select a database file or transaction log:



- b. In the **Disk Selection** pane, select a LUN, SMB share, or VMDK:



- c. Click the  button.
- d. Repeat substeps a to c for each database's files and logs.
- e. In the **Database Location Results** pane, review the results of your selections:

The screenshot shows the 'Database Location Results' pane with a table. The table has three columns: 'Database', 'From', and 'To'. The rows are: 'QATX86W2K8\SQL2014', 'QATX86W2K8', 'db1_new' (with 'Invalid/Local. Data...' in the 'From' column), 'PRIMARY', and 'db1_new' (with 'C' in the 'From' column and 'F' in the 'To' column). The last row is highlighted in blue.

Database	From	To
QATX86W2K8\SQL2014		
QATX86W2K8		
db1_new	Invalid/Local. Data...	
PRIMARY		
db1_new	C	F

Tip: To change your selections, click **Undo All** or select a database and click **Reconfigure**.

5. On the **SnapInfo settings** page, choose whether you want a single SnapInfo directory for your databases or multiple SnapInfo directories.

A single SnapInfo directory works well for many configurations. Multiple SnapInfo directories might be needed if you have some databases on LUNs and some databases on SMB shares. In that case, the databases on LUNs need LUNs for the corresponding SnapInfo directories and the databases on SMB shares need SMB shares for the corresponding SnapInfo directories.

6. On the remaining **SnapInfo Settings** pages, specify where you want to place the SnapInfo directory or directories.
7. On the **Setup SnapManager Share** page, specify a share to ensure that backups of transaction logs are available to all replicas within an Availability Group.

SnapManager uses the transaction log backups for up-to-the-minute restore operations, database reseeding, and for using clones as replicas.

8. On the **Database Migration Settings** page, choose settings for physical database verification, database deletion, and updates to database table statistics:

For this field...	Do this...
Run DBCC physical integrity verification before/after migration	Clear these fields. The integrity verification runs on the production database and is resource intensive, which can negatively impact the performance of your database. You should run verification on the initial backup set instead.
Delete Copy of Migrated Database at Original Location	Clear this field. You should always delete the database yourself, in case there are problems during the migration.
Run UPDATE STATISTICS on tables before detaching the databases	Leave this field selected.

The following image shows the Database Migration Settings page with the recommended settings selected:

Run DBCC physical integrity verification before migration

Run DBCC physical integrity verification after migration

Select whether you want to delete copy of the migrated database at the original location after the migration completes successfully.

Delete Copy of Migrated Database at Original Location

Select whether you want to update statistics on all tables before detaching databases.

Run UPDATE STATISTICS on tables before detaching the databases

9. On the **iSCSI Initiator Information** page, choose whether to set the iSCSI service as a dependency and click **Next**:

- If you use iSCSI, you should keep the option selected.
- If you do not use iSCSI, clear the option.

Setting the iSCSI service as a dependency for all SQL Server services helps protect SQL data if there is a problem with iSCSI.

10. In the **E-Mail Notification Settings** page, configure settings for email notifications, event logging, and AutoSupport notifications.

Most of the fields on this page are self-explanatory. The following table describes fields for which you might need guidance:

Field	Description
Send e-mail notification	Enables email notifications to the specified address about the success or failure of SnapManager operations. If you select this field, click Advanced to tune the notification settings—for example, to receive notifications only when operations fail.
Log SnapManager events to storage system syslog	If AutoSupport is enabled on the storage system, posts SnapManager events to the storage system's event log. Technical support can use this information to troubleshoot issues.
Send AutoSupport notification	If AutoSupport is enabled on the storage system, enables email notifications to technical support about SnapManager events or storage system problems that might occur.
On failure only	Limits the SnapManager events that are posted to the storage system event log and sent through AutoSupport to failure events only.

11. In the **Monitoring and Reporting Settings** page, choose whether you want to receive email notifications that contain the status of backup, verification, and clone operations during the specified time frame.

To receive these email notifications, you must have enabled email notifications on the previous page.

12. In the **Finish** page, review the settings and click **Finish**.

13. In the **Configuration Status** dialog box, click **Start Now**.

SnapManager migrates the databases and updates your SnapManager configuration. You can view details of the operation in the Configuration Report.

14. After the operation completes, click **OK** and then click **Close**.

After you finish

You can rerun the Configuration wizard at any time to make changes to your database configurations.

When you add new databases, you should run the Configuration wizard to ensure that the databases are stored in valid locations and to create a mapping between those databases and their respective SnapInfo directories.

Preparing storage systems for SnapMirror and SnapVault replication

You can use SnapManager with ONTAP SnapMirror technology to create mirror copies of backup sets on another volume, and with ONTAP SnapVault technology to perform disk-to-disk backup replication for standards compliance and other governance-related purposes. Before you perform these tasks, you must configure a *data-protection relationship* between the source and destination volumes and *initialize* the relationship.

A data protection relationship replicates data on primary storage (the source volume) to secondary storage (the destination volume). When you initialize the relationship, ONTAP transfers the data blocks referenced on the source volume to the destination volume.

Note: If you are using SnapDrive 7.1 for Windows, you cannot configure both a SnapMirror relationship and a SnapVault relationship on the same ONTAP source volume. You must configure these relationships on different source volumes. This limitation does not exist starting with SnapDrive 7.1.1 for Windows.

SnapManager does not support cascade relationships between SnapMirror and SnapVault volumes (**Primary > Mirror > Vault**). Use fanout relationships only (**Primary > Mirror, Primary > Vault**).

Related tasks

[Understanding the differences between SnapMirror and SnapVault](#) on page 25

[Preparing storage systems for SnapMirror replication](#) on page 26

[Preparing storage systems for SnapVault replication](#) on page 27

Understanding the differences between SnapMirror and SnapVault

SnapMirror is disaster recovery technology, designed for failover from primary storage to secondary storage at a geographically remote site. SnapVault is disk-to-disk backup replication technology, designed for standards compliance and other governance-related purposes.

These objectives account for the different balance each technology strikes between the goals of backup currency and backup retention:

- SnapMirror stores *only* the Snapshot copies that reside in primary storage, because, in the event of a disaster, you need to be able to fail over to the most recent version of primary data you know to be good.
Your organization, for example, might mirror hourly copies of production data over a ten-day span. As the failover use case implies, the equipment on the secondary system must be equivalent or nearly equivalent to the equipment on the primary system to serve data efficiently from mirrored storage.
- SnapVault, in contrast, stores Snapshot copies *whether or not* they currently reside in primary storage, because, in the event of an audit, access to historical data is likely to be as important as access to current data.
You might want to keep monthly Snapshot copies of your data over a 20-year span, for example, to comply with government accounting regulations for your business. Because there is no requirement to serve data from secondary storage, you can use slower, less expensive disks on the vault system.

The different weights that SnapMirror and SnapVault give to backup currency and backup retention ultimately derive from the limit of 255 Snapshot copies for each volume. While SnapMirror retains the most recent copies, SnapVault retains the copies made over the longest period of time.

Preparing storage systems for SnapMirror replication

Before you can use SnapManager to mirror Snapshot copies, you need to configure a *data-protection relationship* between the source and destination volumes, then initialize the relationship. Upon initialization, SnapMirror makes a Snapshot copy of the source volume, then transfers the copy and all the data blocks that it references to the destination volume. It also transfers any other, less recent Snapshot copies on the source volume to the destination volume.

Before you begin

- You must have created the source and destination volumes in peered clusters with peered storage virtual machines (SVMs).
[Cluster and SVM Peering Express Guide](#)
- You must be a cluster administrator.
- For Snapshot copy verification on the destination volume, the source and destination storage virtual machines (SVMs) must have a management LIF as well as a data LIF.

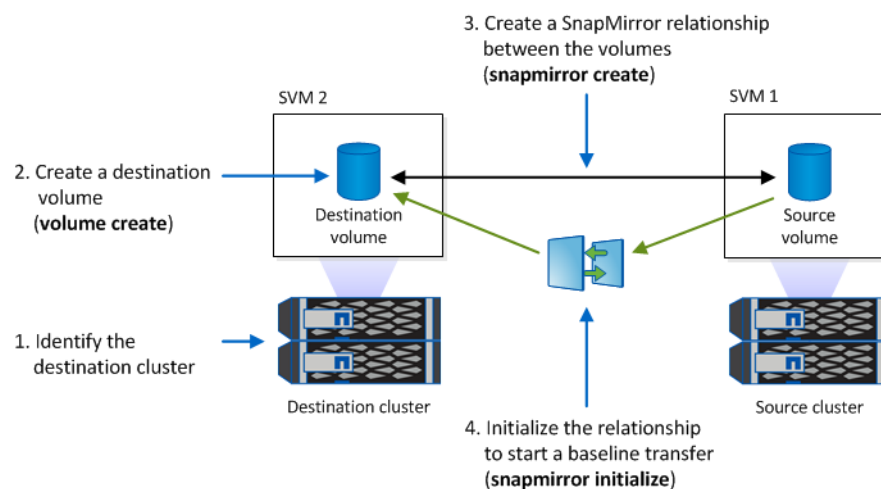
The management LIF must have the same DNS name as the SVM. Set the management LIF role to **data**, the protocol to **none**, and the firewall policy to **mgmt**.

About this task

You can use the ONTAP command-line interface (CLI) or OnCommand System Manager to create a SnapMirror relationship. The following procedure documents CLI usage.

Important: If you are storing database files and transaction logs on different volumes, you must create relationships between the source and destination volumes for the database files and between the source and destination volumes for the transaction logs.

The following illustration shows the procedure for initializing a SnapMirror relationship:



Steps

1. Identify the destination cluster.

2. On the destination cluster, use the `volume create` command with the `-type DP` option to create a SnapMirror destination volume that is either the same or greater in size than the source volume.

Important: The language setting of the destination volume must match the language setting of the source volume.

Example

The following command creates a 2 GB destination volume named `dstvolB` in SVM2 on the aggregate `node01_aggr`:

```
cluster2::> volume create -vserver SVM2 -volume dstvolB -aggregate node01_aggr -type DP
-size 2GB
```

3. On the destination SVM, use the `snapmirror create` command with the `-type DP` parameter to create a SnapMirror relationship.

The `DP` type defines the relationship as a SnapMirror relationship.

Example

The following command creates a SnapMirror relationship between the source volume `srcvolA` on SVM1 and the destination volume `dstvolB` on SVM2, and assigns the default SnapMirror policy `DPDefault`:

```
SVM2::> snapmirror create -source-path SVM1:srcvolA -destination-path SVM2:dstvolB
-type DP
```

Note: Do not define a mirror schedule for the SnapMirror relationship. SnapManager does that for you when you create a backup schedule.

If you do not want to use the default SnapMirror policy, you can invoke the `snapmirror policy create` command to define a SnapMirror policy.

4. Use the `snapmirror initialize` command to initialize the relationship.

The initialization process performs a *baseline transfer* to the destination volume. SnapMirror makes a Snapshot copy of the source volume, then transfers the copy and all the data blocks it references to the destination volume. It also transfers any other Snapshot copies on the source volume to the destination volume.

Example

The following command initializes the relationship between the source volume `srcvolA` on SVM1 and the destination volume `dstvolB` on SVM2:

```
SVM2::> snapmirror initialize -destination-path SVM2:dstvolB
```

Preparing storage systems for SnapVault replication

Before you can use SnapManager to perform disk-to-disk backup replication, you need to configure a *data-protection relationship* between the source and destination volumes, then initialize the relationship. On initialization, SnapVault makes a Snapshot copy of the source volume, then transfers the copy and all the data blocks it references to the destination volume.

Before you begin

- You must have created the source and destination volumes in peered clusters with peered storage virtual machines (SVMs).

Cluster and SVM Peering Express Guide


- You must be a cluster administrator.

About this task

You can use the ONTAP command-line interface (CLI) or OnCommand System Manager to create SnapVault relationships. The following procedure documents CLI usage.

Important: If you are storing database files and transaction logs on different volumes, you must create relationships between the source and destination volumes for the database files and between the source and destination volumes for the transaction logs.

The following illustration shows the procedure for initializing a SnapVault relationship:

**Steps**

1. Identify the destination cluster.
2. On the destination cluster, use the `volume create` command with the `-type DP` option to create a SnapVault destination volume that is the same size as or larger than the source volume.

Important: The language setting of the destination volume must match the language setting of the source volume.

Example

The following command creates a 2 GB destination volume named `dstvolB` in SVM2 on the aggregate `node01_aggr`:

```
cluster2::> volume create -server SVM2 -volume dstvolB -aggregate node01_aggr -type DP
-size 2GB
```

3. On the destination SVM, use the `snapmirror policy create` command to create a SnapVault policy.

Example

The following command creates the SVM-wide policy `SVM1-vault`:

```
SVM2::> snapmirror policy create -vserver SVM2 -policy SVM1-vault
```

Note: Do not define a cron schedule or Snapshot copy policy for the SnapVault relationship. SnapManager does that for you when you create a backup schedule.

4. Use the `snapmirror policy add-rule` command to add rules to the policy that define the following Snapshot copy labels and the retention policy for each label:
 - Daily
 - Weekly
 - Monthly
 - Hourly
 - Unlimited

Important: The labels are case-sensitive.

These are fixed labels for options displayed when you perform a backup. You select one of these options when you archive the backup. You must execute this command once for each of the rules you are adding.

Example

The following command adds a rule to the SVM1-vault policy that defines the “Daily” label and specifies that 30 Snapshot copies matching the label should be kept in the vault:

```
SVM2::> snapmirror policy add-rule -vserver SVM2 -policy SVM1-vault
-snapmirror-label Daily -keep 30
```

5. Use the `snapmirror create` command with the `-type XDP` parameter and the `-policy` parameter to create a SnapVault relationship and assign a vault policy.

The **XDP** type defines the relationship as a SnapVault relationship.

Example

The following command creates a SnapVault relationship between the source volume `srcvolA` on SVM1 and the destination volume `dstvolB` on SVM2, and assigns the policy SVM1-vault:

```
SVM2::> snapmirror create -source-path SVM1:srcvolA -destination-path SVM2:dstvolB
-type XDP -policy SVM1-vault
```

6. Use the `snapmirror initialize` command to initialize the relationship.

The initialization process performs a *baseline transfer* to the destination volume. SnapMirror makes a Snapshot copy of the source volume, then transfers the copy and all the data blocks it references to the destination volume.

Example

The following command initializes the relationship between the source volume `srcvolA` on SVM1 and the destination volume `dstvolB` on SVM2:

```
SVM2::> snapmirror initialize -destination-path SVM2:dstvolB
```

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.

Related tasks

- [SnapManager backup overview](#) on page 30
- [Defining a backup strategy](#) on page 30
- [Backing up your databases for the first time](#) on page 33
- [Verifying the initial backup set](#) on page 34
- [Scheduling recurring backups](#) on page 35
- [Scheduling recurring transaction log backups](#) on page 36
- [Scheduling recurring backup set verifications](#) on page 37

SnapManager backup overview

SnapManager uses NetApp Snapshot technology to create online, read-only copies of databases. It uses an SQL Server utility 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
- Transaction logs
- SnapInfo directories

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 Database Consistency Checker (DBCC), a Microsoft SQL Server utility, to verify the page-level integrity of databases. Verification ensures that you can use backup sets to restore databases as needed.

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

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 4353: Best Practices Guide for Microsoft SQL Server and SnapManager 7.1 for SQL Server with Clustered Data ONTAP](#).

What type of SnapManager backup do you need?

SnapManager supports two types of backups:

Backup type	Description
Full database backup	Backs up database files and truncated transaction logs. SQL Server truncates transaction logs by removing entries already committed to the database. This is the most common backup type.
Transaction log backup	Backs up truncated transaction logs, copying only transactions committed since the most recent backup. If you schedule transaction log backups to work with full database backups, SnapManager can restore databases to a specific recovery point more quickly. For example, you might schedule full database backups at the start and end of the day and transaction log backups every hour.

For both types of backups, you can choose the *copy-only* option to specify that SQL Server not truncate transaction logs. Use this option when you are backing up your databases with another backup application. Keeping transaction logs intact ensures that any backup application can restore the databases. You typically should not use copy-only in any other circumstance.

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 restore operations.

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. A single verification job can verify multiple backup sets at the same time.

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	Adds a time stamp to all Snapshot copy names. This is the default option. Example:  sqlsnap__QATX86W2K8_07-02-2014_10.45.08
Generic	Adds the string “recent” to the name of the most recent Snapshot copy. All other Snapshot copies include a time stamp. Example:  sqlsnap__QATX86W2K8__recent




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”.

Also, when the database resides on a VMDK, you must use the Unique naming convention when you want to clone Snapshot copies.

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:  sqlsnap__QATX86W2K8_07-02-2014_10.45.08
Daily	Adds “Daily” to Snapshot copy names. Example:  sqlsnap__QATX86W2K8_07-02-2014_10.47.00__Daily
Weekly	Adds “Weekly” to Snapshot copy names. Example:  sqlsnap__QATX86W2K8_07-02-2014_10.49.03__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 restore operations*, 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 restore operations* only, which restore a database to the time that SnapManager completed a full backup.

Typically, you 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 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.

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 and Verification** dialog box, keep **Full database backup** selected and define the properties for the backup job:

For this field...	Do this...
Copy-only backup	If you back up your databases using another backup application, select this field.
Run transaction log backup after full database backup	Keep this field selected.
Backup management group	Select a management group.
Delete full backups	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 databases after backup	Clear this field because it is best to verify databases in a separate operation.
SnapMirror options	If you set up a SnapMirror destination volume, select the option to replicate the backup copy to the destination volume.
Backup archiving options	If you set up a SnapVault destination volume, select the option to archive the backup copy to the destination volume.
Run Command	If you want to run a command before or after the backup operation, click the ... button and specify details for the command: where to run the command, the path to the program or script, the SnapManager variables to execute, and the command arguments.
Federated Backup	If you want to back up databases from different instances or different servers, click this field and then add databases to a federated group.
Availability Group Backup	If you have an Availability Group and you want to take backups on all replicas, the primary replica, the secondary replica, or preferred replicas, click this field and specify the replicas.

5. Click **Backup Now**.
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.

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 and Verification** dialog box, select **Verify most recent unverified snapshot backups only** and then define the properties for the backup verification:

For this field...	Do this...
Number of snapshot backups to verify	Keep the default. You should have only one backup set at this point.
Backup management group	Select a management group.
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 Verify on available SnapMirror destination volumes .
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 Verify archive backup on secondary storage .
Run Command	If you want to run a command before or after the operation, click the ... button and specify details for the command: where to run the command, the path to the program or script, the SnapManager variables to execute, and the command arguments.
Availability Group Backup	If you are scheduling the verification from the primary server, you can use this option to perform verification on databases whose backup was taken on a secondary replica copy.

4. Click **Verify Now**.
5. 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 the SQL Server Agent or Windows Scheduled Tasks.

About this task

In a Windows Failover cluster, it is a best practice to schedule jobs using the SQL Server Agent.

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 and Verification** dialog box, keep **Full database backup** selected and define the properties for the backup schedule, as described in [Backing up your databases for the first time](#) on page 33.
4. Click **Schedule**.
5. In the **Schedule Job** dialog box, enter a job name, choose a schedule service (SQL Server Agent or Windows Scheduled Tasks), and click **OK**.
6. Create the schedule using the service that you chose:

If you chose...	Do this...
The SQL Server Agent	<ol style="list-style-type: none"> Click Yes. SQL Server Management Studio opens. Connect to the SQL instance. In the Object Explorer pane, expand the instance. Expand SQL Server Agent, expand Jobs, right-click the job and then click Properties. Click Schedules and then click New. Fill out the New Job Schedule dialog box and click OK.
Windows Scheduled Tasks	<ol style="list-style-type: none"> Click Schedule. Specify the schedule. Click OK. Click Yes to save the job.

After you finish

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

Scheduling recurring transaction log backups

You should schedule transaction log backups alongside full database backups at a frequency that allows you to meet your Recovery Point Objective (RPO).

About this task

In a Windows Failover cluster, it is a best practice to schedule jobs using the SQL Server Agent.

Steps

- In the **Backup** pane, select the databases that you want to include in the backup schedule.
- In the **Actions** pane, click **Backup and Verify**.
- In the **Backup and Verification** dialog box, select **Transaction log backup** and then define the properties for the transaction log backup schedule:

For this field...	Do this...
Copy-only log backup	If you back up your databases using another backup application, select this field.
Verify log backup upon completion	Clear this field because it is best to verify backups in a separate operation.
Delete log backups	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.
Update SnapMirror after operation	If you set up a SnapMirror destination volume, select this field to replicate the backup copy to the destination volume.

For this field...	Do this...
Federated Backup	If you want to back up transaction logs from different instances or different servers, click this field and then add databases to a federated group.
Marking Transaction Options	If you specified a federated backup group, choose the default mark name and description or modify them. You use these marks to restore databases to the same marked transaction across multiple databases for a synchronous restoration.
Run Command	If you want to run a command before or after the backup operation, click the ... button and specify details for the command: where to run the command, the path to the program or script, the SnapManager variables to execute, and the command arguments.
Availability Group Backup	If you have an Availability Group and you want to take backups on all replicas, the primary replica, the secondary replica, or preferred replicas, click this field and specify the replicas.

4. Click **Schedule.**

5. In the **Schedule Job dialog box, enter a job name, choose a schedule service (SQL Server Agent or Windows Scheduled Tasks), and click **OK**.**

6. Create the schedule using the service that you chose:

If you chose...	Do this...
The SQL Server Agent	<ol style="list-style-type: none"> a. Click Yes. SQL Server Management Studio opens. b. Connect to the SQL instance. c. In the Object Explorer pane, expand the instance. d. Expand SQL Server Agent, expand Jobs, right-click the job and then click Properties. e. Click Schedules and then click New. f. Fill out the New Job Schedule dialog box and click OK.
Windows Scheduled Tasks	<ol style="list-style-type: none"> a. Click Schedule. b. Specify the schedule. c. Click OK. d. Click Yes to save the job.

Scheduling recurring backup set verifications

You can schedule recurring backup set-verification jobs using the SQL Server Agent or Windows Scheduled Tasks.

About this task

In a Windows Failover cluster, it is a best practice to schedule jobs using the SQL Server Agent.

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 and Verification** dialog box, select **Verify most recent unverified snapshot backups only** and define the properties for the backup verification schedule as described in [Verifying the initial backup set](#) on page 34.

You might need to modify the **Number of snapshot backups to verify** field, depending on the number of backups SnapManager will take between scheduled verifications.

4. Click **Schedule**.
5. In the **Schedule Job** dialog box, enter a job name, choose a schedule service (SQL Server Agent or Windows Scheduled Tasks), and click **OK**.
6. Create the schedule using the service that you chose:

If you chose...	Do this...
The SQL Server Agent	<ol style="list-style-type: none"> a. Click Yes. SQL Server Management Studio opens. b. Connect to the SQL instance. c. In the Object Explorer pane, expand the instance. d. Expand SQL Server Agent, expand Jobs, right-click the job and then click Properties. e. Click Schedules and then click New. f. Fill out the New Job Schedule dialog box and click OK.
Windows Scheduled Tasks	<ol style="list-style-type: none"> 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 verification job in the SnapManager Scheduled Jobs pane.

Where to go next

After you have configured backups in SnapManager, you can perform full or partial restores as necessary. You can also explore other important SnapManager features, such as cloning, reports, control files, and PowerShell cmdlets.

You can find more information about these features, as well as release-specific information for SnapManager, in the following documentation, available on the NetApp Support Site.

- [*SnapManager 7.2 for Microsoft SQL Server Administration Guide*](#)
Describes how to administer SnapManager after deployment is complete. Topics include how to restore the database, how to import configuration information from a control file, how to clone databases, how to use the SnapManager PowerShell cmdlets, and how to upgrade and uninstall the product.
- [*SnapManager 7.2 for Microsoft SQL Server Release Notes*](#)
Describes new features, important cautions, known problems, and limitations for the SnapManager 7.x product.
- Describes how to quickly configure peer relationships between clusters and SVMs.
- [*ONTAP 9 Volume Disaster Recovery Preparation Express Guide*](#)
Describes how to prepare storage system for SnapMirror replication.
- [*ONTAP 9 Volume Backup Using SnapVault Express Guide*](#)
Describes how to prepare storage system for SnapVault replication.
- [*NetApp Technical Report 4225: Best Practice Guide for Microsoft SQL Server and SnapManager 7.0 for SQL Server with Clustered Data ONTAP*](#)
Describes SnapManager for Microsoft SQL Server best practices.

Copyright information

Copyright © 2018 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

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., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

- additional information
 - about SnapManager features [39](#)
- administration server
 - Windows host requirements [13](#)
- AlwaysOn Availability Group cluster
 - prerequisites and considerations for installing on [16](#)
- AutoSupport
 - configuring [20](#)
- Availability Groups
 - connecting to [19](#)
 - specifying a share for transaction logs [20](#)

B

- backing up databases
 - for the first time [33](#)
 - using a schedule [35](#)
- backup sets
 - verifying the initial set [34](#)
 - verifying with a schedule [37](#)
- backup types
 - overview of [30](#)
- backups
 - using SnapMirror and SnapVault replication [25](#)
- base configuration
 - Windows host requirements [13](#)
- benefits and features
 - overview of [4](#)
- best practices [10](#)

C

- comments
 - how to send feedback about documentation [42](#)
- configuration
 - workflow [6](#)
- considerations
 - installation [16](#)

D

- databases
 - backing up for the first time [33](#)
 - backing up with a schedule [35](#)
 - migrating to NetApp storage [20](#)
 - name requirements [19](#)
 - overview of backing up [30](#)
 - preparing for SnapMirror replication [26](#)
 - preparing storage systems for SnapVault replication [27](#)
 - prerequisites for migrating [19](#)
 - strategy for backing up [30](#)
 - verifying the initial backup set [34](#)
 - verifying with a schedule [37](#)
- deployment
 - preparing, SnapManager [7](#)

- documentation
 - how to receive automatic notification of changes to [42](#)
 - how to send feedback about [42](#)

E

- email notifications
 - configuring [20](#)

F

- features
 - where to find additional information about SnapManager [39](#)
- features and benefits
 - overview of [4](#)
- feedback
 - how to send comments about documentation [42](#)

I

- IMT
 - using to verify support for system configurations [11](#)
 - verifying your configuration [7](#)
- information
 - how to send feedback about improving documentation [42](#)
- installing
 - prerequisites and considerations for [16](#)
- installing SnapManager
 - interactively [16](#)
 - silently [17](#)
 - workflow [6](#)
- Interoperability Matrix Tool
 - See* IMT

M

- management groups
 - overview of [30](#)
- migrating
 - database files [20](#)
 - prerequisites for [19](#)
 - transaction logs [20](#)

N

- naming conventions
 - overview of [30](#)
- NetApp Interoperability Matrix Tool
 - See* IMT

P

- per-server licensing
 - SnapManager for Microsoft SQL Server [10](#)

- per-storage system licensing
 - SnapManager for Microsoft SQL Server [10](#)
- preparations
 - for deployment, SnapManager [7](#)
- prerequisites
 - installation [16](#)
- product overview
 - features and benefits [4](#)

R

- Recovery Point Objective
 - meeting [36](#)
- replication
 - considerations for preparing storage systems for SnapMirror and SnapVault [25](#)
 - preparing storage systems for SnapMirror [26](#)
 - preparing storage systems for SnapVault [27](#)
- requirements
 - service account [15](#)

S

- service accounts
 - requirements [15](#)
- SnapInfo directory
 - mapping to data files and logs [20](#)
- SnapManager
 - supported storage types [12](#)
 - where to find additional information about features [39](#)
- SnapManager for Microsoft SQL Server
 - administration server [10](#), [13](#)
 - archiving Snapshot copies [27](#)
 - backup overview [30](#)
 - backup strategy [30](#)
 - base configuration [10](#), [13](#)
 - connecting to SQL Server instances [19](#)
 - deployment workflow [6](#)
 - features and benefits [4](#)
 - installing
 - interactively [16](#)
 - installing silently [17](#)
 - licensing [10](#)
 - mirroring Snapshot copies [26](#)
 - storage layout requirements [7](#)
 - verification server [10](#), [13](#)
 - VMDK support [13](#)
 - Windows Firewall requirements [13](#)
 - Windows host requirements [13](#)
- SnapMirror
 - considerations for preparing storage systems for replication [25](#)

- differences from SnapVault [25](#)
 - preparing storage systems for replication [26](#)
- Snapshot copies
 - preparing for SnapVault replication [27](#)
 - preparing storage systems for SnapMirror replication [26](#)
- SnapVault
 - considerations for preparing storage systems for replication [25](#)
 - differences from SnapMirror [25](#)
 - preparing storage systems for replication [27](#)
- SQL Server instances
 - connecting SnapManager to [19](#)
- storage layout requirements
 - for SnapManager for Microsoft SQL Server [7](#)
- storage systems
 - considerations for preparing for SnapMirror and SnapVault replication [25](#)
- storage types
 - supported by SnapManager [12](#)
- suggestions
 - how to send feedback about documentation [42](#)
- support
 - storage types [12](#)
 - verifying SnapManager support for system configurations [11](#)
- system configurations
 - verifying SnapManager support for [11](#)
- systems
 - considerations for preparing for SnapMirror and SnapVault replication [25](#)

T

- transaction log backups
 - using a schedule [36](#)
- transaction logs
 - backing up with a schedule [36](#)
- Twitter
 - how to receive automatic notification of documentation changes [42](#)

V

- verification server
 - selecting [20](#)
 - Windows host requirements [13](#)
- VMDK support
 - SnapManager for Microsoft SQL Server [13](#)