



Technical Report

Microsoft SQL Server and SnapManager for SQL Deployment Guide

Cheryl George, Mohamed Niyaz, NetApp
July 2014 | TR-4302

Executive Summary

This deployment guide assists customers in deploying solutions as per best practices to meet specific business requirements. This guide is documented considering the deployed solution based on experiences with existing NetApp customers, real-world simulations, and NetApp engineering lab validations. Deployment guides help customers through the entire project lifecycle, including requirement assessment, solution design, installation, and administration.

Note: This guide specifically focuses on NetApp® clustered Data ONTAP® deployment scenarios.

TABLE OF CONTENTS

1	Business Requirements	4
2	Solution Deployment Details	4
2.1	Storage Environment	4
2.2	Additional NetApp Software Used	4
2.3	Virtualization	5
2.4	Database Storage Layout	6
2.5	Installation and Configuration	8
2.6	Install and Configure SnapDrive on SQL Server	10
2.7	Install SnapManager for SQL Server	12
2.8	Test and Validate Solution	12
3	Solution Operation	14
3.1	Capacity Management	14
3.2	Backup and Recovery Operations	15
4	Conclusion	15
	Appendix A: Installation and Configuration Details	16
	Detailed Steps to Install SnapDrive for Windows	16
	Create an RDM LUN on a Guest OS	21
	Create VMDK on VMFS Datastore	22
	Create VMDK on NFS Datastore	22
	Detailed Steps to Install SnapManager for SQL Server	23
	Appendix B: Build the SQL 2012 Availability Group for High Availability and Disaster Recovery	24
	Failover Clustering Feature	25
	Install a Standalone Instance of SQL Server 2012 on Each Cluster Node	27
	Enable AlwaysOn Availability Group Feature	28
	Install SnapDrive for Windows on Each Node of the WSFC Cluster	29
	Install SnapManager for SQL on Each Node of the WSFC Cluster	29
	Create Initial Database Backup	29
	Create SQL 2012 Availability Group	30
	Create Availability Group Listener	31
	Test and Validation Details	31
	Basic Tests to Validate the Solution	31
	Solution Operation Details	35
	Create SQL Server Database Backup	35

Restore SQL Server Database	37
Clone SQL Server Database from SnapManager for SQL Server Backup	39
References	41
Version History	42

LIST OF TABLES

Table 1) Features and software used in this solution.	4
Table 2) SQL Server implementation protocols.	6
Table 3) Installation and configuration procedure.	8
Table 4) Required components.	10
Table 5) Prerequisites prior to installing SnapDrive for Windows.	11
Table 6) Prerequisites prior to installing SnapManager for SQL Server.	12
Table 7) SMSQL readiness tests performed.	13
Table 8) Connectivity and validation tests performed.	13
Table 9) Backup and recovery tests.	13
Table 10) SnapManager for SQL Server backup tests.	14
Table 11) SnapManager for SQL Server restore tests.	14
Table 12) Replica settings.	24

LIST OF FIGURES

Figure 1) SMSQL creates Snapshot copy, then leverages NetApp FlexClone and SnapMirror technologies.	5
Figure 2) SMSQL volume layout.	7
Figure 3) Nodes added to Windows Server Failover Cluster (WSFC).	26
Figure 4) Enable TCP/IP protocol for MS SQL Server.	28
Figure 5) Enable AlwaysOn availability group feature.	29
Figure 6) SQL 2012 availability group "AG2012" with nodes SQL 1, SQL 2, SQL 3, and SQL 0 added.	31

1 Business Requirements

Today's business applications are more data centric than in the past, requiring fast and reliable access to intelligent information structures that are often provided by a high-performance relational database system. Many organizations use Microsoft® SQL Server® as a back-end datastore for mission-critical business applications. The latest release, Microsoft SQL Server 2012, delivers performance, scalability, availability, and security. SQL Server implementations have become more complex and require more reliability than before. Database service-level agreements require predictable performance, and outages are disruptive and costly for database consumers. The underlying physical storage architectures supporting SQL Server are expected to scale in order to meet the growing capacity, but frequently bottlenecks arise, and backup processes get slower as databases grow. SQL Server databases are growing significantly larger to meet the needs of today's organizations, while business requirements dictate shorter backup and restore windows.

2 Solution Deployment Details

This section discusses the details of the various components used in the solution and the NetApp storage layout used for SQL Server, highlighting the components and the overall architecture.

2.1 Storage Environment

NetApp FAS Storage Controllers

The NetApp FAS series is the storage platform used in this solution. This platform offers the following advantages for this solution:

- Capacity management, which provides the ability to grow as demands on collaboration increase
- Integration of the layers of storage, Microsoft Windows®, and Microsoft SQL Server to simplify and automate data backup and restore
- Data protection to meet the high-availability requirements of SQL Server

2.2 Additional NetApp Software Used

Table 1 lists an array of NetApp software to augment the storage platform, which assists in deployment, backup, recovery, replication, management, and data protection.

Table 1) Features and software used in this solution.

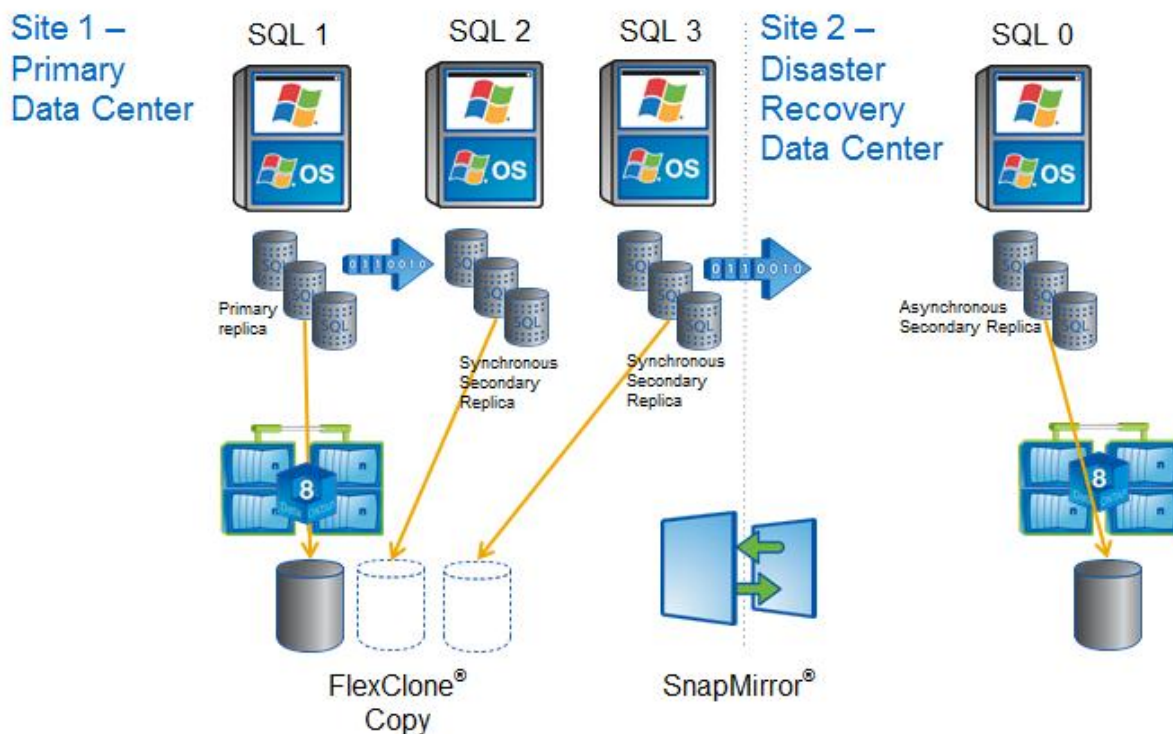
Software and Feature	Function	Benefit
FlexVol® technology®	Creates flexibly sized LUNs and volumes across a large pool of disks.	Fast, simple, and flexible storage provisioning and high-capacity utilization.
Snapshot™ technology	Makes incremental, data-in-place, point-in-time copies of a LUN or volume with minimal performance impact.	Enables quick and frequent backups, nondisruptively and efficiently.
SnapRestore®	Rapidly restores single files, directories, or entire LUNs and volumes from any Snapshot technology backup.	Enables near-instantaneous recovery of files, databases, and complete volumes.
SnapDrive® for Windows	Provides host-based data management of NetApp storage from Windows Server®.	Simplifies host-consistent Snapshot copy creation and automates error-free restores.
SnapManager® for	Provides host-based data management	Simplifies application-consistent

Software and Feature	Function	Benefit
Microsoft SQL Server	of NetApp storage for Microsoft SQL Server databases.	Snapshot copies, automates error-free data restores, and enables application-aware disaster recovery.

SnapManager for SQL Server complements SQL Server availability and broadens the data protection of availability groups to achieve nondisruptive operations that support all workflows, including backup, restore, cloning, and DR for AGs as a single unit. With SnapManager, customers can accelerate availability group setup, rapidly back up and restore, make clones of all databases in availability groups, and quickly resynchronize databases within availability groups. In addition, databases in an availability group can be mirrored to remote locations using NetApp SnapMirror® technology.

Figure 1 illustrates the disaster recovery for SQL Server 2012 with AlwaysOn availability groups.

Figure 1) SMSQL creates Snapshot copy, then leverages NetApp FlexClone and SnapMirror technologies.



For more information, refer to [SnapManager for Microsoft SQL Server](#).

2.3 Virtualization

Server virtualization is a major component of data center virtualization and plays a key role in the virtualization initiative. NetApp has been on the forefront of solving complex business problems with its innovative technology breakthroughs and end-to-end solutions approach. The virtualization platform, through its ability to virtualize SQL Server, assists in efficient use of hardware resources that can be combined with the other key advantages of server virtualization, which include better availability, lower cost, and increased flexibility. You can realize multiple benefits from using SQL Server in a virtualized environment with NetApp storage technology, including:

- **Effective use of server hardware.** Migrating the entire SQL Server from dedicated physical servers that have relatively low utilization rates can lead to significantly higher server utilization.
- **Savings.** You save on power and space.
- **Reduced server hardware requirements.** The number of physical servers required to support SQL Server can also be reduced.

SQL Server supports virtualization using Hyper-V® or VMware®. SnapManager for SQL Server (SMSQL) is qualified with guest initiator LUNs and pass-through disks in Hyper-V environments. In VMware environments, SnapManager for SQL Server is validated with guest initiator LUNs (easy for deployment) and iSCSI/FC RDMs or VMFS or NFS.

Table 2 lists the five protocols that SQL Server can use to communicate between server and storage.

Table 2) SQL Server implementation protocols.

Environment Type	Fibre Channel (FC)	Fibre Channel over Ethernet (FCoE)	Internet Small Computer System Interface (iSCSI)	Network File System (NFS)	Server Message Block (SMB)
Physical	Yes	Yes	Yes	Yes	Yes
VMware guest	Raw device mapping (RDM)	RDM	Guest or RDM	VMware disk (VMDK)	Yes
Hyper-V guest	RDM	RDM	Guest or RDM	No	Yes

2.4 Database Storage Layout

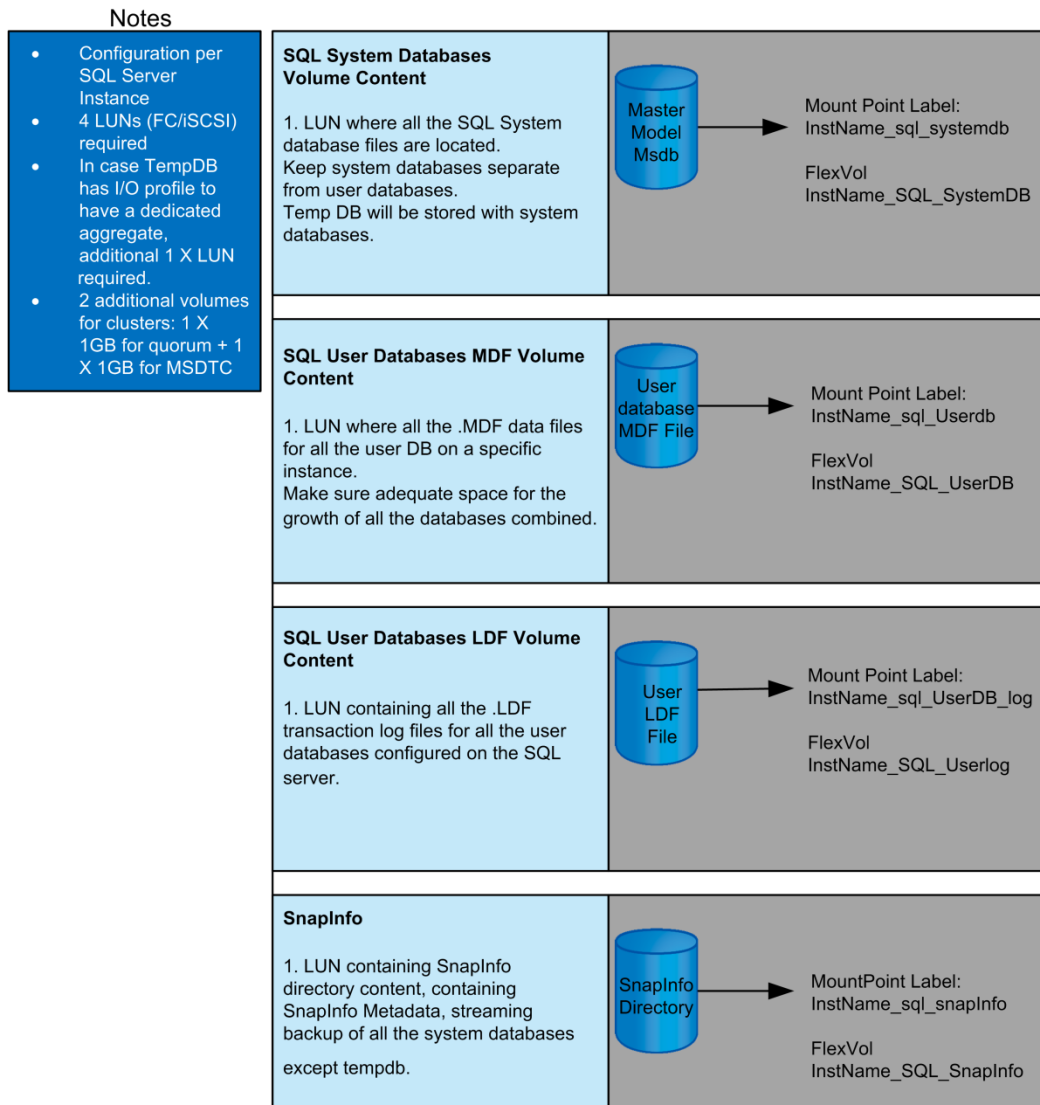
This section describes how the storage from the NetApp storage controller will be configured. FlexVol volumes will be provisioned for the SQL Server instance. Following is an example of the SQL Server design for NetApp storage and considerations of environments that use SnapManager for SQL Server.

- Do not use SQL Server partitions beyond the default configuration.
- One aggregate for the SQL Server instance.
- A dedicated vol/LUN for the SQL Server system databases: master, model, msdb. Keep system databases in a separate LUN and FlexVol volume because colocating system databases with user databases in a single LUN prevents NetApp Snapshot backups of the user databases.
- A dedicated vol/LUN for the SQL Server system databases: tempdb, depending on its I/O. No backup will be configured for TempDB database because it is rebuilt every time SQL Server restarts.
- A dedicated vol/LUN for each user database to allow NetApp Snapshot backups of user databases.
- A dedicated LUN for transaction log files to separate the random I/O of the data files from the sequential I/O to the log files to improve SQL Server performance.
- A dedicated vol for databases that reside on SMB share.
- A dedicated vol/LUN for the SMSQL SnapInfo directory.

Note: For SnapManager for SQL Server, there can be one SnapInfo directory for all the databases, which is implied in this design, or one per database.

Figure 2 explains the volume layout.

Figure 2) SMSQL volume layout.



LUN Sizing

SQL Server System Database FlexVol Volume and LUN Sizing

System databases are not configured for backups based on Snapshot; backups of these databases are streamed to the SnapInfo area on a separate LUN. Snap reserve can be set to 0%, and the LUN in the FlexVol volume can be almost as large as the FlexVol volume itself.

A 100GB LUN for system database in a 120GB FlexVol volume will be set up for each SQL Server instance.

SQL Server User Database Data (.mdf) FlexVol Volume and LUN Sizing

A separate LUN for storing user MDF database files will be created per instance. Database DB change rate is assumed to be 10% as a standard for all user DB.

The size of the user databases MDF FlexVol volume is calculated as follows:

= (database MDF LUN size x1) + (max .mdf size in LUN x daily change rate expressed as a decimal x # days Snapshot copies to keep)

Note: Fractional reserve will be set to 0% for all the volumes, which removes the requirement of having LUN x 2 free space in the volume.

SQL Server User Database Transaction Log (.ldf) FlexVol Volume and LUN Sizing

A separate LUN for storing user LDF database files will be created per instance. Database DB log change rate is assumed to be 10% as a standard for all user DB logs.

The size of the user databases LDF FlexVol volume is calculated as follows:

= (database LDF LUN size) + (max .ldf size in LUN x daily change rate expressed as a decimal x # days Snapshot copies to keep)

Note: Fractional reserve will be set to 0%, which removes the requirement of having LUN x 2 free space in the volume.

SMSQL SnapInfo FlexVol Volume and LUN Sizing

The SnapInfo LUN will contain the following:

- Streaming backups of system databases, except TempDB, which is not backed up
- Backups of the active portion of user database transaction logs
- SnapManager for SQL Server (SMSQL) backup set metadata

The data in the SnapInfo LUN is arranged in a folder structure beneath a SnapInfo folder in the root of the drive. The amount of data in the SnapInfo area is dependent on the size of a backup and the number of days' backups that are retained online.

The size of the SnapInfo LUN is calculated as follows:

= ((system database size + (maximum DB LDF size X daily log change rate %)) X Snapshot retention)/(1-LUN overhead space %)

The SnapInfo LUN sizing calculation above assumes the following:

- System databases backup. This figure does not include TempDB. TempDB is recreated every time SQL Server is restarted and is not backed up.
- Backups are retained for 7 days on primary storage (online backups based on Snapshot).
- 10% LUN overhead space.

The size of the SnapInfo FlexVol volume is calculated as follows:

= ((size of SnapInfo LUN) x (1+fractional reserve)) + (system database size+ (maximum DB LDF size X daily log change rate %)) X Snapshot retention)

Note: Fractional reserve is set to 0% for all the volumes and will not be used.

2.5 Installation and Configuration

This section provides an overview of the installation sequence, as well as specific configuration parameters, with a focus on the solution's NetApp components.

Table 3 lists the primary tasks for installing and configuring the solution.

Table 3) Installation and configuration procedure.

Steps	Task	Description
1.	Review the solution design	<ul style="list-style-type: none">• Review and sign off on the requirements and design

Steps	Task	Description
		<ul style="list-style-type: none"> Make sure that business objectives and IT deliverables are aligned
2.	Prepare data centers	<ul style="list-style-type: none"> Rack space and power preparations Cabling, network ports, and SAN ports TCP/IP address and DNS host names <p>Note: The power supply requirements are relevant to the hardware that is used. For information about power supply requirements specific to the customer environment, refer to the appropriate technical specification documentation.</p> <p>Note: For detailed information about the site preparation requirements, refer to the NetApp Site Requirements Guide.</p>
3.	Install server hardware	<ul style="list-style-type: none"> Install server hardware Comply with both internal standards and hardware vendor best practices
4.	Install and configure FAS storage arrays	<ul style="list-style-type: none"> Physical installation Install Data ONTAP Create aggregates, storage virtual machine (SVM), and volumes
5.	Install and configure Microsoft Windows Server	<ul style="list-style-type: none"> Validate Microsoft Windows operating system prerequisites for Microsoft SQL Server 2012 (refer to Planning a SQL Server Installation) Perform any required updates and patches
6.	Install and configure Microsoft SQL Server	<ul style="list-style-type: none"> Validate Microsoft prerequisites for Microsoft SQL Server 2012 Hardware and Software Requirements for Installing SQL Server 2012 Install SQL Server 2012 from the Installation Wizard (Setup) Perform any required updates and patches
7.	Install NetApp SnapDrive for Windows on all nodes of the SQL Server 2012 AlwaysOn availability group	<ul style="list-style-type: none"> Install SnapDrive prerequisites Install SnapDrive software and provision LUNs from FAS controller
8.	Install NetApp SnapManager for SQL Server	<ul style="list-style-type: none"> Install SnapManager for SQL Server on all nodes in the SQL Server 2012 Availability Group Configure and schedule backups using SMSQL
9.	Test and validate the solution	<ul style="list-style-type: none"> Pretest of solution readiness Fault-tolerance test and validation: refer to appendix Validation of backup and restore: refer to appendix

For detailed installation and configuration steps, refer to “Installation and Configuration Details” in the appendix.

2.6 Install and Configure SnapDrive on SQL Server

Microsoft SQL Server 2012 depends heavily on storage for a stable and reliable configuration. NetApp offers the ability to capture the state of the SQL Server databases at various points in time called Snapshot copies.

Before setup begins, verify the compatibility of all hardware and software involved using the [NetApp Interoperability Matrix Tool](#).

1. Prepare Windows host part of the SQL Server 2012 AlwaysOn availability group in the SnapDrive configuration.
2. Verify that the host meets the minimum requirements for use with SnapDrive.
3. Determine whether the Microsoft iSCSI Software Initiator program is installed.
4. Determine whether SnapDrive was previously installed.
5. Determine which FC or iSCSI HBA or MPIO components are already installed.

SnapDrive supports three protocols for creating and managing LUNs: iSCSI, FC, and FCoE. Before installing SnapDrive for Windows, install these components on the host computer as listed in Table 4.

Table 4) Required components.

Scenario	Tasks
The iSCSI protocol and software initiator will be used to create and manage LUNs.	<ul style="list-style-type: none">• Install the Microsoft iSCSI Software Initiator.• Install the iSCSI Host Utilities on the hosts. <p>Note: NetApp highly recommends installing Data ONTAP DSM for multipathing.</p>
The iSCSI protocol and hardware initiator will be used to create and manage LUNs.	<ul style="list-style-type: none">• Install the iSCSI HBA.• Install the iSCSI HBA driver and firmware. <p>Note: NetApp highly recommends installing Data ONTAP DSM for multipathing. Install the iSCSI Host Utilities on the hosts if Data ONTAP DSM is not installed.</p>
The FC protocol will be used to create and manage LUNs.	<ul style="list-style-type: none">• Install FCP HBA or CNA for FCoE.• Install the FC driver and firmware. <p>Note: NetApp highly recommends installing Data ONTAP DSM for multipathing. Install the Windows Host Utilities on the hosts if Data ONTAP DSM is not installed.</p>
A Common Internet File System (CIFS) protocol for CIFS shares used by SMSP storage optimization modules.	<ul style="list-style-type: none">• License CIFS, and it must be set up and started on the NetApp storage system.• Create a CIFS share to host SQL Server databases.

For detailed SDW installation steps, refer to [Installation and Configuration Details](#) in the appendix.

Prepare Each Storage System in SnapDrive Configuration

1. After verifying that licenses for FC, iSCSI, or both are enabled on the storage system, start the services by entering the `fcv start` command or the `iscsi start` command at the storage system command line.

Note: For more information, refer to the appropriate Data ONTAP administration guide on the [NetApp Support](#) site.

2. Prepare a volume on the storage system to hold SnapDrive LUNs.

SnapDrive Prerequisites in Clustered Data ONTAP Environments

Before installing SnapDrive for Windows, the prerequisites in Table 5 must be met.

Table 5) Prerequisites prior to installing SnapDrive for Windows.

Description of Prerequisites
<p>The following licenses are required on the storage system:</p> <ul style="list-style-type: none">• Fibre Channel Protocol (FCP) or iSCSI (depending on the configuration): Use FC/iSCSI-accessed LUNs.• FlexClone® technology: Enable volume clone functionality on flexible volumes.• SnapRestore technology: Restore LUNs from Snapshot copies.• SnapDrive and SnapManager suite (either on the host or on the system): Use whichever license enables SnapDrive functionality when the SDW license is not on the host.• SnapVault® technology (optional).• SnapMirror technology (optional). <p>To determine which licenses are enabled on a storage system, complete these steps:</p> <ol style="list-style-type: none">1. Log in to the storage system through the console or telnet.2. Type <code>license</code> to display the list of licenses installed. <p>Note: This can also be done through System Manager.</p>
<p>A SnapDrive user service account on the storage system is required.</p> <p>Note: This account is required to connect SnapDrive to the storage system.</p>
<p>The transport protocol (HTTP, HTTPS, or RPC) that SnapDrive will use to communicate with the storage system must be determined.</p> <p>Note: NetApp recommends using HTTPS. The HTTPS protocol allows using the Data ONTAP interface for all interactions between the storage system and the host, including sending passwords securely. For SnapDrive to use the HTTP or HTTPS protocol, the <code>httpd.admin.enable</code> option must be set on the storage system.</p> <p>Note: The RPC protocol is not supported by SnapDrive on clustered Data ONTAP systems.</p>
<p>NET Framework 4.0 is required.</p>
<p>NetApp Windows Host Utilities 6.0.2 or later are required.</p>
<p>The required hotfix for Windows 2012 is 2859162, and for Windows 2008 R2 the hotfixes are 2522766, 2528357, 2494016, 2520235, 2531907, and 974930.</p>
<p>Cluster- and node-management logical interfaces (LIFs) are required.</p>
<p>An SVM with the following settings and configurations is required:</p> <ul style="list-style-type: none">• Data volumes with junction paths defined• SVM management LIF with the following parameter settings: <code>data=role, protocols=none, firewall policy=management</code>• Vsadmin password is set and the account is unlocked• On each node in the cluster, there is a data LIF for SAN protocols that is separate from any data LIFs

Note: Set appropriate storage controller options for the volume.

Note: Synchronize the storage controller clocks with the Active Directory® servers and configure DNS settings on the storage controller. Verify that name resolution is working before installing SnapDrive for Windows on the hosts.

For help with supported configurations for NetApp, refer to the [NetApp Interoperability Matrix Tool](#).

2.7 Install SnapManager for SQL Server

NetApp SnapManager for SQL Server provides the underlying host intelligence required to present storage over the iSCSI SAN storage protocol and to coordinate SQL Server aware backups using NetApp Snapshot technology.

Prior to installing SMSQL, make sure the prerequisites in Table 6 are met.

Table 6) Prerequisites prior to installing SnapManager for SQL Server.

Description of Prerequisites
Microsoft SQL Server application has been installed on the host where NetApp SnapManager for SQL Server (SMSQL) will be installed.
SQL Server browser service must be running on the host on which SMSQL will be installed.
The installation account used to install SMSQL must be a member of the local administrators group on the host where SnapManager for SQL Server will be installed.
The following components must be installed or configured on the storage system: <ul style="list-style-type: none">• iSCSI or Fibre Channel (FC) protocols• SnapManager license• SnapRestore license• SnapMirror license• FlexClone license• CIFS license
SnapDrive for Windows (SDW) has been installed on the SQL Server host on which SnapManager will be installed. Note: For more information about installing SnapDrive, refer to the SnapDrive for Windows Installation and Administration Guide .
The SnapDrive preferred IP address (if the storage system has multiple IP addresses), Microsoft iSCSI Software Initiator, and SnapDrive for Windows must be installed.
Storage must be configured, and the LUNs must be presented to the Windows Server.
The SMSQL service account must have a SQL Server login and sysadmin role assigned to the SQL Server login for the SQL Server instance on which SMSQL will be installed.
.NET Framework 4.0 has been installed on the host on which SMSQL will be installed.

For help with supported configurations for NetApp, refer to the [NetApp Interoperability Matrix Tool](#).

Prior to installing SMSQL, refer to the section “Preparing to Install or Upgrade SnapManager” in [SnapManager 7.0 for Microsoft SQL Server Installation and Administration Guide](#).

Download the SnapManager for SQL Server software from <http://support.netapp.com/> and follow the installation steps outlined in [Installation and Configuration Details](#) in the appendix.

2.8 Test and Validate Solution

This phase entails validating solution readiness using various backup and restore operations.

Pretest of Solution Readiness

After setting up and configuring the designed SQL Server environment in this reference architecture, it is important to test every component for its readiness in terms of functionality, resiliency, and availability.

Table 7 lists the tests that can be performed at each level to check the readiness of the architecture to meet its design specifications.

Table 7) SMSQL readiness tests performed.

Operation	Basic	Functionality	High Availability
Connect to SQL Server database engine	x		
Database access in SQL Server Management Studio	x		
Add a Windows authentication login in SQL Server Management Studio	x		
SQL Server database migration using SMSQL Configuration Wizard		x	

Storage Connectivity and Validation Tests

Table 8 lists the tests performed so that the basic storage connectivity is tested and validated before the solution is stress tested.

Table 8) Connectivity and validation tests performed.

Operation	Basic	Functionality	High Availability
SnapDrive create LUN		x	
SnapDrive destroy LUN		x	
SnapDrive expand LUN		x	
SnapDrive LUN connect		x	
SnapDrive LUN disconnect		x	
SnapDrive create Snapshot copy		x	
SnapDrive delete Snapshot copy		x	

Backup and Recovery Tests

Table 9 through Table 11 lists the SnapManager for SQL Server restore tests.

Table 9 lists the backup and recovery tests validated with the solution.

Table 9) Backup and recovery tests.

Operation	Basic	Functionality	High Availability
Periodic backups of SQL Server databases using		x	

Operation	Basic	Functionality	High Availability
SnapManager for SQL Server Snapshot copies: full backup should be taken at every 1-hour interval; log backup is taken every 15 minutes. Availability group (AG)-level full backup is taken every 24-hour interval, and AG-level log backup is taken every 30 minutes.			
Frequent recovery point backup		x	
Backup verification using SnapManager for SQL Server (every 24 hours)		x	x

Table 10 lists the SnapManager for SQL Server backup tests.

Table 10) SnapManager for SQL Server backup tests.

Operation	Description
Task	To test the backup of SnapManager for SQL Server backup and data retention job: Perform the tasks to create a schedule of SQL Server Snapshot copy backups to include hourly, daily, and weekly copies with granular and configurable retention of each. View SQL Server notification and logs of these backups. Monitor the time to do the average backup tasks.
Observed results	Observe the backup, backup cleanup, and total backup times.

Table 11 lists the SnapManager for SQL Server restore tests.

Table 11) SnapManager for SQL Server restore tests.

Operation	Description
Task	To test the granularity and scalability of SnapManager for SQL Server backup and data retention: Perform point-in-time restoration of an entire SQL Server database.
Observed results	Observe the point-in-time restore and up-to-the-minute restore.

For detailed SnapManager for SQL Server backup and restore steps, refer to [Solution Operation Details](#) in the appendix.

3 Solution Operation

3.1 Capacity Management

This section describes how to grow and manage the capacity of the NetApp storage systems as business needs grow.

FlexVol Volume Management

Flexible volumes can be sized according to capacity requirements. The size of a flexible volume can be increased or decreased. The resize can be accomplished from the command line (RSH, telnet, console) or from System Manager.

To resize the volume from the command line, enter the following command:

```
vol size -vserver <vserver name> -volume <vol_name> -new-size [+ | -] <New size>
```

In this command, `volname` is the name of the volume, and `size` is the space to be added to or removed from the volume. Size also includes a modifier to identify whether the space added is in kilobytes (k), megabytes (m), gigabytes (g), or terabytes (t). The plus or minus sign indicates whether space will be added to or subtracted from the volume.

Example: The volume name is `SQLDbvol1`, and the size is 250GB. The new desired size is 300GB. The command to change the size of the volume is:

```
vol size -vserver <vserver name> -volume vol_ SQLDbvol1 -new-size +50g
```

Expanding LUNs

To expand the LUNs, follow these steps.

1. Under SnapDrive in the left MMC pane, expand the instance of SnapDrive you want to manage, then expand Disks and select the disk you want to manage.
2. From the menu choices at the top of MMC, navigate to Action > Resize Disk.
3. Next to "Maximum size" in the Resize Disk window, leave "Reserve space for at least one Snapshot copy" selected.

Note: When you select this option, the disk size limits displayed are accurate only when they first appear on the Select LUN Properties panel.

4. In the "New size" box, either type a value or use the slider bar to increase or decrease the amount of space the disk uses.
5. Select "Take a Snapshot before resizing the LUN" checkbox to take a Snapshot copy before you resize your disk.
6. Click OK.
7. Create a new Snapshot copy of the resized disk.

If you increase the size of the LUN, you might need to close and reopen the computer management MMC (`compmgmt.msc`) before the increased LUN size becomes visible in the Disk Management snap-in.

3.2 Backup and Recovery Operations

SnapManager backup uses Snapshot copy functionality to create online, read-only copies of databases.

For Initial Configuration

Before using SnapManager to back up and restore a SQL Server database, you must use the SnapManager Configuration Wizard to migrate the databases and transaction logs from your SQL Server instances to the LUNs that you configured on your storage system with SnapDrive.

To View or Change the Database Configuration

After the initial configuration, you can rerun the Configuration Wizard at any time to review or make changes to your SQL Server data store configuration.

For detailed SnapManager for SQL Server backup, restore, and clone steps, refer to [Solution Operation Details](#) in the appendix.

4 Conclusion

Microsoft SQL Server is a powerful and cost-effective database solution that has been deployed to meet a wide variety of enterprise and departmental needs. However, there are challenges that enterprises

commonly face in delivering SQL Server databases and related applications. With this solution, customers are able to:

- Achieve increased storage efficiency. Our storage technologies include many efficiency and data protection features to help customers reduce the complexity and expense of their storage infrastructure.
- Reduce storage and data management costs through intelligent management of a unified architecture that supports iSCSI, Fibre Channel, and Fibre Channel over Ethernet SAN as well as SMB (CIFS).
- Increase availability and provide comprehensive data protection for SQL Server data and applications through nondisruptive operations on NetApp storage technology. Our support for Microsoft database availability groups along with SnapMirror and MetroCluster™ software enables business continuity for most any customer scenario. Clustered Data ONTAP can help minimize the pain of data migration that is required when performing a noninplace upgrade from earlier versions of SQL Server to SQL Server 2012.
- Simplify data manageability as customers consolidate SQL Server systems, all of which can be supported on the same shared infrastructure with seamless scaling.
- Accelerate the deployment, testing, and rollout of business applications based upon SQL Server data through automated rapid provisioning. Our storage efficiency technologies, rapid cloning, and clone lifecycle management capabilities make it possible to simultaneously support many development and testing environments without breaking the budget.

NetApp with Microsoft SQL Server 2012 is simple to deploy and manage, and it offers powerful capabilities that are designed to satisfy your growing database needs, with exceptional performance at the lowest cost of ownership.

Appendix A: Installation and Configuration Details

Detailed Steps to Install SnapDrive for Windows

Downloading SnapDrive for Windows

Download SnapDrive for Windows from:

<http://support.netapp.com/NOW/cgi-bin/software?product=SnapDrive&platform=Windows>

Consult your NetApp systems engineer to acquire a license to start using SnapDrive capabilities.

Install SnapDrive for Windows

Install SnapDrive for Windows on all the Microsoft SQL Server nodes to provision and manage storage LUNs. Before setup begins, verify the compatibility of all hardware and software involved by using the [NetApp Interoperability Matrix Tool](#).

To install SnapDrive for Windows (SDW), complete the following steps:

1. Browse to the location of the SnapDrive installation package and double-click the executable file to launch the SnapDrive Installation Wizard.
2. On the Welcome to the SnapDrive Installation Wizard page, click Next.
3. Read and accept the license agreement and click Next.
4. On the SnapDrive License page, select the type of licensing to use.
5. Enter the license key. If host-side licensing is used, enter the license key per server. Click Next.

Note: When storage system licensing is selected, SnapDrive can be installed without entering a license key. SnapDrive operations can be performed only on storage systems that have a SnapDrive or SnapManager license installed.

Note: With cluster-based systems, the storage system licensing for SnapDrive is bundled with the other SnapManager product licenses. The bundle is a single license called the SnapManager_suite license.

6. On the Customer Information page, enter the user name and organization name. Click Next.
7. On the Destination Folder page, select a host directory in which to install SnapDrive.

Note: By default, this directory is C:\Program Files\NetApp\SnapDrive\.

8. If the VMware ESX® guest OS is detected, the Installation Wizard prompts for the IP address and a user name with the appropriate vCenter™ or ESX server privileges. On the VirtualCenter or ESX Server Web Service Credentials screen, type the IP address of the vCenter or ESX server and the user name and password for SnapDrive to authenticate for web service. To use vMotion®, use vCenter. Selecting Enable VirtualCenter or ESX Server Settings enables SnapDrive to use RDM pass-through LUNs. Select this option to use RDM pass-through disks. By default, this option is not selected.

SnapDrive® - Installation Wizard

VirtualCenter or ESX Server Web Service Credentials

Specify account information for the installed services.

☒ Enable VirtualCenter or ESX Server Settings

Enables LUN provisioning and Snapshot copy management support with VMware ESX Server Guest OS using FC HBAs or ESX iSCSI(RDM) initiators. Specify VirtualCenter or ESX Server user account username and password. Ensure that the specified account is a member of the VirtualCenter or ESX Server local root group.

IP address / Name:
10.238.162.99

User Name:
administrator

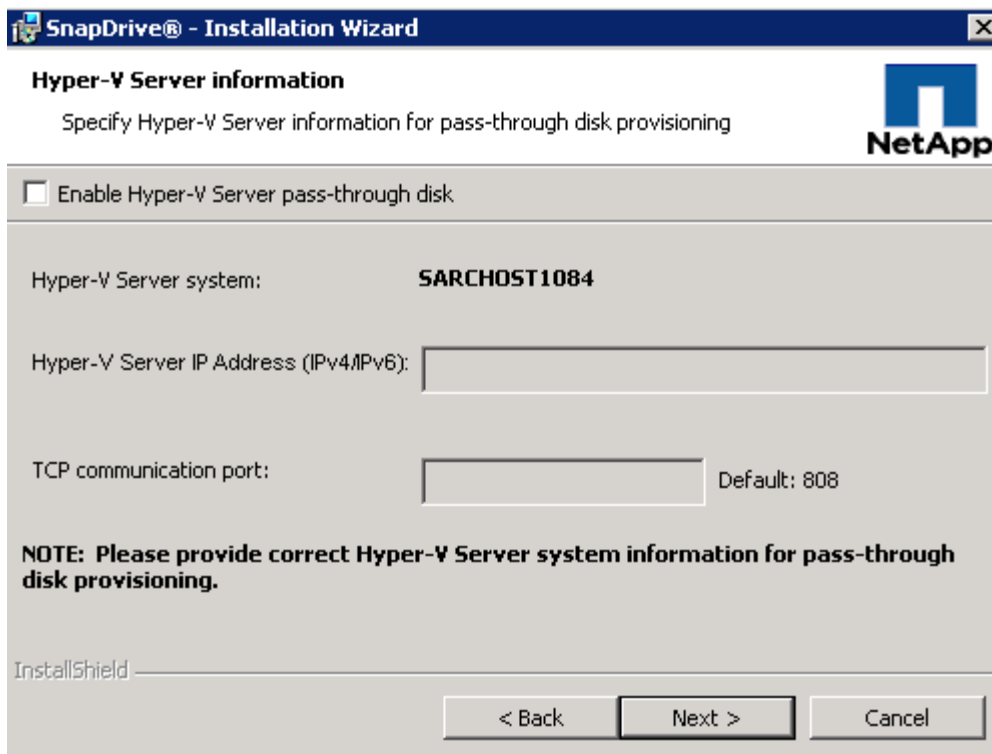
Password:
••••••••

Confirm Password:
••••••••

InstallShield

< Back Next > Cancel

If the Hyper-V guest OS is detected, the Installation Wizard prompts for the IP address and a user name with the appropriate Hyper-V server privileges. Type the IP address of the Hyper-V server.



SnapDrive® - Installation Wizard

Hyper-V Server information
Specify Hyper-V Server information for pass-through disk provisioning

☐ Enable Hyper-V Server pass-through disk

Hyper-V Server system: **SARCHOST1084**

Hyper-V Server IP Address (IPv4/IPv6):

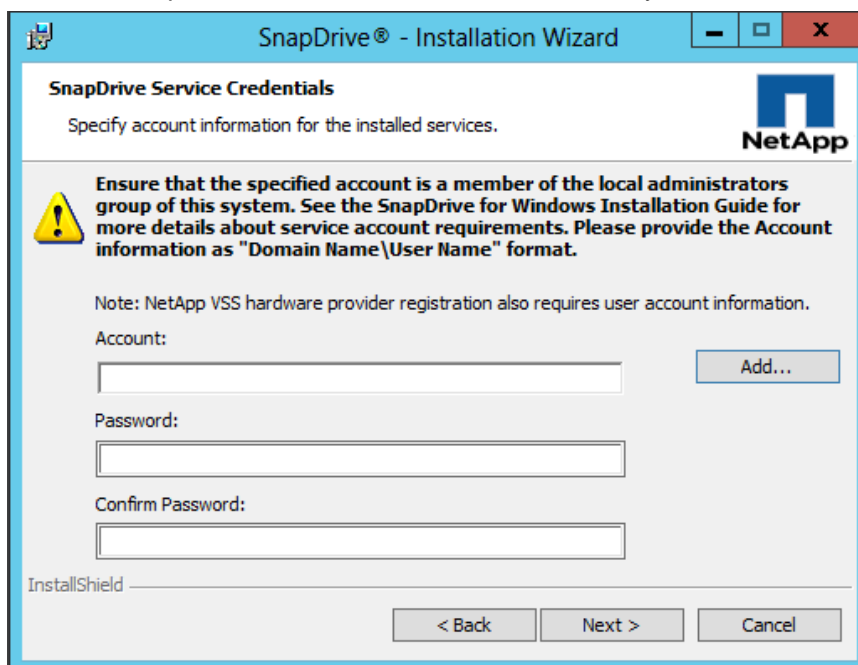
TCP communication port: Default: 808

NOTE: Please provide correct Hyper-V Server system information for pass-through disk provisioning.

InstallShield

< Back Next > Cancel

9. On the SnapDrive Service Credentials page, enter the account credentials. Alternatively, click Add to select a specific user account from Active Directory. Click Next.



SnapDrive® - Installation Wizard

SnapDrive Service Credentials
Specify account information for the installed services.

Ensure that the specified account is a member of the local administrators group of this system. See the SnapDrive for Windows Installation Guide for more details about service account requirements. Please provide the Account information as "Domain Name\User Name" format.

Note: NetApp VSS hardware provider registration also requires user account information.

Account:

Password:

Confirm Password:

InstallShield

< Back Next > Cancel

Note: The specified account must be a member of the local administrators group on this system.

10. On the SnapDrive Web Service Configuration page, keep the default port settings and click Next.

SnapDrive® - Installation Wizard

SnapDrive Web Service Configuration
Specify SnapDrive Web Service Configuration

SnapDrive Web Service Tcp/Ip Endpoint (Port)

SnapDrive Web Service HTTP Endpoint (Port)

SnapDrive Web Service HTTPS Endpoint (Port)

InstallShield

< Back **Next >** Cancel

11. On the Preferred Storage System IP Address page, you can specify the IP address you want to use to communicate with the storage system. Click Next.

SnapDrive® - Installation Wizard

Preferred Storage System IP Address
Configure SnapDrive to use a preferred IP Address

☐ Enable preferred storage system IP Address

**Configure SnapDrive to use a preferred IP Address for management traffic.
If storage system has only one interface, setting a preferred IP Address can prevent issues if more interfaces are added later.**

Storage System Name:

Preferred IP Address:

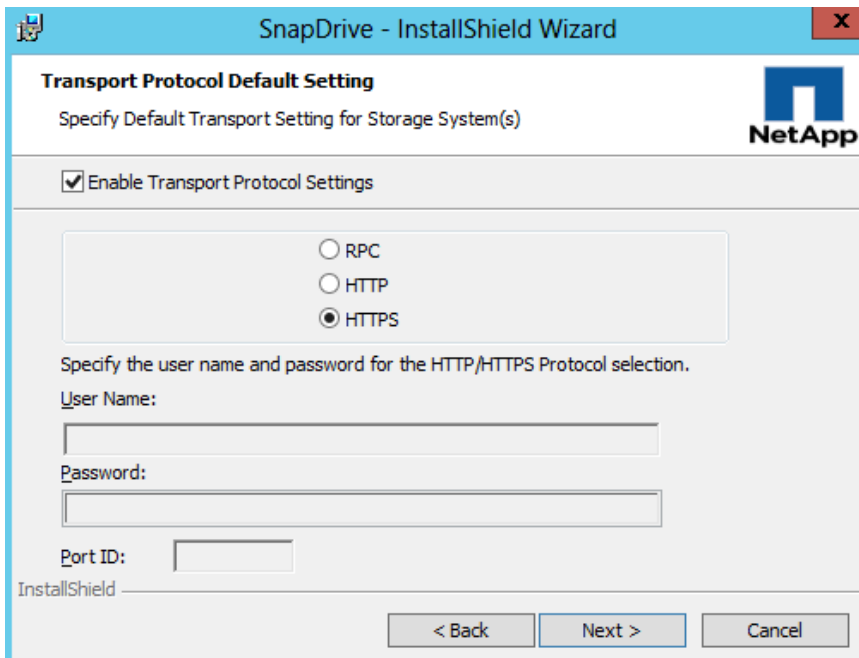
InstallShield

< Back **Next >** Cancel

12. In the Transport Protocol Default Setting page, select HTTPS and click Next.

Note: NetApp recommends using HTTPS. The HTTPS protocol allows the use of the Data ONTAP interface for all interactions between the storage system and host, including sending passwords securely.

Note: The RPC protocol is not supported for SnapDrive when clustered systems are used.



13. On the OnCommand configuration page, clear the Enable Protection Manager Integration checkbox and click Next.

Note: Protection Manager can be configured after the SnapDrive installation is complete if required in the customer environment.

14. On the Server Information page, clear the Configuration Option checkbox and verify that all of the fields are unavailable. Click Next.

Note: VMware integration and pass-through disk setup for Hyper-V can be configured after the SnapDrive installation is complete. Refer to [SnapDrive Documentation](#) for enabling and disabling vCenter or ESX logon from SnapDrive MMC if using VMware or to enable Hyper-V server pass-through disk if using Hyper-V configuration.

15. On the Ready to Install page, click Install.
16. When the SnapDrive Installation Completed page is displayed, click Finish.

Set Transport Protocol Settings in SDW

To configure transport protocol settings in SDW, complete the following steps:

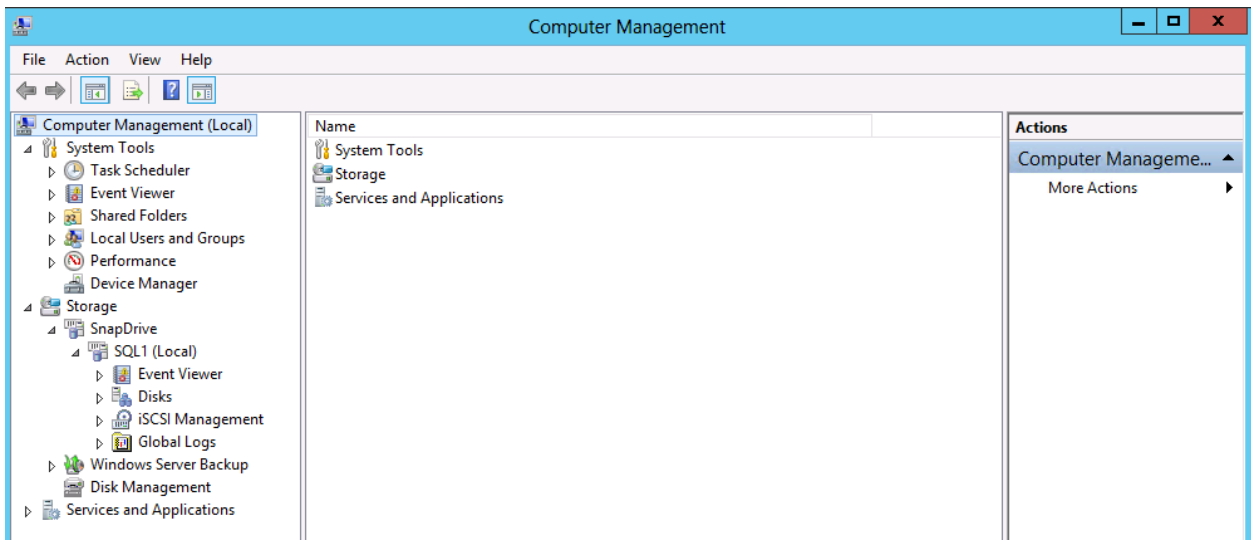
1. Log in to the host system and launch SDW.
2. From the SnapDrive console, right-click the host name and select Transport Protocol Settings.
3. From the Storage Systems tab, click Add.
4. In the Add Storage System dialog box, add the IP address of the SVM and click OK.
5. Verify that the SVM has been successfully added to the SnapDrive host.

Accessing and Managing SnapDrive for Windows

SnapDrive for Windows can be managed from the Microsoft Windows Computer Management MMC console.

From the Start menu, select Control Panel > Administrative Tools > Computer Management. Or select Start > Run, enter `compmgmt.msc`, and click OK.

In the left panel, under Storage, select SnapDrive to list the SnapDrive options.



Enable vCenter or ESX Logon from SDW

1. In the SnapDrive for Windows pane, select the instance of SnapDrive for which you want to enable vCenter or ESX logon. From the menu choices at the top, navigate to Action > vCenter Server or ESX Server Login Settings.
2. The vCenter Server or ESX Server Log On window is displayed.
3. To enable vCenter or ESX logon, in the vCenter Server or ESX Server Log On window, select the "Enable vCenter Server or ESX server settings" checkbox.
4. Type the IP address or host name, user name, and password for the vCenter or ESX instance to which you want to log in.
5. Click OK.

Create LUNs in VMware Environments

You can use SnapDrive to create LUNs in the VMware environments.

Before creating datastores, you must install and configure any adapters that your storage requires. Rescan the adapters to discover newly added storage devices.

Create an RDM LUN on a Guest OS

A raw disk mapping (RDM) can be used to present a LUN directly to a virtual machine from a SAN. You can use SnapDrive to create FC, iSCSI, or ESX iSCSI accessed RDM LUNs on a guest OS.

To add an RDM to a virtual machine, this VM needs to be first shut down.

1. Log in to the VMware vSphere® Client and select the host from Hosts and Clusters in the Inventory panel.
2. Right-click that selected host and click Edit Settings.
3. On the Hardware tab, click Add and choose Hard Disk.
4. Select Raw Device Mapping (RDM) on the page for Select a Disk. Click Next.
5. On the Select Target LUN page, choose the appropriate LUN. Click Next.
6. Choose "Store with virtual machine" or if you want store the link to the RDM in a specific datastore. Click Next.

7. On the Compatibility Mode page, select Physical compatibility mode, which allows the VM to pass SCSI commands directly to the storage system LUN. This allows it to leverage SAN-specific features such as interaction with the SAN's own Snapshot copy functions.
8. For Advanced Options, choose a virtual device node that is on a different SCSI bus than the current virtual disks. The RDM must be located on a separate SCSI controller. Click Next.
9. Confirm settings and click Finish on the Ready to complete screen.
10. A new SCSI controller and hard disk are added to the virtual machine configuration.

Note: Upon VM boot, the OS checks for new disk and formats/mounts the disk.

Create VMDK on VMFS Datastore

1. Log in to the VMware vSphere Client and select the host from Hosts and Clusters in the Inventory panel.
2. Click the Configuration tab and click Storage in the Hardware panel.
3. Click Datastores and click Add Storage.
4. Select the disk/LUN storage type to create a datastore on a Fibre Channel, iSCSI, or local SCSI disk, or mount an existing VMFS volume. Click Next.

Note: Adding a datastore on FC or iSCSI will add this datastore to all hosts that have access to the storage media.

5. Select a device to use for your datastore and click Next.

Note: Select the device that does not have a datastore name displayed in the VMFS Label column. If a name is present, the device contains a copy of an existing VMFS datastore.
6. If the disk you are formatting is blank, the Current Disk Layout page automatically presents the entire disk space for storage configuration. If the disk is not blank, review the current disk layout in the top panel of the Current Disk Layout page and select a configuration option. Use all available partitions from the bottom panel. Click Next.
7. In the Properties page, enter a datastore name and click Next.
8. If needed, adjust the file system and capacity values. By default, the entire free space on the storage device is available. Click Next.
9. In the Ready to Complete page, review the datastore configuration information and click Finish.

Create VMDK on NFS Datastore

1. Log in to the VMware vSphere Client and select the host from Hosts and Clusters in the Inventory panel.
2. Click the Configuration tab and click Storage in the Hardware panel.
3. Click Datastores and click Add Storage.
4. Select the Network File System as the storage type and click Next.

Note: Adding a datastore on FC or iSCSI will add this datastore to all hosts that have access to the storage media.

5. Enter the server name, the mount point folder name, and the datastore name.
6. (Optional) Select Mount NFS read only if the volume is exported as read only by the NFS server. Click Next.
7. In the Network File System Summary page, review the configuration options and click Finish.

Creating LUNs for Hyper-V

Before starting to provision LUNs by using SnapDrive, confirm that the FCP or iSCSI service is started on the storage system.

Note: You might need to license FCP or iSCSI protocol based on your requirements. For license details, consult your NetApp systems engineer or your NetApp software subscription package.

Two types of LUNs can be provisioned to the host by using SnapDrive for Windows:

- Dedicated
- Shared

Dedicated LUNs are dedicated to the server to which they are connected or mapped. Shared LUNs are used with Microsoft Cluster Service (MSCS). SnapDrive for Windows is cluster-aware and allows all the cluster nodes to connect to a single LUN when a shared LUN is provisioned.

1. Open the Computer Management Console.
2. Expand SnapDrive and expand the server name.
3. Right-click Disks and select Create Disk.
4. Enter the storage system name or IP address in the Create Disk Wizard. If you have already added the storage system in the storage systems management window, you can select the storage system from the drop-down list.
5. Select the volume where this LUN will be hosted. Enter a LUN name in the LUN Name field and enter a meaningful description for the LUN. Click Next.
6. Select Dedicated in the LUN type panel and click Next.
7. In the Select LUN Properties window, select a drive letter or mount point.
8. Select Limit or Do Not Limit for the option “Do you want the maximum disk size to accommodate at least one Snapshot copy?” In this case, we selected Do Not Limit to make the best use of thin provisioning. Enter the size of the disk to be created and click Next.
9. In the Select Initiators window, select the initiators. If you need to achieve multipathing, select all the initiators. The selected initiators must be of the same protocol. (A selection cannot have one FC initiator and one iSCSI initiator.) Click Next.
10. In the Select Initiator Group Management window, select Automatic and then click Next.
11. Click Finish to create the SnapDrive provisioned LUN.
12. From the SnapDrive GUI, you can locate the drive that you just created.

Detailed Steps to Install SnapManager for SQL Server

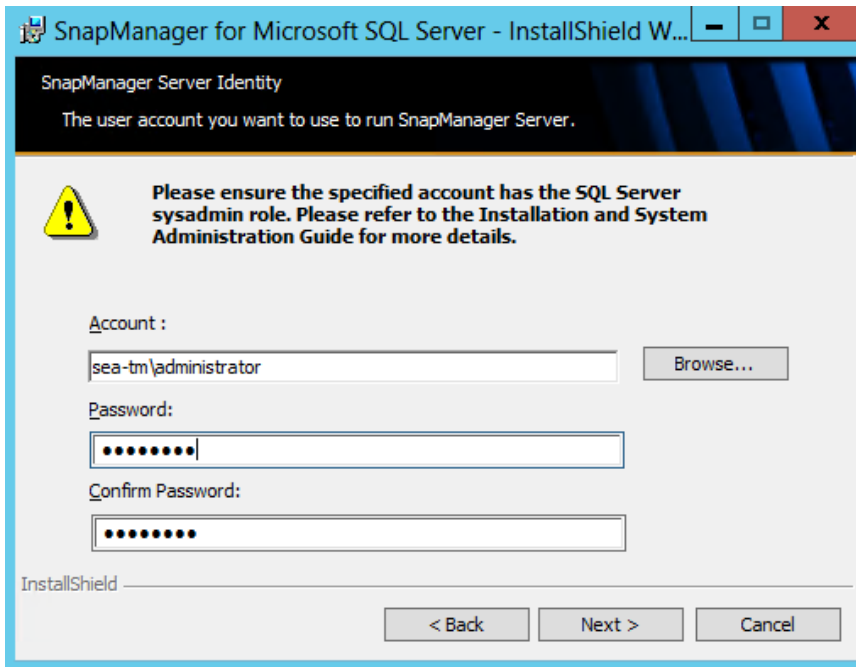
Install and Configure SnapManager for SQL Server

To install SMSQL, complete the following steps:

1. Download the SMSQL software installation package from the NetApp Support site [Software Downloads](#) page.
2. From the host, right-click the SMSQL installation package and select Run as Administrator to start the SMSQL installation.
3. On the Welcome page of the SnapManager for Microsoft SQL Server – InstallShield Wizard, click Next.
4. Specify the following customer information and click Next:
 - a. User name
 - b. Optional: information about the organization
 - c. License type

Note: If a storage-based SMSQL license is used, select the Per Filer license type. If a host-based SMSQL license is used, select the Per Server license type and enter the license key string.
5. Accept the default folder location for the installation or click Change to specify an alternate folder location. Click Next.

6. On the SnapManager Server Identity page, complete the following steps and click Next:
 - a. Enter a user name for the SMSQL service account. To select a different service account, click Browse.
 - b. Enter and confirm a password for the SMSQL service account.



7. Click Install.
8. Click Finish to complete the installation.

Add Server in SnapManager for SQL Server

1. Right-click SnapManager for SQL Server and select Add Servers to be Managed.
2. Enter SQL1 and click Add.

Appendix B: Build the SQL 2012 Availability Group for High Availability and Disaster Recovery

This section discusses the steps required to create a SQL 2012 AlwaysOn availability group for a local high-availability and remote disaster-recovery solution with this design pattern using nonshared storage for each SQL Server instance.

Table 12 lists the replica settings configured.

Table 12) Replica settings.

Data Center	Replica	Role	Availability Mode	Failover Mode
Primary data center	SQL 1	Primary	Synchronous commit	Automatic
Primary data center	SQL 2	Secondary	Synchronous commit	Automatic

Data Center	Replica	Role	Availability Mode	Failover Mode
Primary data center	SQL 3	Secondary	Synchronous commit	Manual
Disaster recovery data center	SQL 0	Secondary	Asynchronous commit (but a secondary synchronous replica is permitted; consider the network latency between the data centers and its effect on performance to the application)	Forced manual

Failover Clustering Feature

Windows Server Failover Clustering (WSFC) provides infrastructure features that support the high-availability and disaster-recovery scenarios for Microsoft SQL Server. Add the Failover Clustering feature to SQL 1, SQL 2, and SQL 3 in the primary data center and SQL 0 located in the secondary data center. For the AlwaysOn availability group, the availability group and the availability group listener are registered as WSFC cluster resources.

To install the Failover Clustering feature, refer to [Install the Failover Clustering Feature](#) and [Understanding Requirements for Failover Clusters](#).

Note: Confirm that the Windows Server edition supports the Windows Server Failover Clustering feature.

Note: All the cluster nodes must be in the same Active Directory Domain Services domain.

Note: The domain user also needs to be a member of the local Administrators group and this domain account should also have administrator permissions on each cluster node and **Create Computer Objects** and **Read All Properties** permissions for the container used for the domain computer accounts. For more information, refer to the [Failover Cluster Step-by-Step Guide: Configuring Accounts in Active Directory](#).

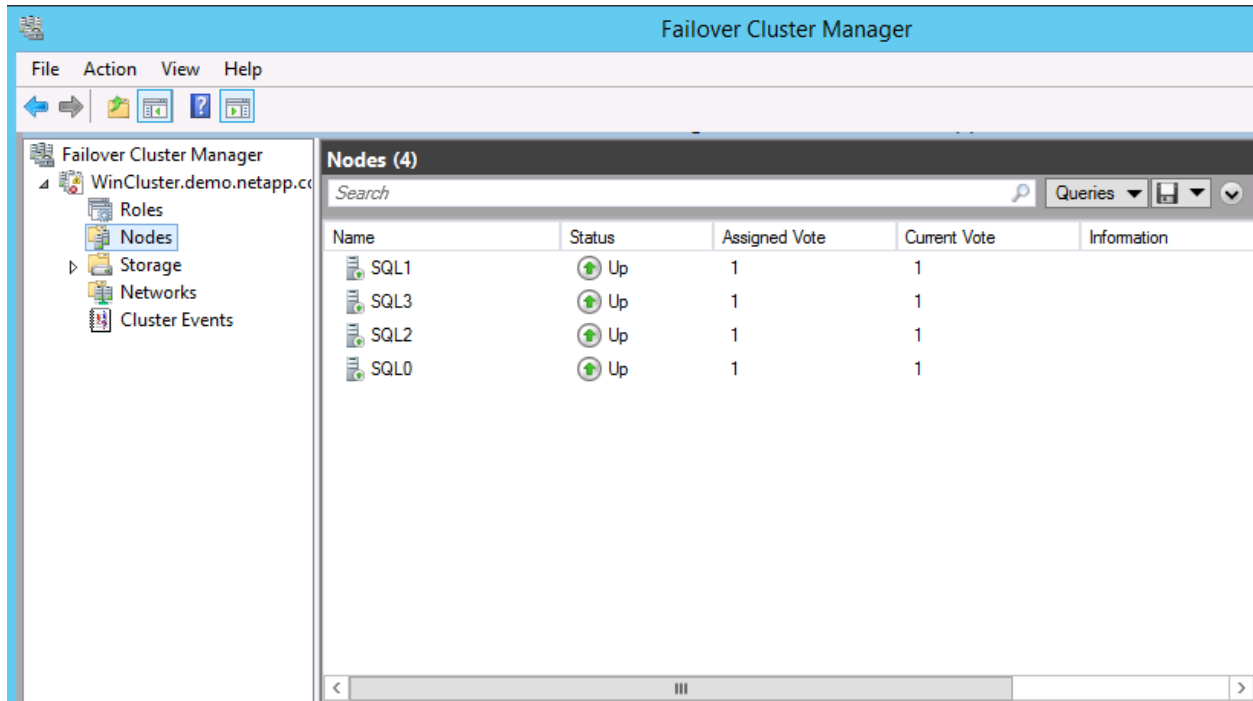
Create a New Windows Failover Cluster

Use the procedure listed in [Create a New Failover Cluster](#) for creating a new Windows cluster and add nodes SQL 1, SQL 2, SQL 3, and SQL 0 to this cluster.

Note: Confirm that the Remote Registry and Server services are started on each node to successfully add remote nodes to the Windows cluster.

The Failover Cluster Manager automatically connects to the cluster when the wizard finishes, as shown in Figure 3.

Figure 3) Nodes added to Windows Server Failover Cluster (WSFC).



To Set Quorum Mode Configuration

WSFC uses a quorum-based approach for monitoring overall cluster health and maximizing node-level fault tolerance. WSFC supports four quorum models. However, only two quorum models (Node Majority and Node and File Share Majority) apply when using a nonshared-storage solution in which there is a total of three nodes in the primary data center. Hence, you can use the Node Majority quorum model (with one vote assigned to each node in the primary data center and zero votes to the node in the DR data center). In that case, when you have two nodes in the primary data center, use the Node and File Share Majority quorum model with a file share witness in which the remote file share is also configured as a voting witness and connectivity from any node to that share is also counted as an affirmative vote.

The section “To change the quorum configuration in a failover cluster by using failover cluster snap-in” in the [Failover Cluster Step-by-Step Guide: Configuring the Quorum in a Failover Cluster](#) shares the steps to set Node Majority as the quorum model for this Windows cluster.

Set NodeWeight Property

The NodeWeight property of the WSFC node represents the vote for that particular node. Adjust each node's NodeWeight setting (a value of 0 or 1) so that the node's vote is not counted toward the quorum.

Following is the voting scheme:

- One vote to each node (SQL 1, SQL 2, and SQL 3) in the primary data center.
 - Zero votes to node SQL 0 in the disaster recovery data center because you do not want node SQL 0 to contribute to a decision to take the cluster offline when there is nothing wrong with the primary site.
1. To configure the NodeWeight from a node in a WSFC using Windows PowerShell®, on SQL 0, click Start > Administrative Tools > Windows PowerShell Modules.
 2. To set node SQL 0's vote to 0:

```
(Get-ClusterNode "SQL 0").NodeWeight=0
```

Note: Confirm that the other servers, SQL 1, SQL 2, and SQL 3, have the NodeWeight as 1.

For more information on adjusting node votes, refer to [Configure Cluster Quorum NodeWeight Settings](#).

Recommendations for quorum voting for AlwaysOn high-availability and disaster-recovery solutions are provided in the [“Recommended Adjustments to Quorum Voting”](#) section of [WSFC Quorum Modes and Voting Configuration](#).

Install a Standalone Instance of SQL Server 2012 on Each Cluster Node

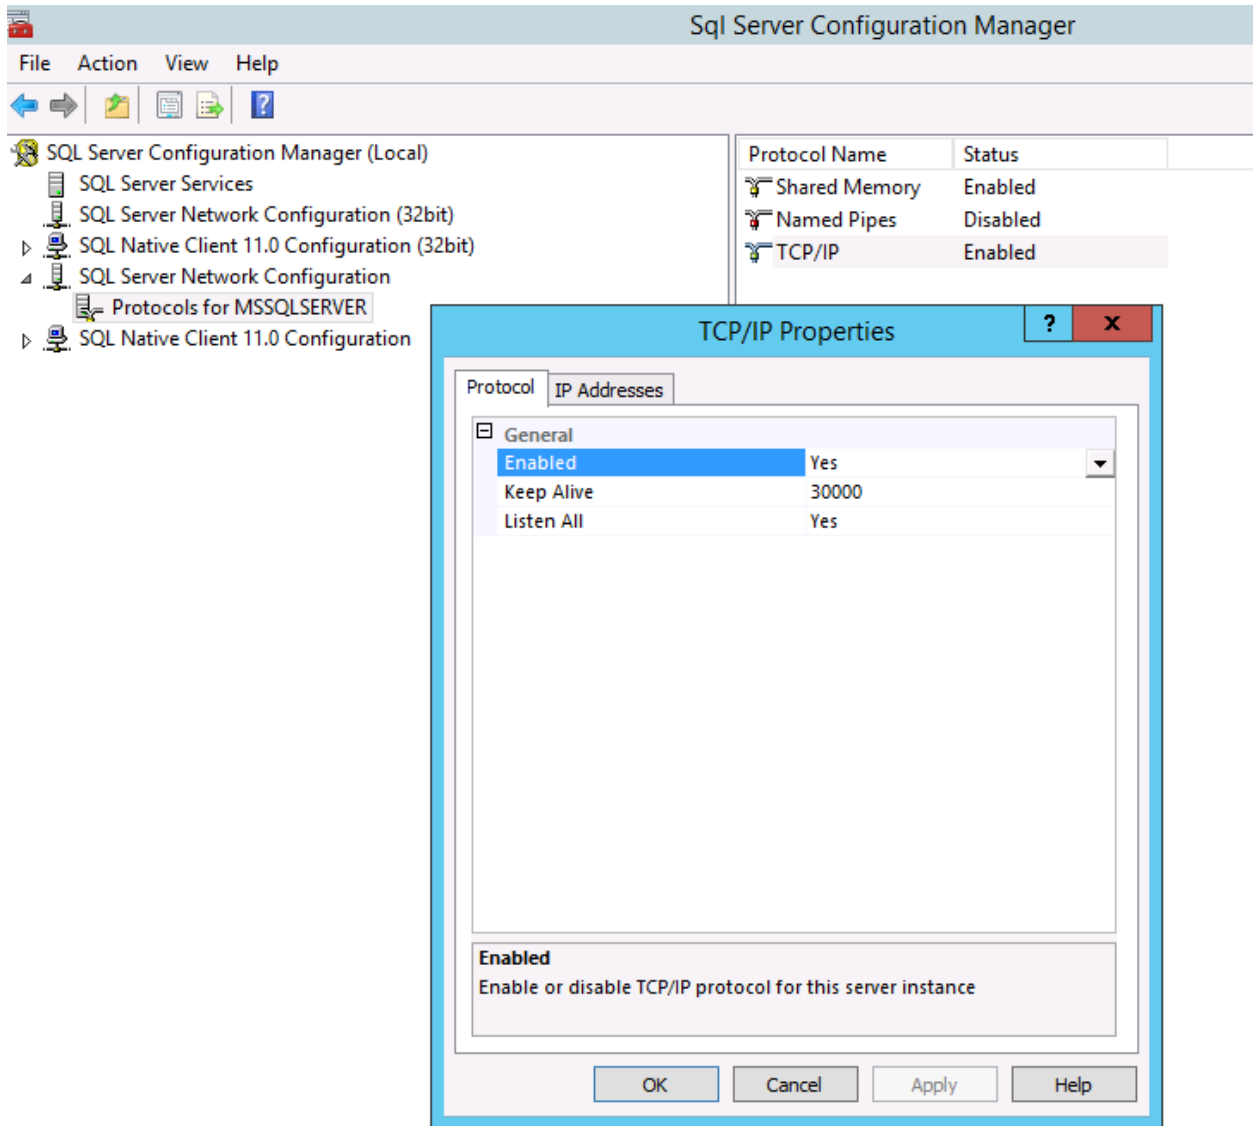
Prior to installing SQL Server 2012, verify that the necessary prerequisites have been met on each cluster node. For more information, refer to “Prerequisites, Restrictions, and Recommendations for AlwaysOn Availability Groups (SQL Server).” Then install SQL Server 2012 Enterprise edition to leverage the AlwaysOn availability group feature with the standalone instance. For more information, refer to “Installation for SQL Server 2012.”

Note: In the Database Engine Configuration step of the installation, NetApp recommends using identical file paths on each node.

Note: Also confirm that all of the SQL Server instances use the same SQL Server collation in order to host the availability group replicas.

By default, SQL Server does not accept remote connections. To change this default setting, click Start > Programs > Microsoft SQL Server 2012 > Configuration Tools > SQL Server Configuration Manager. Expand SQL Server Network Configuration, click Protocols for MSSQLSERVER, and double-click the TCP/IP entry. Change the Enabled option to Yes and click OK.

Figure 4) Enable TCP/IP protocol for MS SQL Server.

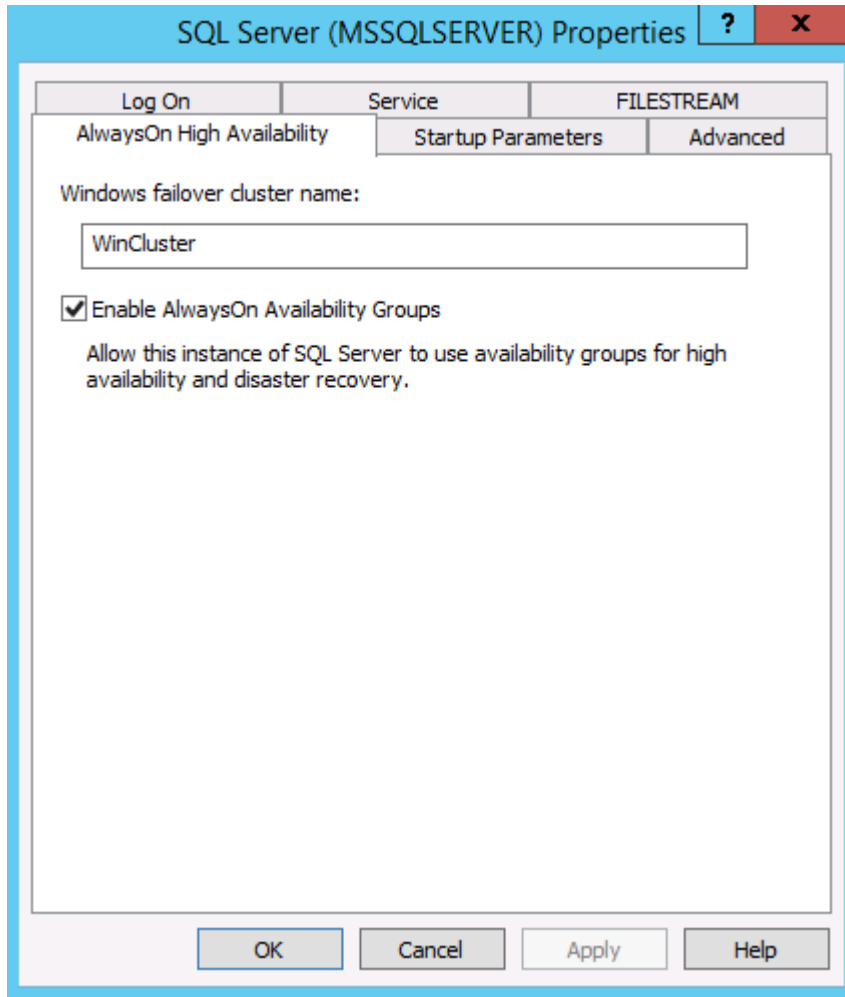


Enable AlwaysOn Availability Group Feature

To enable the AlwaysOn availability groups feature on each node, refer to [Enable and Disable AlwaysOn Availability Groups](#).

Note: Verify that the Windows failover cluster name field contains the name of the local failover cluster. If this field is blank, this server instance does not support AlwaysOn availability groups. Either the local computer is not a cluster node, the Windows failover cluster has been shut down, or this edition of SQL Server 2012 does not support AlwaysOn availability groups.

Figure 5) Enable AlwaysOn availability group feature.



Note: Restart the SQL Server (MSSQLSERVER) service for the changes to take effect.

Install SnapDrive for Windows on Each Node of the WSFC Cluster

SnapDrive for Windows (SDW) needs to be installed on each node of the WSFC cluster.

For more information, refer to the section “[Detailed Steps to Install SnapDrive for Windows](#)” in Appendix A: “Installation and Configuration Details.”

Install SnapManager for SQL on Each Node of the WSFC Cluster

SnapManager for SQL (SMSQL) needs to be installed on each node of the SQL 2012 WSFC cluster.

For more information, refer to the section “[Detailed Steps to Install SnapManager for SQL](#)” in Appendix A: “Installation and Configuration Details.”

Create Initial Database Backup

In order to create an availability group, the database must be in full recovery mode and at least one full backup must have been performed, for which we use SnapManager for SQL Server (SMSQL). SMSQL uses NetApp Snapshot technology to deliver near-instantaneous and space-efficient backups. After SMSQL has been installed, it needs to be configured to perform a full database and log backup. For

performing a full database backup, refer to the section “[Create SQL Server Database Backup](#)” in Appendix A: “Installation and Configuration Details.”

Create SQL 2012 Availability Group

Log on to the server SQL 1 that will host the primary replica and follow the procedure listed in [Use the Availability Group Wizard \(SQL Server Management Studio\)](#) to create a SQL 2012 availability group with databases added only to the primary instance SQL 1.

Then you can leverage NetApp SnapManager for SQL (SMSQL) to restore databases to secondary replica servers as follows.

Leverage NetApp Technology to Servers in the Same Primary Data Center

Typically, the backup and restore process for large databases during the creation of a SQL 2012 availability group can be time consuming, and the backup file takes up a large amount of disk space. SMSQL can provide faster backup and restore these databases (in no-recovery mode) by using NetApp Snapshot and FlexClone technologies on the secondary replicas.

For more details, refer to section 6.6, “Back Up and Restore Database to Replica Instance in Same Data Center,” in [TR-4106: Accelerate SQL Server 2012 AlwaysOn Availability Groups Deployment on NetApp Storage](#).

Repeat the preceding procedure on another secondary replica SQL 3.

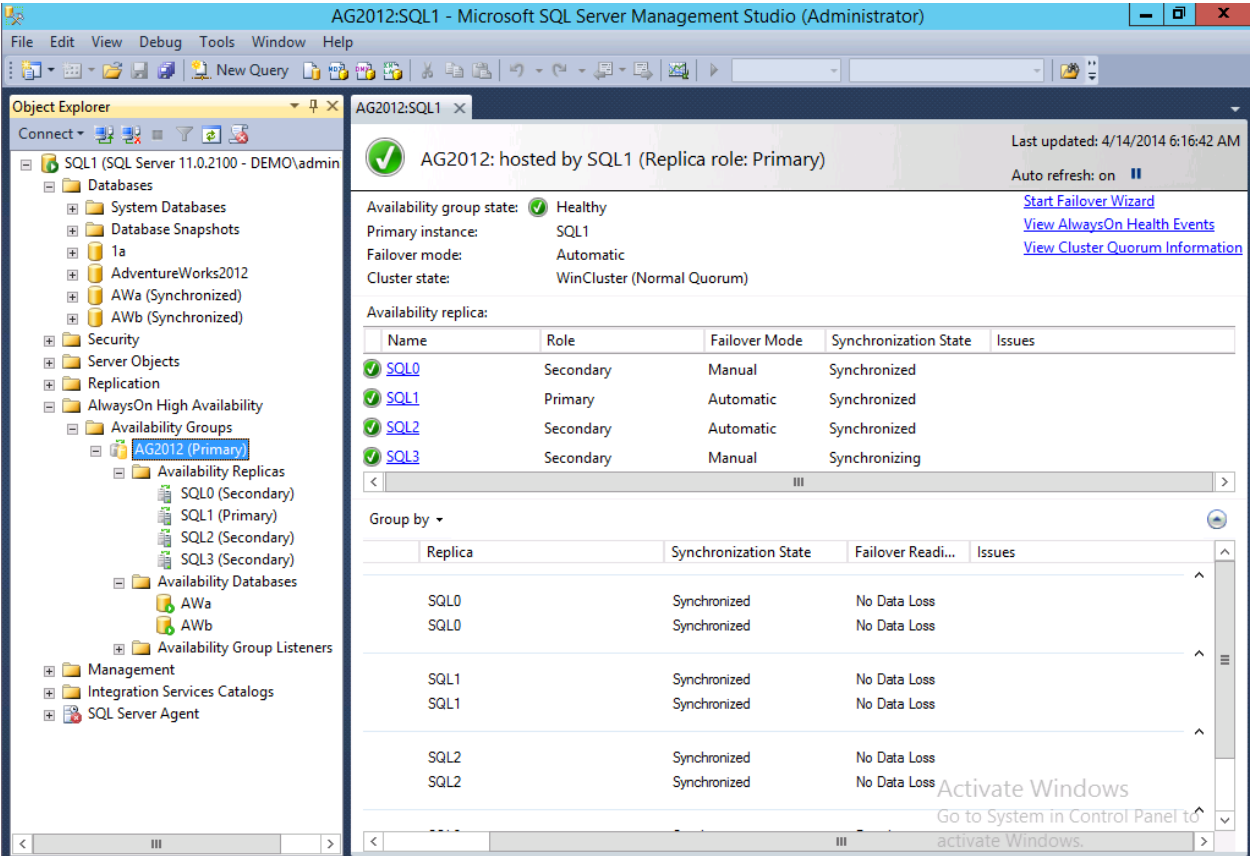
Leverage NetApp Technology to Server in Disaster Recovery Data Center

Network bandwidth limitations make it much more of a challenge to back up and restore databases between SQL Servers in separate data centers. SMSQL and NetApp SnapMirror make this task faster and simpler.

For complete details of the steps that need to be performed on SQL 0, refer to section 7, “Back Up and Restore Database to Replica Instance in Different Data Centers,” in [TR-4106: Accelerate SQL Server 2012 AlwaysOn Availability Groups Deployment on NetApp Storage](#).

Note: SnapMirror relationships of the database volumes (database, log, and SnapInfo volumes) must be created.

Figure 6) SQL 2012 availability group "AG2012" with nodes SQL 1, SQL 2, SQL 3, and SQL 0 added.



Create Availability Group Listener

In order to have fast application failover access to SQL Server, availability group listener must be created on the current primary node. Refer to [Create or Configure an Availability Group Listener \(SQL Server\)](#).

Test and Validation Details

Basic Tests to Validate the Solution

This section outlines the test procedures for the basic tests that should be performed to validate the deployment.

Connect to SQL Server Database Engine	
Task	<p>To connect to the SQL Server instance using the SQL Server Management Studio:</p> <ol style="list-style-type: none">1. Log in to Windows as a member of the Administrators group.2. Go to the Start menu > All Programs > Microsoft SQL Server 2012 > SQL Server Management Studio.3. In the Server name box, type the name of the instance of the database engine. For the default instance of SQL Server, the server name is the computer name.4. Click Connect.

Database Access in SQL Server Management Studio

Task	<p>To log in to the SQL Server Management Studio and access the database, complete the following steps:</p> <ol style="list-style-type: none"> 1. Log in to Windows as a member of the Administrators group. 2. Go to the Start menu > All Programs > Microsoft SQL Server 2012 > SQL Server Management Studio. 3. In Object Explorer, expand your server instance; expand Database to confirm that you are able to view the databases.
------	---

Add a Windows Authentication Login into the SQL Server Management Studio

Task	<p>To log in to the SQL Server Management Studio and access the database, complete the following steps:</p> <ol style="list-style-type: none"> 1. Go to the Start menu > All Programs > Microsoft SQL Server 2012 > SQL Server Management Studio. 2. In Object Explorer, expand your server instance, expand Security, right-click Logins, and then double-click the user account used by SMSQL. 3. On the Server Roles page, click sysadmin. 4. Click OK.
------	---

SQL Server Database Migration Using the SMSQL Configuration Wizard

Task	<p>After SMSQL is installed, the SMSQL Server Management Console must be connected to the SQL Server instance installed on the SQL Server host, and then the SMSQL Configuration Wizard must be run to configure SMSQL for the instance.</p> <p>To configure and migrate the SQL Server databases to a LUN, complete the following steps:</p> <ol style="list-style-type: none"> 1. Log in to the SQL Server instance where SnapManager is installed. 2. In the SnapManager for SQL Server Management Console, right-click SnapManager for SQL Server and select Add Servers to be Managed. 3. In the Add SQL Instance to be Managed dialog box, specify the SQL Server instance as (local). Select Use Windows Authentication and click Add. 4. SnapManager displays a notification indicating that the SQL Server instance is not configured for SnapManager. Click OK. 5. On the Welcome page of the SnapManager for SQL Configuration Wizard, click Next. 6. Specify the following verification server values and click Next: <ol style="list-style-type: none"> a. Verification Server. Select a verification server. By default, the local SQL Server instance is used for verification. b. Connection. Select the Use Windows Authentication option. c. Assign a Drive Letter or Path to Access a Mounted LUN in Snapshot copy. Select a drive letter or NTFS mount point directory. By default, SMSQL creates the NTFS mount point directory <code>C:\Program Files\NetApp\SnapManager for SQL Server\SnapMgrMountPoint</code>.
------	--

Note: The verification server can use a remote SQL Server instance for verification. If a remote SQL Server instance is used for verification, it must be the same version and architecture as the local SQL Server instance. The remote SQL Server instance must also have the same version of SDW and SMSQL installed.

7. Because we have manually changed the location for the default mount point directory, you will receive a pop up notification when you click OK.
8. Review the location of the databases and the list of available LUNs. After all databases have been assigned to their desired locations, click Next.

Note: If a database must be moved to a LUN, select the database in the Database Selection pane and then select the LUN in the Disk Selection pane. Click <=> to associate them.

9. Select the appropriate SnapInfo directory type for the environment and click Next.

Note: By default, the Single SnapInfo Directory option is selected.

10. In the Disk Selection pane, select the LUN where the SnapInfo directory will be created. Click <=> to assign the SnapInfo LUN. After the SnapInfo LUN has been assigned, review the path shown in the SnapInfo Directory field and click Next.

11. For a SQL Server 2012 availability group, select Enable SnapManager Share and specify the share location. Click Next.

12. Select the migration options for database verification, source data deletion, and statistics update that are appropriate for the environment and click Next.

By default, the following options are selected:

- Run DBCC Physical Integrity Verification Before Migration
- Delete Copy of Migrated Database at Original Location
- Run UPDATE STATISTICS on Tables Before Detaching the Databases

Note: Selecting the Delete Copy of Migrated Database at Original Location option can result in the inability to roll back the migration of the database if there is a problem with the migration. Manually deleting the original SQL Server data files should be done after confirmation that the SQL Server databases are healthy and operational in their new location.

Note: Selecting the DBCC verification options or the updating statistics option increases the time required for the data migration to complete. For the fastest migration experience, NetApp recommends not enabling these options.

13. Set the iSCSI service as a dependency for all SQL Server services to help protect the SQL Server data if there is a problem with iSCSI. By default, SQL Server Instances is selected.

Note: If the iSCSI service is not running, SQL Server services will not start. If the iSCSI service is stopped, SQL Server services will also be shut down.

Note: If iSCSI is not being used, clear the option to set the SQL Server service dependency.

Click Next.

14. NetApp strongly recommends configuring the following automatic event

notifications:

- a. Select the Send E-mail Notification option and specify the following settings:

SMTP Server: Enter the fully qualified domain name (FQDN) or IP address of the SMTP relay server.

From: Enter the e-mail address of the sender. The default value of SMSQLAutoSender might not work with some SMTP relays. In that case, a dummy e-mail address can be used.

To: Enter the e-mail addresses of the recipients.

Subject: Customize a subject for the subject line in the e-mail notifications. The default value is SnapManager for SQL Server.

Note: For SMTP e-mail notifications to work properly, the host antivirus and firewall settings must be configured to allow SMSQL to send e-mail by using SMTP.

- b. Select the Log SnapManager Events to Storage System Syslog option.
- c. Select the Send AutoSupport Notification option.

Click Advanced to configure the advanced event notification settings. By default, the Send Operation Results Summary is selected. To reduce the number of e-mails generated by SMSQL, select the Only Send Notification When Operation Fails option.

Click OK and then click Next to continue.

15. On the Configure Monitoring and Reporting Settings page, specify the following monitoring and reporting settings based on the organizational requirements and click Next:

- a. Select the Enable Monitoring and Reporting option. This option is not selected by default.

Note: SMSQL issues a notification that the Monitoring and Reporting option is enabled. Click OK.

- b. Select which of the following operations will be monitored: backups, verification, and clone resync. These options are selected by default.
- c. Set the SMSQL reporting interval. The default value is one day.
- d. Set the SMSQL reporting time. The default value is 4:00 a.m.

16. Verify the settings configured for SMSQL. If any changes must be made, click Back until the page with the options that must be changed is reached. Click Finish.

17. On the Configuration Status page, click Start Now to start the SMSQL configuration.

18. After the SMSQL configuration is complete, SMSQL notifies the user whether the operation was successful or not. Click OK.

19. The configuration status is displayed. Click Close.

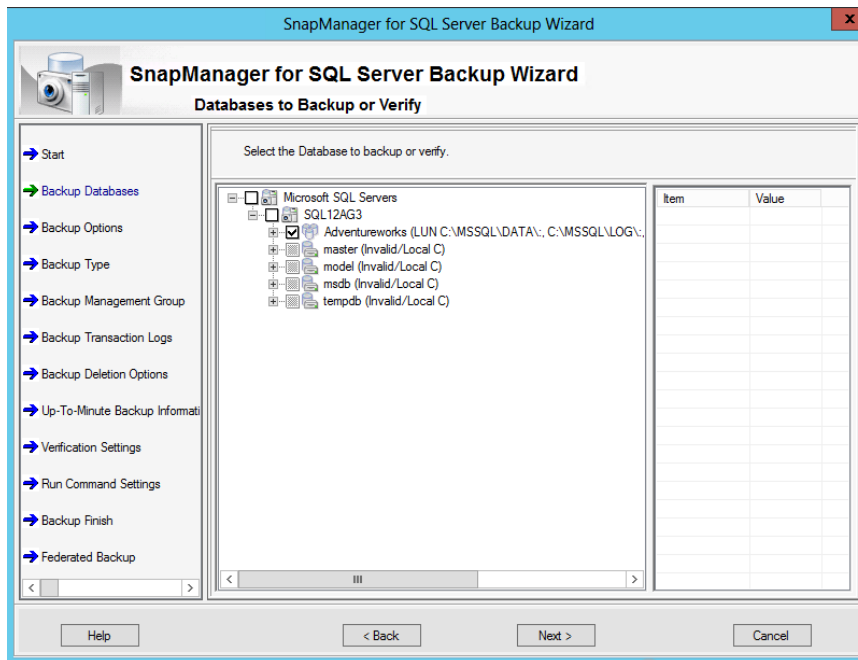
Note: If the configuration operation was not successful, note the error messages that are displayed and then review the SMSQL configuration logs and Windows application event logs.

Solution Operation Details

Create SQL Server Database Backup

To create a SQL Server database backup by using SnapManager for SQL Server (SMSQL), complete the following steps:

1. Log in to the SQL Server instance on which SMSQL is installed.
2. In the Actions pane, click Backup Wizard to start the SMSQL Server Backup Wizard.
3. On the Welcome page, click Next.
4. Select the database instance to be backed up and click Next.



5. Select the Back Up Databases and Transaction Logs (default) option and click Next.
6. Select the Continue with Selected Databases option to back up the database and click Next.
7. Select the backup type and click Next.
Note: By default, the Full option is selected.
8. Select the backup management group option and click Next.
Note: By default, the Standard Backup option is selected.
9. Select whether to automatically back up the transaction log after the full database backup.
Note: By default, Yes is selected.
Note: A transaction log backup does not occur for databases that are configured to use the simple recovery model.
10. Set the backup retention policy for the backup job by configuring the following options and click Next:
 - a. **Delete Full Backups.** This option is selected by default.
 - b. **In Excess of.** This option is also selected by default with a default value of 7.
 - c. **Older Than.** This option is not selected by default. It can be used to delete all backups older than the specified value (in days).

- Note:** Before setting the retention policy, verify that there is enough space in the SQL Server data and SnapInfo flexible (FlexVol) volumes to accommodate the desired number of Snapshot copies. Each FlexVol volume can contain up to a maximum of 255 Snapshot copies.
11. Set the retention period for transaction log backup data stored in the SnapInfo directory and click Next.

Note: By default, the Backups Generated in the Last (Days) option is selected, with a default value of 7.
 12. Select one of the database verification options.

Note: By default, Yes is selected.
 13. The Verification Setting button becomes active if verification is selected. Click the Verification Setting tab and configure the following fields:
 - a. **Server.** The local server is specified by default. When using a remote verification server, confirm that the versions of SQL Server, SnapManager for SQL Server, and SnapDrive are the same as those running on the local SQL Server instance. Also make sure that the remote verification server is connected to the NetApp storage on which the SMSQL Snapshot copies are stored.
 - b. **Connection.** Select the Use Windows Authentication option, which is selected by default.
 - c. **Mount in an Empty NTFS Directory.** This option is selected by default. The default mount point directory value is C:\Program Files\NetApp\SnapManager for SQL Server\SnapMgrMountPoint.
 14. Click the SnapMirror Options tab, and then click Verification on Destination Volumes to perform verification on a SnapMirror volume.
 15. Click the DBCC Options tab and select the Database Console Commands (DBCC) options based on operational requirements, then click OK.

Note: By default, the NO_INFOMSGS and PHYSICAL_ONLY options are selected.

Note: For more information about DBCC options, refer to the [Microsoft SQL Server documentation](#) on Microsoft TechNet.
 16. Select the SnapMirror and SnapVault options for the database backup and click Next. The default value for the Remote Backup Management Group option is Daily.

Note: SnapMirror must be licensed on both source and destination NetApp storage systems. The SnapMirror relationships for the database volumes must be initialized and healthy before any SnapMirror options are selected.

Note: SnapVault must be licensed, and Network Data Management Protocol (NDMP) must be enabled, on both source and destination NetApp storage systems. The SnapVault datasets must be in a compliant state before any SnapVault options are selected.

Note: SMSQL 7.0 supports native SnapVault with clustered Data ONTAP 8.2 through SnapDrive, which does not require Protection Manager. However, in 7-Mode, Protection Manager and DataFabric® Manager are still used.

Note: Verification on either a SnapVault or SnapMirror destination requires that a FlexClone license be installed on the secondary storage system. The remote verification server must have the same version of SQL Server, SMSQL, and SnapDrive as the local SQL Server instance. The remote verification server must also be connected to the secondary storage system.
 17. Click Backup Setting to configure and modify advanced backup settings.
 18. In the Backup Settings dialog box, from the Full Database Backup tab, configure the full database backup naming convention and database verification options.

Note: By default, the Use Unique (Time Stamp) Naming Convention option is selected.

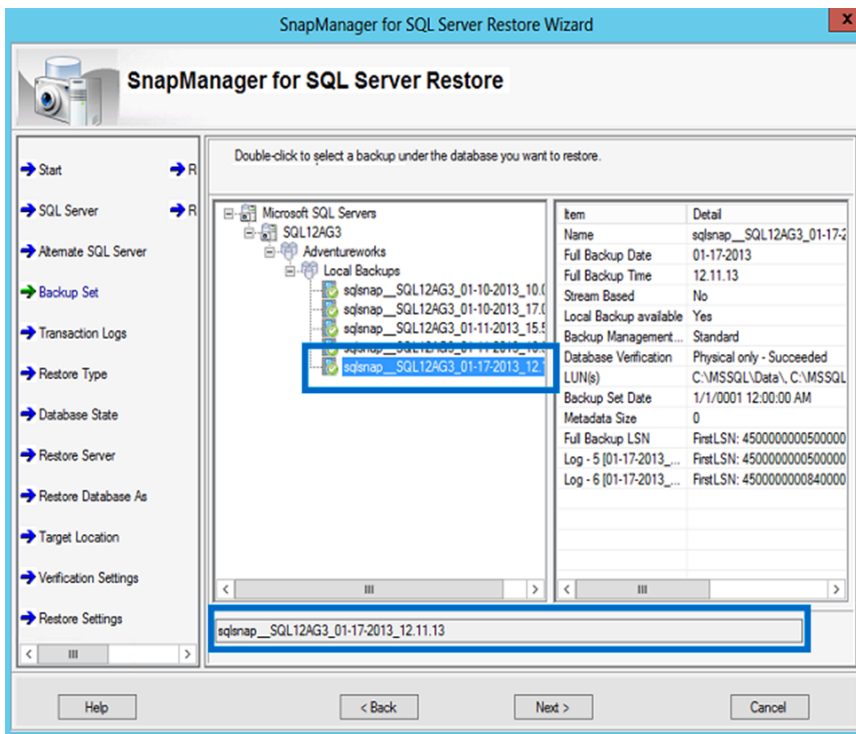
Note: NetApp recommends using generic naming conventions when using scripting with SMSQL backups.

19. Click the Transaction Log Backup tab and configure the transaction log backup options. After configuring the backup settings, click OK to save the configuration, and then click Next.
Note: The default values for this backup setting are displayed in this example.
20. Select whether to run a custom program or script with the current SMSQL operation.
Note: By default, No is selected.
21. Review the database backup settings. Click Schedule to schedule the backup job and then proceed to step 22, or click Finish to continue with the backup.
22. In the Schedule Job dialog box, specify the following values and then click OK:
 - a. **Schedule Job Name.** Enter the name of the backup. Select the Replace Job If It Exists option to overwrite a job in scheduled tasks.
 - b. **SQL Server Agent/Server Name.** Specify a SQL Server instance that will run the SMSQL backup job. The local server name is specified by default.
 - c. **Windows Scheduled Task (Job Will Be Created on This Local Machine).** Specify the user account and password for the job in scheduled tasks.
Note: NetApp recommends using the SQL Server Agent option because the SQL Server Agent is a scheduling engine that Microsoft SQL Server database administrators (DBAs) know how to use.
23. Click Finish to start the database backup operation.
24. In the Backup Status dialog box, click Start Now.
25. The Backup Status dialog box displays the status of the clone operation. After the backup operation is complete, a notification message is displayed. Click OK to close the message and then click Close.
Note: If the backup operation is unsuccessful, document any error message that is displayed. Review the Windows application event log and SnapManager backup report log for the cause of the failure.

Restore SQL Server Database

To restore the SQL Server database by using SMSQL, complete the following steps:

1. Log in to the SQL Server instance on which SnapManager is installed.
2. In the Actions pane, click Restore Wizard to start the SnapManager for SQL Server Restore Wizard.
3. On the Welcome page of the SnapManager for SQL Server Restore Wizard, click Next.
4. Select the backup type for the source of the database restore, and then click Next.
Note: By default, Restore SnapManager Backups that Were Created on the Same Server is selected.
5. On the Select Backup page, double-click the backup for restore so that the backup name is shown at the bottom of the dialog box. Click Next.



6. To select the log backups to be restored for each database, complete the following steps and then click Next:
 - a. Select the transaction log backups to be restored.
 - b. Select the databases for which to apply the transaction log restore.
 - c. Select the point-in-time settings for the transaction log restore.

Note: By default, the transaction log backups for the databases in the backup set are selected.

7. On the Database State After Restore page, select from the following options the one that is appropriate for your environment and click Next:
 - a. Leave the databases operational, but unavailable for restoring additional transaction logs

Note: This option is selected by default.

- b. Leave the databases nonoperational but available for restoring additional transaction logs
 - c. Leave the databases in read-only mode and available for restoring additional transaction logs

Note: The Undo File Directory option is used with this option.

8. Select a database to restore.

- a. To restore as a database with a different name from that of the original database, click the ellipsis (...) button beside Restore as Database.
 - b. The Individual Database Restore As dialog box opens. In the Restore as Database field, enter the database name to which you want the backup restored. This database name must not already exist on the SQL Server instance to which you will be restoring the database. Click OK to apply the change and close the dialog box.
 - c. To restore the database to an alternate location, click the ellipsis (...) button beside Restore to Other Location.

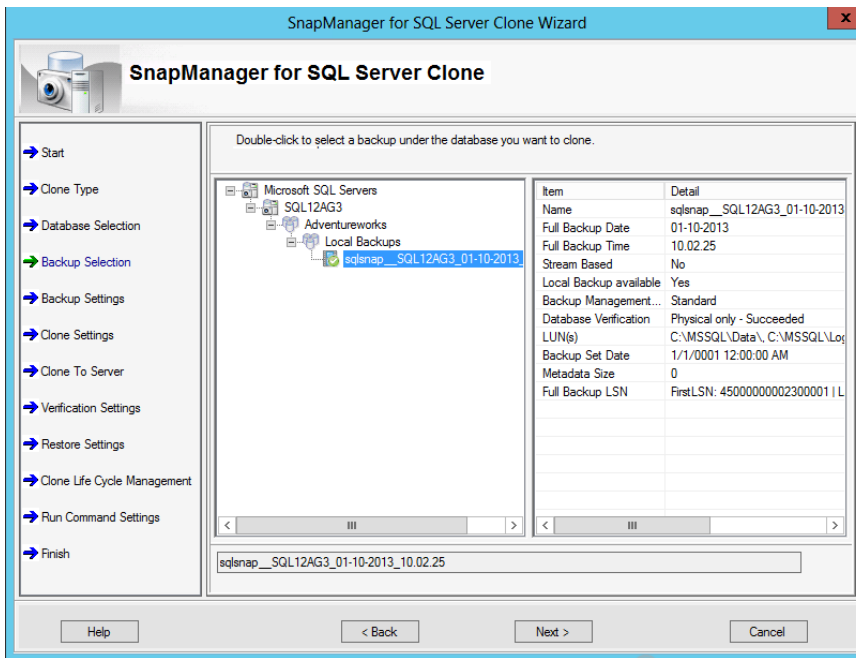
Note: To edit the location, select and modify the Restore to Other Location field for each row or click the ellipsis (...) button and browse for the location.

- Note:** If the alternate location does not have enough space, the restore will fail. If this occurs, delete the partially copied database files.
9. Select the destination of the SQL Server database restore and click Next.
Note: By default, the SQL Server instance of the database in the backup set is selected.
Note: Click the ellipsis (...) button to select an alternate instance as the destination for the restore.
 10. On the Restore Settings page, click Restore Setting to configure advanced restore options.
 11. If the database is still online, select Restore Databases Even if Existing Databases Are Online. Click OK to close the dialog box and then click Next.
Note: The Restore Databases Even if Existing Databases are Online option is not required if the database is not online.
 12. Select whether to run a custom program or script with the current SMSQL operation.
Note: By default, No is selected.
 13. Review the restore settings. Click Finish to continue with the restore operation.
 14. In the Restore Status dialog box, click Start Now.
 15. When the restore operation is complete, the Restore Status dialog box displays the status of the restore operation. Click OK to close the message and then click Close.
Note: If the restore operation is unsuccessful, document any error message that is displayed and review the SnapManager restore report log and the Windows application event log for additional information.

Clone SQL Server Database from SnapManager for SQL Server Backup

To create a clone of a SQL Server database from backup by using SnapManager for SQL Server (SMSQL), complete the following steps:

1. Log in to the SQL Server instance on which SnapManager is installed.
2. In the Actions pane, click Clone Wizard to start the SnapManager for SQL Server Clone Wizard.
3. On the Welcome page, click Next.
4. Select the clone type and click Next.
Note: By default, the Clone Existing Databases from Existing Backup Set is selected.
5. To select a backup to clone, expand the SQL server name and then the database to display the backup sets for that database. Double-click the backup name to load the backup information and click Next.



6. To select log backups to be applied for each database, complete the following steps and click Next:
 - a. Step 1: Select the transaction log backups to be restored.
 - b. Step 2: Select the databases for which to apply the transaction log restore.
 - c. Step 3: Select the point-in-time settings for the transaction log restore.

Note: By default, the transaction log backups for the databases in the backup set are selected.
7. Select a database to clone and click the ellipsis (...) button beside the Clone as Database field.
8. In the Restore as Database field of the Individual Database Restore As dialog box, enter the name of the database to which you want the backup restored. Click OK.

Note: This database name must not already exist on the SQL Server instance to which you are restoring the database.
9. Specify a server on which to clone the selected database and click Next:
 - a. **Clone to Server [Instance].** The default value is the local SQL Server instance.
 - b. **Access Snapshot LUN Option on Target Instance.** The default value is the SMSQL verification mount point directory (C:\Program Files\NetApp\SnapManager for SQL Server\SnapMgrMountPoint).

Note: If you select a remote SQL Server instance as the destination for the database clone, confirm that it has the same versions of SQL Server, SMSQL, and SnapDrive installed as the local SQL Server instance on which the backup was created. It must also have connectivity to the storage on which the SMSQL backup was created.
10. On the Restore Setting page, select the advanced clone restore settings options appropriate for your environment and then click Clone Restore Settings:
 - a. **Clone on Available SnapMirror Destination Volume.** This option is not available if the SQL Server data is not on SnapMirror volumes. Select this option to create the clone on the SnapMirror destination volume.
 - b. **Change Clone Database Paths Based on New Database Name.** Select this option to change the clone database path names to the clone database name.
11. In the Restore Settings dialog box, select the appropriate restore settings for your environment. Click OK to close the dialog box and then click Next.

Note: By default, Create Transaction Log Backup Before Restore is selected.

12. On the Clone Lifecycle Management page, specify the appropriate clone lifecycle management options and click Next:
 - a. **Clone Resynchronize Option.** This option resynchronizes the database clone with the parent database according to the specified schedule.
 - b. **Clone Auto Deletion Option.** This option automatically deletes the clone, based on the specified schedule.
13. To designate the state of the database after restore, select the appropriate option for your environment and click Next.

Note: By default, the Leave the Databases Operational, but Unavailable for Restoring Additional Transaction Logs option is selected.

Note: If the Leave the Database in a Read-only Mode and Available for Restoring Additional Transaction option is selected, the Undo File Directory field option also becomes available.
14. Select whether to run a custom program or script with the current SMSQL operation.

Note: By default, No is selected.
15. Review the clone restore settings. Click Finish to start the database clone operation.
16. In the Clone Status dialog box, click Start Now.
17. The Clone Status dialog box displays the status of the clone operation. After the clone operation is complete, a notification message is displayed. Click OK to close the message, and click Close to exit the Clone Status dialog box.

Note: If the clone operation is unsuccessful, document any error message that is displayed. Review the Windows application event log and SnapManager clone report log for the cause of the failure.

References

This section lists useful resources to assist in planning and managing your SQL Server storage environment:

- [NetApp Storage Systems](#)
- [Data ONTAP Documentation](#)

Additional documentation available from NetApp Support site:

- [NetApp SnapDrive for Windows](#)
- [SnapManager for Microsoft SQL Server \(SMSQL\)](#)
- [Microsoft SQL Server Customer Advisory Team: resources for complex enterprise SQL Server implementations](#)
- [Microsoft SQL Server Storage Engine Blog](#)
- [MSDN: Product documentation including SQL Server Books Online](#)
- [SQL Server Hardware and Software Requirements](#)

Microsoft SQL Server:

- [Description of disaster recovery options for Microsoft SQL Server](#)
- [INF: Disaster Recovery Articles for Microsoft SQL Server](#)
- [Disaster Recovery](#)
- [Introduction to Backup and Restore Strategies in SQL Server](#)
- [You cannot restore system database backups to a different build of SQL Server](#)

Version History

Version	Date	Document Version History
Version 1.0	May 2014	Initial release
Version 1.0.1	July 2014	Updated with complete deployment procedure

Refer to the [Interoperability Matrix Tool](#) (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®



www.netapp.com

© 2014 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, AutoSupport, DataFabric, Data ONTAP, FlexClone, FlexVol, MetroCluster, SnapDrive, SnapManager, SnapMirror, SnapRestore, Snapshot, and SnapVault are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Active Directory, Hyper-V, Microsoft, PowerShell, SQL Server, Windows, and Windows Server are registered trademarks of Microsoft Corporation. ESX, vMotion, VMware, and VMware vSphere are registered trademarks and vCenter is a trademark of VMware, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-4302-0714