Technical Report

Deploying SQL Server 2012 over SMB3 CA Shares on Clustered Data ONTAP

Marc Waldrop, Pat Sinthusan, NetApp

December 2013 | TR-4247

Abstract

NetApp delivers a storage environment with always-on availability by utilizing nondisruptive operations (NDO) in the clustered Data ONTAP® operating system. Microsoft® SQL Server® 2012 database files can now reside on continuously available (CA) shares because of a feature in clustered Data ONTAP 8.2.1. This means that Microsoft SQL Server does not go offline for routine upgrades and maintenance or add-ons. In addition, it completes hardware refreshes.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Overview

1.1 Clustered Data ONTAP

With the release of NetApp® clustered Data ONTAP, NetApp was the first to market enterprise-ready, unified scale-out storage. Developed from a solid foundation of proven Data ONTAP technology and innovation, clustered Data ONTAP is the cornerstone of virtualized shared storage infrastructures that are architected for nondisruptive operations over the lifetime of the system.

A helpful way to start understanding the benefits offered by clustered Data ONTAP is to consider server virtualization. Before server virtualization, system administrators frequently deployed applications on dedicated servers to maximize application performance and to avoid the instabilities often encountered when combining multiple applications on the same operating system instance. Although this design approach was effective it also had the following drawbacks:

- **It did not scale well.** Adding new servers for every new application is extremely expensive.
- **It was inefficient.** Most servers are significantly underutilized, meaning that businesses were unable to churn out the maximum benefit from their hardware investment.
- **It was inflexible.** Reallocating standalone server resources for other purposes was time consuming, staff intensive, and highly disruptive.

Server virtualization directly addresses all of the preceding limitations by decoupling the application instance from the underlying physical hardware. Multiple virtual servers can share a pool of physical hardware, which means that businesses can now consolidate their server workloads to a smaller set of more effectively utilized physical servers. In addition, the ability to transparently migrate running virtual machines across a pool of physical servers enables businesses to reduce the impact of downtime due to scheduled maintenance activities.

As with server virtualization, clustered Data ONTAP enables you to combine multiple physical storage controllers into a single logical cluster that can nondisruptively service multiple storage workload needs. With clustered Data ONTAP you can:

- Combine different types and models of NetApp storage controllers (known as nodes) into a shared physical storage resource pool (referred to as a cluster).
- Support multiple data access protocols (CIFS, NFS, Fibre Channel, iSCSI, and FCoE) concurrently on the same storage cluster.
- Consolidate various storage workloads to the cluster. Each workload can be assigned its own Storage Virtual Machine (SVM), which is essentially a dedicated virtual storage controller, and its own data volumes, LUNs, CIFS shares, and NFS exports.
- Support multitenancy with delegated administration of SVMs. Tenants can be different companies, business units, or even individual application owners, each with its own distinct administrators whose admin rights are limited only to the assigned SVM.
- Use quality-of-service capabilities to manage resource utilization between storage workloads.
- Nondisruptively migrate live data volumes and client connections from one cluster node to another.
- Nondisruptively scale the cluster out by adding nodes. Nodes can likewise be nondisruptively removed from the cluster—you can nondisruptively scale a cluster up and down during hardware refresh cycles.
- Leverage multiple nodes in the cluster to simultaneously service a given SVM's storage workloads. This means that businesses can scale out their SVMs beyond the bounds of a single physical node in response to growing storage and performance requirements, all nondisruptively.
- Apply software and firmware updates and configuration changes without cluster, SVM, and volume downtime.

As IT operations become increasingly critical to business, any downtime can significantly impact business operations. Downtime causes lost business, poor customer satisfaction, and competitive weakness. Storage infrastructures must be functional in today's 24/7 environments. Nondisruptive operations (NDO) in clustered Data ONTAP are fundamental and facilitate the storage infrastructure in remaining operational during hardware and software maintenance and refresh operations.

In particular, the database industry and SQL Server are being hit hard by demands for more data and ever-increasing requirements for improved availability. The trend toward improved business analytics and reporting driven by big data, data warehousing, and business intelligence has also contributed to increased database storage demands. The typical data center lacks the agility that is necessary to meet the growing demands of the business units it supports. With today's more cloud-centric environments and agile-focused development processes, the traditional and slow data centers cannot respond faster.

SMB Protocol

The Server Message Block (SMB) protocol is a network file-sharing protocol. The SMB protocol (also known as the Common Internet File System, or CIFS) has been supported on the Windows Server® operating system since the Windows NT® release. It mainly supports client-server file access between the file shares on the network and client systems.

The ability to store user databases on SMB network shares was first officially supported in SQL Server 2008 and Windows Server 2008. SQL Server 2012 carries forward this capability and adds several new high-availability options when used with Windows Server 2012. Earlier versions did not support using SMB file shares to store SQL Server databases. There were several reasons for this, but the primary reservations were that the raw disk performance was not fast enough to support demanding database I/O requirements and the network connections to file shares were not as reliable as direct-attached or storage-area-network (SAN) storage.

However, since that time, both the network technology and the SMB protocol have seen significant improvements. Today, 10GB Ethernet is not uncommon, and 40GB Ethernet will be available in the near future. As for reliable network connections through the SMB protocol, Windows released the SMB 3.0 protocol Windows Server 2012. The latest version of the SMB protocol provides many significant improvements, including:

- **SMB transparent failover**. SMB transparent failover enhances the availability of Windows File Services. If a hardware or software failure occurs or the connectivity to a clustered node is lost, all SMB 3.0 clients can transparently reconnect to another cluster node with no downtime. Both the client and server must run the SMB 3.0 protocol to enable SMB transparent failover.
- **SMB scale out**. SMB scale out uses cluster shared volumes (CSV) to provide simultaneous file share access through all nodes in a clustered file server. SMB scale out provides better utilization of network bandwidth and load balancing of the file server clients. SMB scale out can optimize network utilization for server applications such as SQL Server.
- **SMB persistent handles.** SMB protocol was unable to handle issues for a long time in scenarios in which a client became disconnected (through any means) from its open files. With the release of SMB2, this protocol made a positive difference in regard to resiliency with the introduction of durable handles. Durable handles allow a client to survive a brief interruption of the network connection. However, they have drawbacks, such as the need for a specific oplock or lease level, which means a downgrade of the oplock or lease level impacts the durable nature of the open file handle. Durable handles are also passive in their reconnect model. When a handle goes into a "disconnected" state, without any I/O to that open file handle, the state stays disconnected. Finally, the biggest drawback is that durable handles were not meant to survive failover events, whether on a client or a server.

  Now persistent handles have been introduced. They are not tied to an oplock or lease level, and they can survive failover events on the client or server end. Last but not least they are not passive in their reconnect model.

- **SMB cluster client failover**. Cluster client failover allows a client-side cluster failure event (that is, two Microsoft Windows Servers in a cluster and one node fails) to occur and provides the capability for the surviving node to take over and continue running the applications of the failed node.
- **SMB witness**. Essentially, this is a mechanism that allows the cluster-based failover of CIFS shares. Rather than waiting for a network timeout to occur, a client can receive notifications from a witness server informing it that a failure has occurred and that the connection it thought it had is no longer available. The witness server will notify the client about the alternate connection.

  Note: Do not confuse this witness server with the Microsoft Cluster "file share witness"; they are not the same.

For more details on SMB3 support in clustered Data ONTAP, refer to TR-4191: Best Practices Guide for Clustered Data ONTAP 8.2 Windows File Services.

The unified architecture of the NetApp platform is fully capable of being used to store databases on network shared storage that is accessed through SMB.

With clustered Data ONTAP 8.2.1, SQL Server database files can now reside on continuously available (CA) shares. Therefore, upgrading clustered Data ONTAP software, adding storage controllers or shelves to a cluster, adding HBAs, adding Flash Cache™ intelligent caching, and upgrading components can be done nondisruptively. So can redistributing data across controllers to improve performance, moving data across controllers to rebalance capacity, and redistributing data across storage tiers within a cluster to optimize disk performance.

2 Intended Audience

The focus of this technical report is to provide an overview of the way to configure your environment in order to run SQL Server over the SMB3 protocol and to do so in a manner that enables the SMB protocol to provide nondisruptive operations.

This document is intended for system administrators, storage architects, and SQL Server database administrators who are responsible for deploying such a solution in a customer environment. We assume that the reader is familiar with the various components of the solution.

2.1 Terminology

Table 1 lists the terminology definitions and examples that are used in this document.

Table 1) Terminology.

Terminology	Definition	Example in This Document
Server	Windows® host server where SQL Server is installed.	SQL3.demo.netapp.com
Controller	NetApp FAS storage controller.	Cluster1
SQL Server Service account	The domain account that starts SQL Server and/or SQL Server Agent.	demo\SQLAdmin
Installer account	The domain account that installs SQL Server.	demo\DBA
SMB share or CIFS share	The folder where system or user databases reside.	\\redmondshare\hr_systemdb\

3 Storage Events and SMB Operations

When using the SMB protocol for applications, you need to understand the capabilities of each protocol version to survive events in the environment. Table 2 lists the storage events and the capability of the different SMB protocol versions to provide nondisruptive operations. The information provided in Table 2 does not take into account any resiliency within the SQL Server software itself.

Table 2) Storage events and capability of different SMB protocol versions to provide nondisruptive operations.

SMB Protocol Version	Volume Moves	LIF Migrate	Aggregate Relocate
SMB1	✓	×	×
SMB2.x	✓	✓	×
SMB3 (continuously available share disabled)	✓	✓	×
SMB3 (continuously available share enabled)	✓	✓	✓

SMB3 with the **continuously available share property enabled** provides a very high level of resiliency to issues that occur on the storage end of the architecture. The limitations of those protocol versions other than SMB3 (CA) are not due to restrictions within clustered Data ONTAP, but rather on constraints in the protocol versions themselves.

Following is a brief description of the storage events.

- **Volume move**. Moving a volume from one node in the cluster to another in the same cluster.
- **LIF migrate**. Moving a virtual network interface from one node in the cluster to another in the same cluster.
- **Aggregate relocate**. Reassigning an aggregate from one node to its HA partner in the cluster.
- **Storage failover**. Taking over a node in an HA pair whether planned or unplanned. Storage failover is a combination of LIF migrate and aggregrate relocate. Only SMB **continuously available shares** can provide NDO during storage failover.

For more details on the preceding features, refer to the clustered Data ONTAP product documentation. Explaining these features extensively is beyond the scope of this document.

4 System Configuration

To set up both the SQL Server system and user databases on NetApp storage, the following software products were used:

- Windows 2012 Datacenter Edition
- SQL Server 2012 sp1 Enterprise Edition x64
- Data ONTAP PowerShell™ Toolkit 1.6 (optional)
- Clustered Data ONTAP release 8.2.1 or higher
- SMB3
- Continuously available share
- OnCommand® System Manager 3.0
- SQL Server domain group that has installer account (demo\DBA) and SQL Server service account (demo\SQLAdmin) as the members of the group

To make sure of the configuration compatibility across the NetApp stack, refer to the NetApp Interoperability Matrix Tool (IMT).

5 Creating SMB Shares for Database Files

5.1 Create SVM

To create SMB shares storage for SQL Server database files, SVM must be created. The following section provides the step-by-step instructions on how to create SVM using OnCommand System Manager Wizards.

1. Start OnCommand System Manager.
2. Double-click the cluster to create a Storage Virtual Machine.

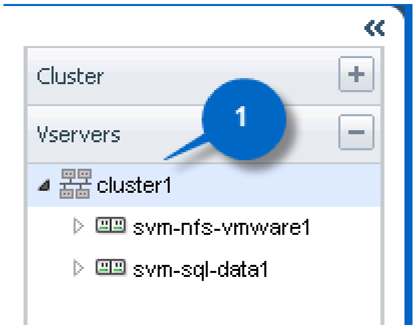

3. Click Vservers.

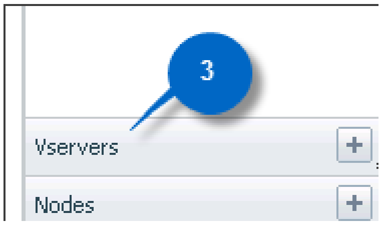

4. Click Create.

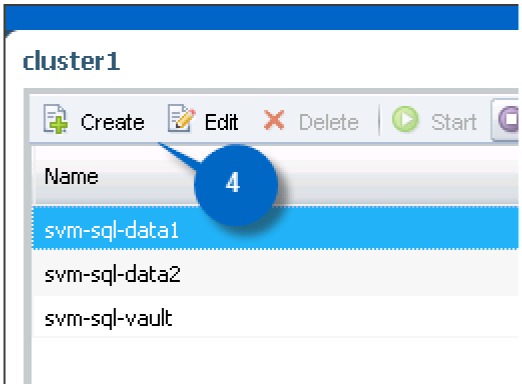

5. The Vserver Setup window appears.
6. Enter your Vserver name, select CIFS for Data Protocols, and select NTFS for Security Style.
7. Click Submit & Continue.

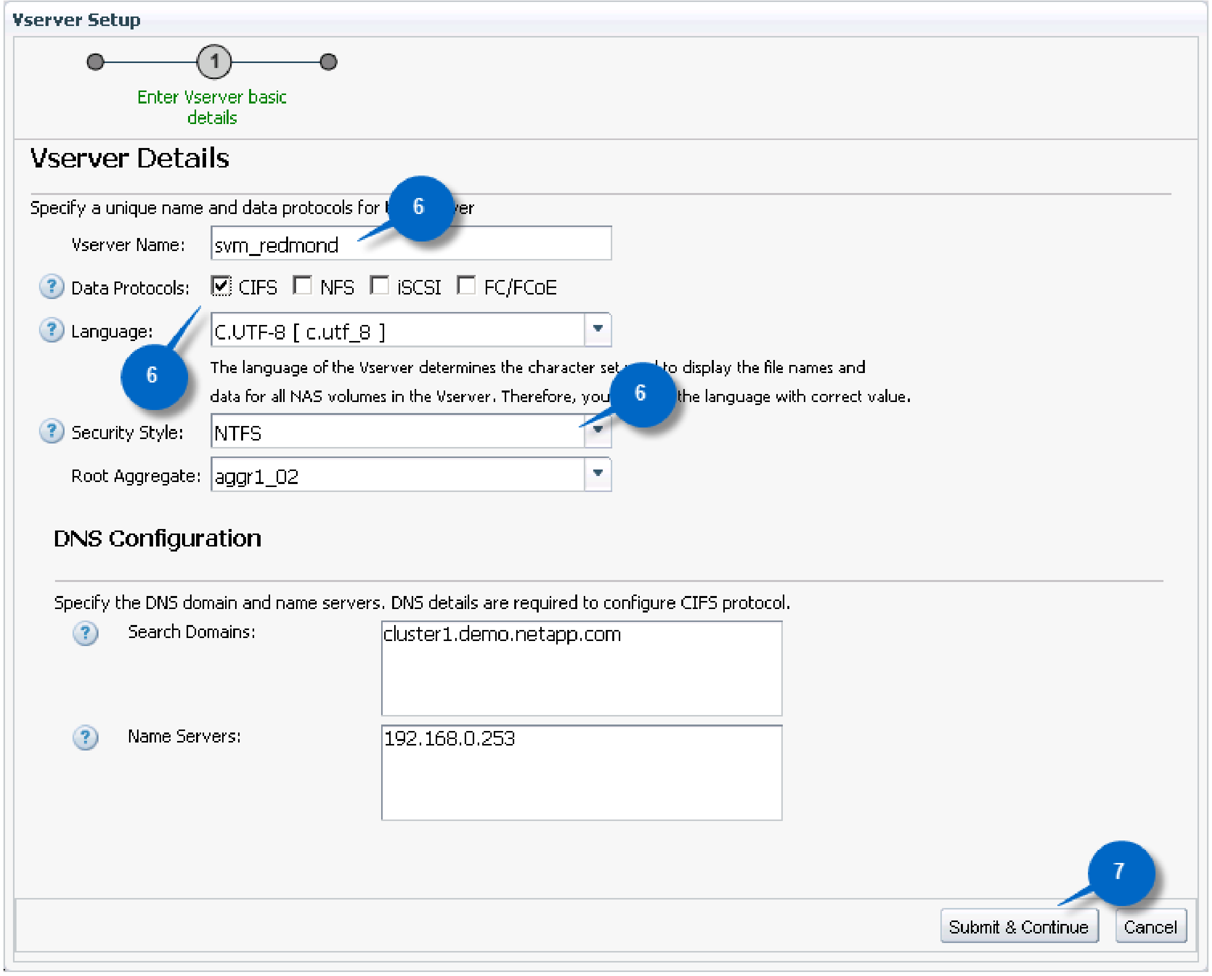

Vserver Setup

Vserver Details

Specify a unique name and data protocols for t ver

Vserver Name: svm_redmond

Data Protocols: ☑ CIFS □ NFS □ iSCSI □ FC/FCoE

Language: C.UTF-8 [c.utf_8]

The language of the Vserver determines the character set to display the file names and data for all NAS volumes in the Vserver. Therefore, you the language with correct value.

Security Style: NTFS

Root Aggregate: aggr1_02

DNS Configuration

Specify the DNS domain and name servers. DNS details are required to configure CIFS protocol.

Search Domains: cluster1.demo.netapp.com

Name Servers: 192.168.0.253

The next set of steps provides information about how to configure the CIFS protocol. This procedure creates a Data LIF to serve CIFS traffic. It also sets up CIFS and creates an Active Directory® computer account. The name you specify is what the SQL Servers use to connect to the CIFS shares.

8. Enter the IP Address, Netmask, and Gateway information for the Data LIF.
9. Enter the CIFS Server Name—the name you need to refer to when setting up SQL Server data and log files to reside on an SMB share (that is, \\redmondshare.demo.netapp.com\).
10. Enter the Active Directory Domain name—the domain you need to make the CIFS server a member of.
11. Enter the Administrator Name and Password that can join the CIFS server to the domain. The account specified here will need permissions to create objects in Active Directory.
12. Click Submit & Continue.

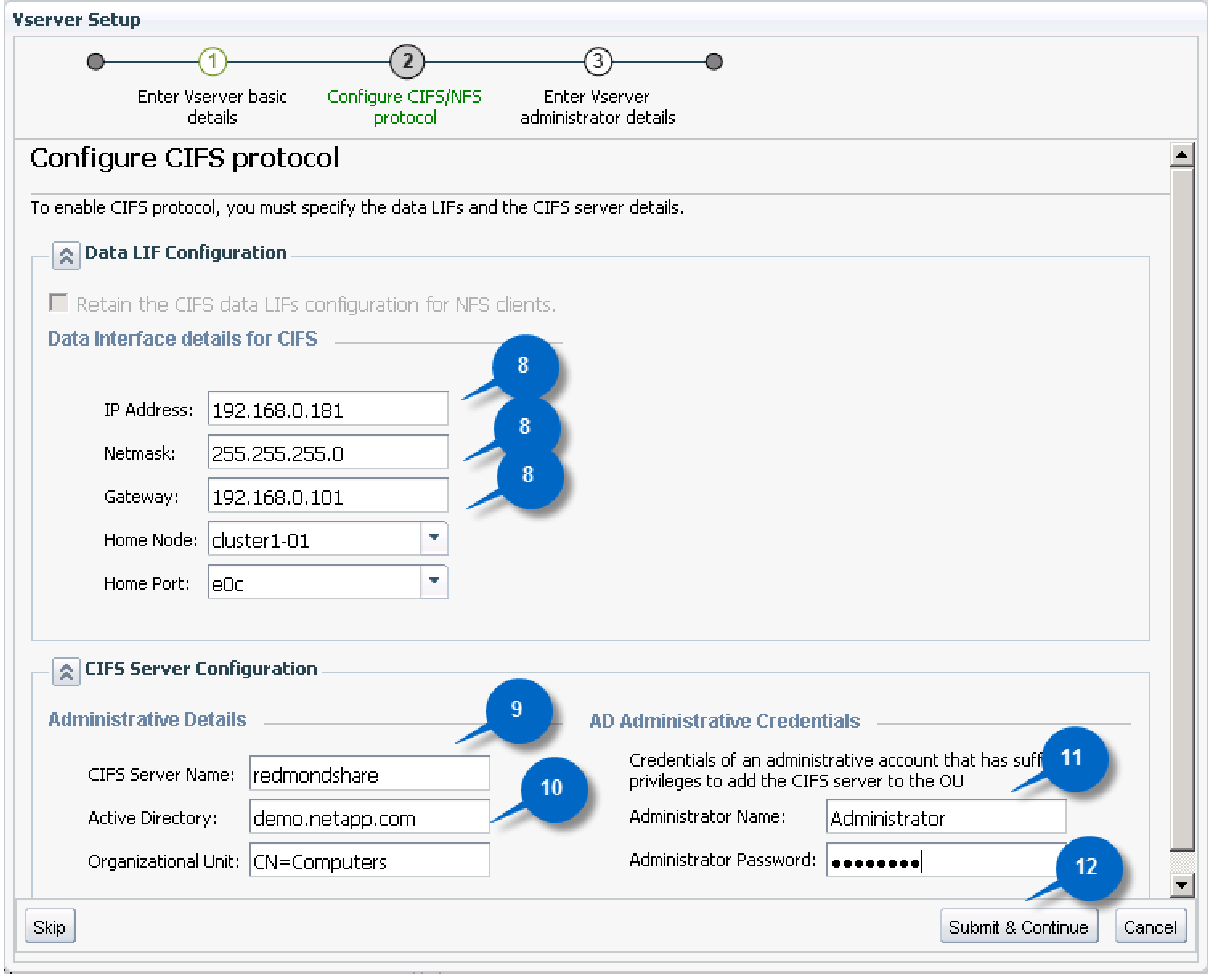
Vserver Setup
1 Enter Vserver basic details
2 Configure CIFS/NFS protocol
3 Enter Vserver administrator details

Configure CIFS protocol

To enable CIFS protocol, you must specify the data LIFs and the CIFS server details.

Data LIF Configuration

□ Retain the CIFS data LIFs configuration for NFS clients.

Data Interface details for CIFS

IP Address:	192.168.0.181
Netmask:	255.255.255.0
Gateway:	192.168.0.101
Home Node:	cluster1-01
Home Port:	e0c

CIFS Server Configuration

Administrative Details

CIFS Server Name:	redmondshare
Active Directory:	demo.netapp.com
Organizational Unit:	CN=Computers

AD Administrative Credentials

Credentials of an administrative account that has suff privileges to add the CIFS server to the OU

Administrator Name:	Administrator
Administrator Password:	●●●●●●●●

13. Enter the vsadmin password and confirm the password.
14. You must create a new management LIF for the Vserver to allow SnapDrive® for Windows and SnapManager® for SQL to connect to the CIFS server. Therefore, select Create a new LIF for Vserver management.
15. Enter the IP Address, Netmask, and Gateway.
16. Click Submit & Continue.

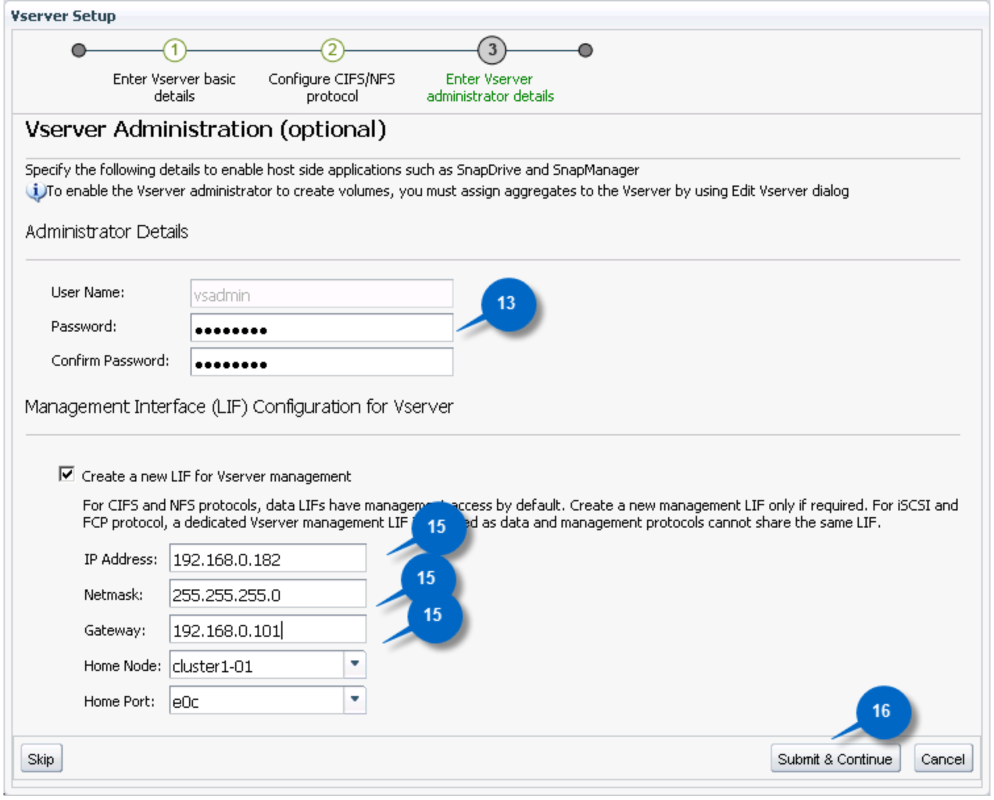

Vserver Setup
1
Enter Vserver basic details
2
Configure CIFS/NFS protocol
3
Enter Vserver administrator details
13
15
15
15
16
Skip
Submit & Continue
Cancel

Vserver Administration (optional)

Specify the following details to enable host side applications such as SnapDrive and SnapManager

To enable the Vserver administrator to create volumes, you must assign aggregates to the Vserver by using Edit Vserver dialog

Administrator Details

User Name: vsadmin

Password: ••••••••

Confirm Password: ••••••••

Management Interface (LIF) Configuration for Vserver

☑ Create a new LIF for Vserver management

For CIFS and NFS protocols, data LIFs have managem access by default. Create a new management LIF only if required. For iSCSI and FCP protocol, a dedicated Vserver management LIF i ed as data and management protocols cannot share the same LIF.

IP Address: 192.168.0.182

Netmask: 255.255.255.0

Gateway: 192.168.0.101

Home Node: cluster1-01

Home Port: e0c

17. In the New Vserver Summary page, click OK.

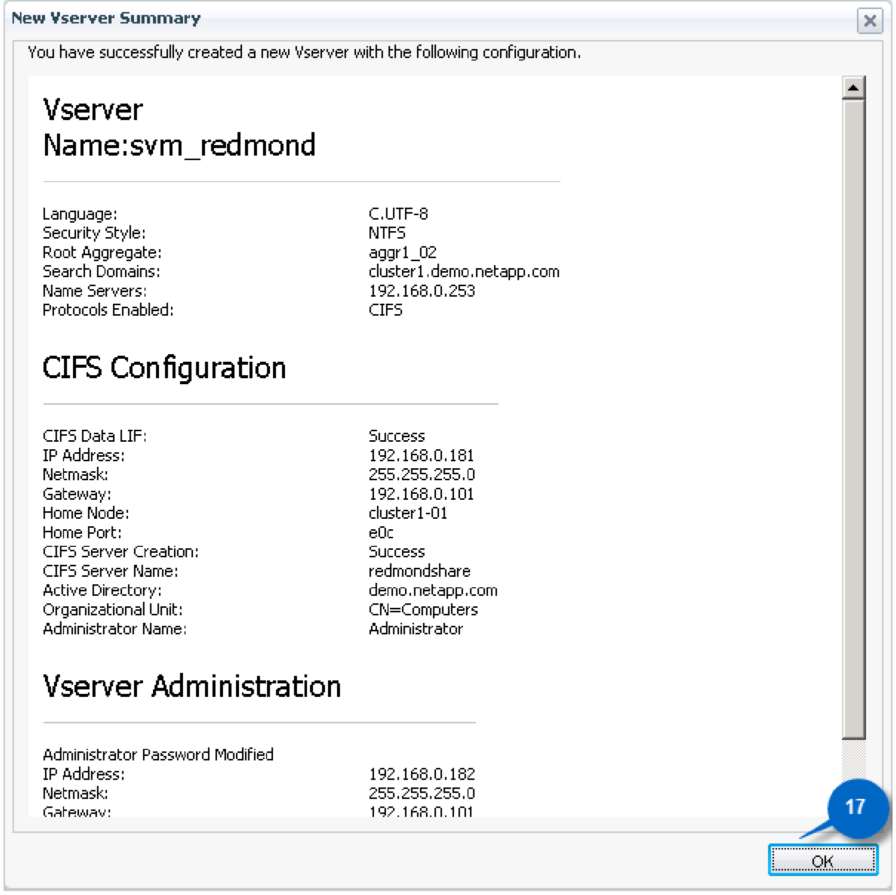

New Vserver Summary

You have successfully created a new Vserver with the following configuration.

Vserver Name:svm_redmond

Language:	C.UTF-8
Security Style:	NTFS
Root Aggregate:	aggr1_02
Search Domains:	cluster1.demo.netapp.com
Name Servers:	192.168.0.253
Protocols Enabled:	CIFS

CIFS Configuration

CIFS Data LIF:	Success
IP Address:	192.168.0.181
Netmask:	255.255.255.0
Gateway:	192.168.0.101
Home Node:	cluster1-01
Home Port:	e0c
CIFS Server Creation:	Success
CIFS Server Name:	redmondshare
Active Directory:	demo.netapp.com
Organizational Unit:	CN=Computers
Administrator Name:	Administrator

Vserver Administration

Administrator Password Modified	
IP Address:	192.168.0.182
Netmask:	255.255.255.0
Gateway:	192.168.0.101

Note: After the SVM has been created, the CIFS server (in this example, redmondshare) should be registered in the Domain Name Server (DNS).

5.2 Creating Volumes

For the Windows host to store data, the volumes and shares must be created. You can create volumes by using either OnCommand System Manager or the Data ONTAP PowerShell Toolkit. Following is the procedure to create volumes using OnCommand System Manager.

1. Select the Virtual Storage Machine that you created in OnCommand System Manager.
2. Select Storage > Volumes.
3. Select Create.
4. The Create Volume window appears.
5. Enter your volume name and select Aggregate.
6. Select NAS and specify the size of the volume.
7. Click Create.

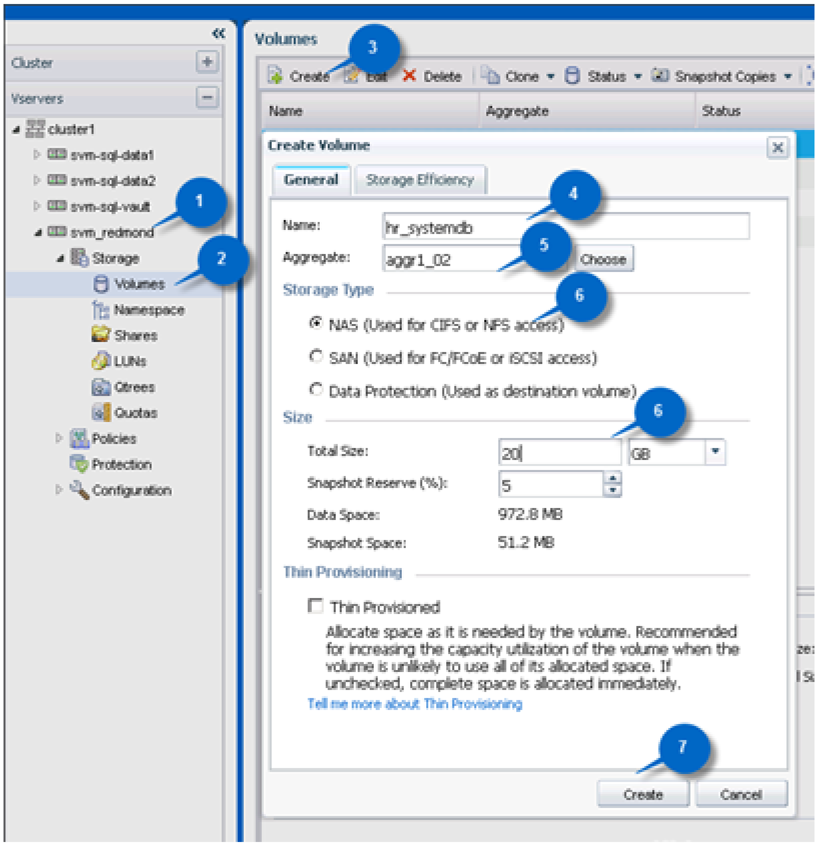

Create Volume

General | Storage Efficiency

Name: hr_systemdb

Aggregate: aggr1_02 Choose

Storage Type

● NAS (Used for CIFS or NFS access)

○ SAN (Used for FC/FCoE or iSCSI access)

○ Data Protection (Used as destination volume)

Size

Total Size: 20 GB

Snapshot Reserve (%): 5

Data Space: 972.8 MB

Snapshot Space: 51.2 MB

Thin Provisioning

□ Thin Provisioned

Allocate space as it is needed by the volume. Recommended for increasing the capacity utilization of the volume when the volume is unlikely to use all of its allocated space. If unchecked, complete space is allocated immediately.

Tell me more about Thin Provisioning

Create | Cancel

After the volumes have been created, you need to share the volumes so that SQL Server can access them. You can accomplish this task by using either OnCommand System Manager or the Data ONTAP PowerShell Toolkit. The following section discusses granting permission to the shares for a domain group called demo\SQLAdmins. This domain group has demo\SQLAdmin and demo\Administrator accounts as the members. These step-by-step instructions describe how to share the volume using OnCommand System Manager.

1. Select your Server Virtual Machine.
2. Select Storage > Volume > Shares.
3. Select Create Share.
4. The Create Share window appears.
5. Select Browse, select the volume you want to share, and click OK.
6. Enter the Share Name that will be used later in the SQL configuration (that is, \\redmondshare\hr_systemdb).
7. Select Enable continuous availability for Hyper-V.
8. Click Create.

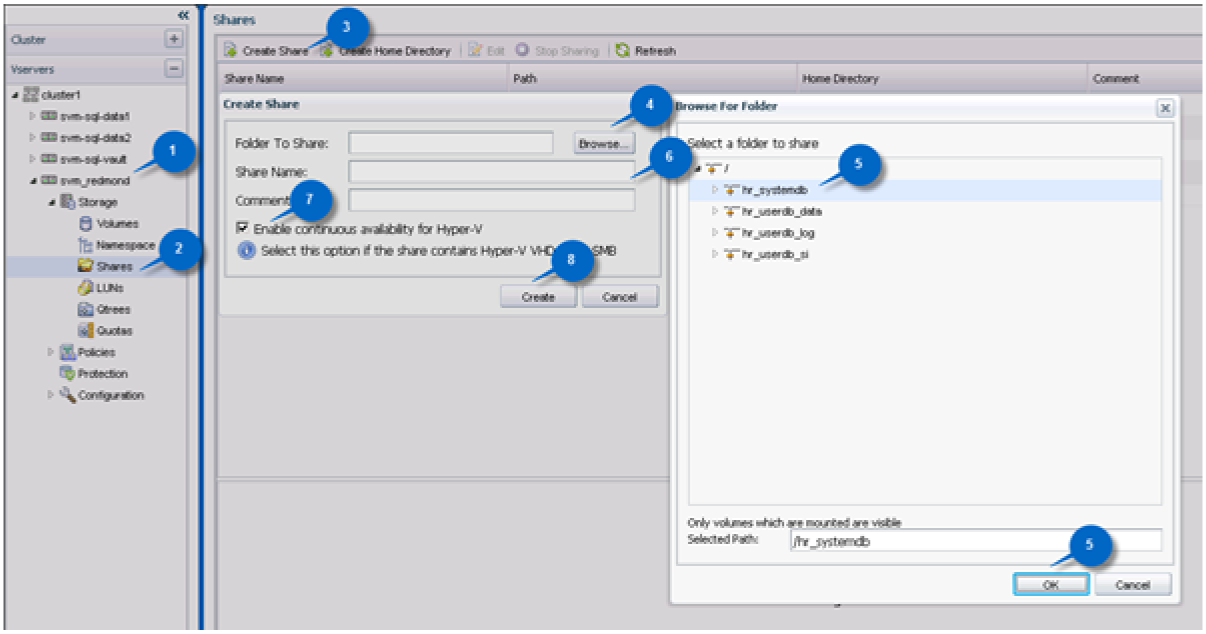

☑ Enable continuous availability for Hyper-V

Selected Path: /hr_systemdb

9. Select the volume that has been shared and select Edit. The Edit Settings window appears.
10. Select the Permissions tab and click Add. The Add Users and Groups dialog box is displayed.
11. Type the domain group for which you want access to the share, and then click OK.

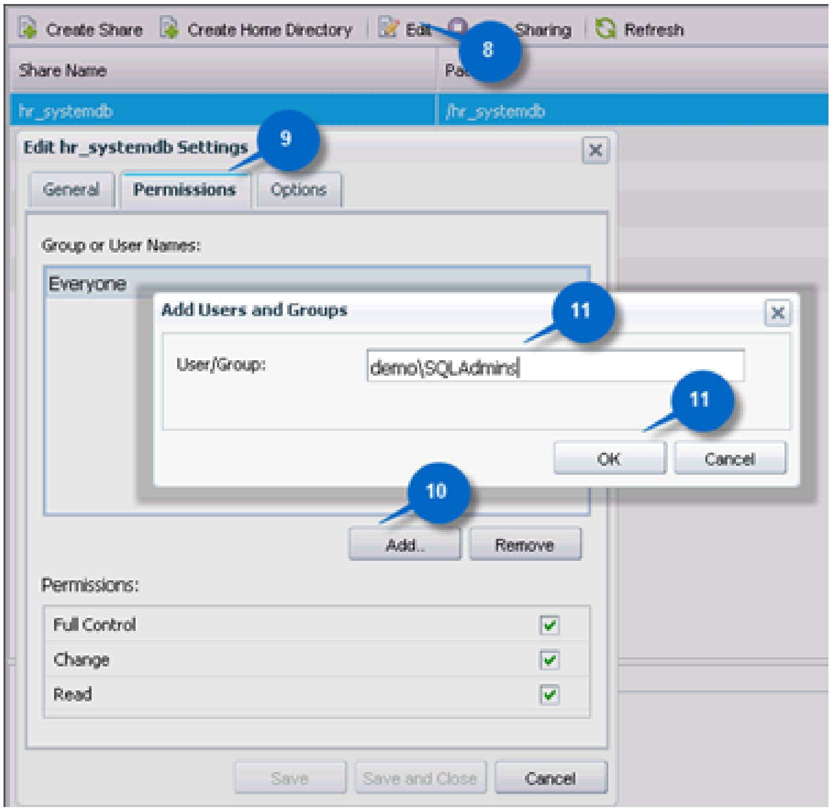

Share Name	Path
hr_systemdb	/hr_systemdb

User/Group: demo\SQLAdmins

Full Control	☑
Change	☑
Read	☑

12. Select Full Control.
13. Select Everyone, and then click Remove.

14. Click Options.

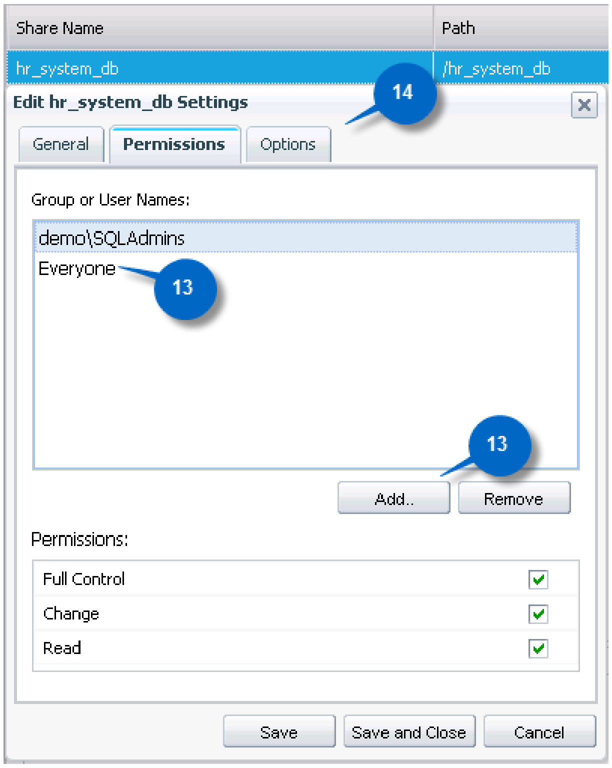

Share Name	Path
hr_system_db	/hr_system_db

Edit hr_system_db Settings

Group or User Names:

demo\SQLAdmins
Everyone

Permissions:

Full Control	☑
Change	☑
Read	☑

15. Select Browsable, Enable Oplocks, and Notify Change.
16. Click Save and Close.

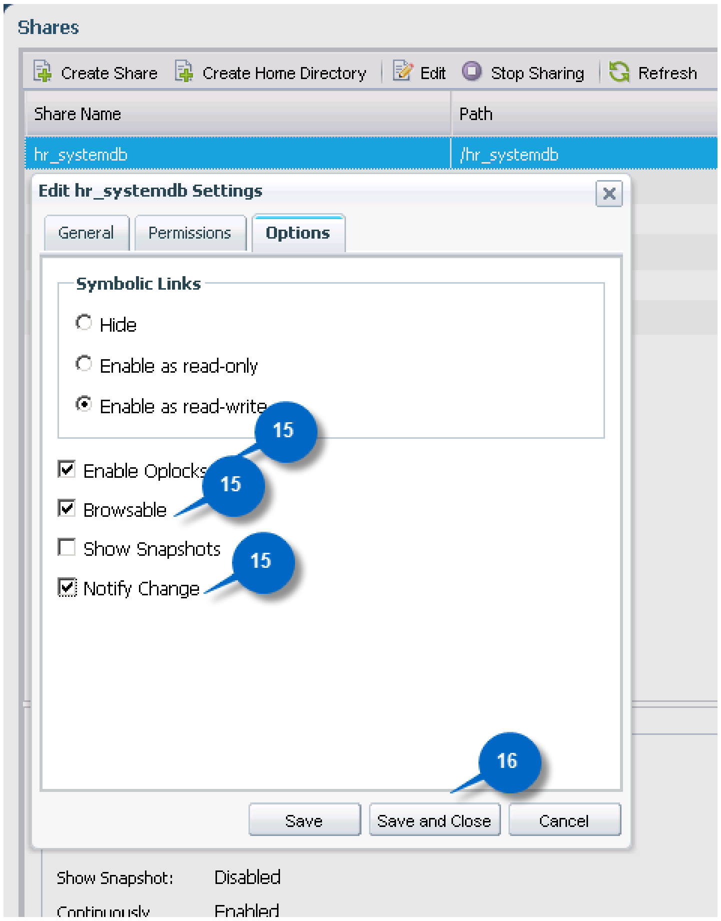

Share Name	Path
hr_systemdb	/hr_systemdb

Symbolic Links

○ Hide

○ Enable as read-only

● Enable as read-write

☑ Enable Oplocks

☑ Browsable

□ Show Snapshots

☑ Notify Change

Show Snapshot: Disabled

Continuously Enabled

Following is the PowerShell script to create volumes and shares and grant appropriate permission to a domain group.

```
powershell.exe -ExecutionPolicy Unrestricted -NoLogo -NonInteractive
if ((Get-Module| select -exp name) -notcontains 'DataOntap'){Import-module DataOntap}

$cserver = "192.168.0.101"
$vserver = "svm_redmond"
$aggr = "aggr1_01"
$dc = "192.168.0.253"
$netmask = "255.255.255.0"
$domain = "demo.netapp.com"
$UserGroup = "Demo\SQLAdmins"
$AdminPwd = "Netapp1!"
$size = "20g"
$vols = @("hr_userdb_data", "hr_userdb_log", "hr_userdb_si")
$node = "cluster1-01"
$datalif = @{lifname="cifs_data_lif";Policy="data";ip="192.168.0.181";port="e0c"}
$mgmtlif = @{lifname="cifs_mgmt_lif";Policy="mgmt";ip="192.168.0.182";port="e0g"}

$password = ConvertTo-SecureString "Netapp1!" -AsPlainText -Force
$cred = New-Object -TypeName System.Management.Automation.PSCredential -ArgumentList
"admin",$password

#Connect to Cluster
Connect-NcController $cserver -Credential $cred
#Add Volume
Foreach ($vol in $vols) {
        New-NcVol -Name $vol -Aggregate $aggr -Size $size -VserverContext $vserver -JunctionPath
/$vol;
        #Create Shares
        Add-NcCifsShare -Name $vol -Path /$vol –VserverContext $vserver `
        -ShareProperties "oplocks","browsable","changenotify","continuously-available" `
        | Add-NcCifsShareAcl -UserorGroup $UserGroup –Permission "full_control";
        #Remove Everyone Account
        Remove-NcCifsShareAcl –VserverContext $vserver $vol Everyone
}
```

6 Installing SQL Server over SMB

SQL Server 2012 allows you to install with both system and user databases on SMB file shares. This allows you to build end-to-end NAS solutions with the data management capabilities of NAS, such as volume autogrow and the ability to shrink NAS volumes, which the SQL Server can take advantage of. In order to install SQL Server over SMB shares, the installer must meet these requirements.

- The installer has been granted security privileges to the SVM. This process can be done using the following PowerShell script.

```
powershell.exe -ExecutionPolicy Unrestricted -NoLogo -NonInteractive
if ((Get-Module| select -exp name) -notcontains 'DataOntap'){Import-module DataOntap}

$cserver = "192.168.0.101"
$vserver = "svm_redmond"
$installer = “demo\DBA”
$password = ConvertTo-SecureString "Netapp1!" -AsPlainText -Force
$cred = New-Object -TypeName System.Management.Automation.PSCredential -ArgumentList
"admin",$password

#Connect to Cluster
Connect-NcController $cserver -Credential $cred
Add-NcCifsPrivilege -Name $installer -vservercontext $vserver -Privilege sesecurityprivilege
```

- The installer has read/write access to the share.
- The installer is a member of local administrator to the window host.

The installation process is the same as for the typical block storage. The only exception is that the data root directory for system database files can be pointed to a CIFS share that has been created during the database engine configuration step.

Figure 1) Specified SMB shares where system database files will reside.

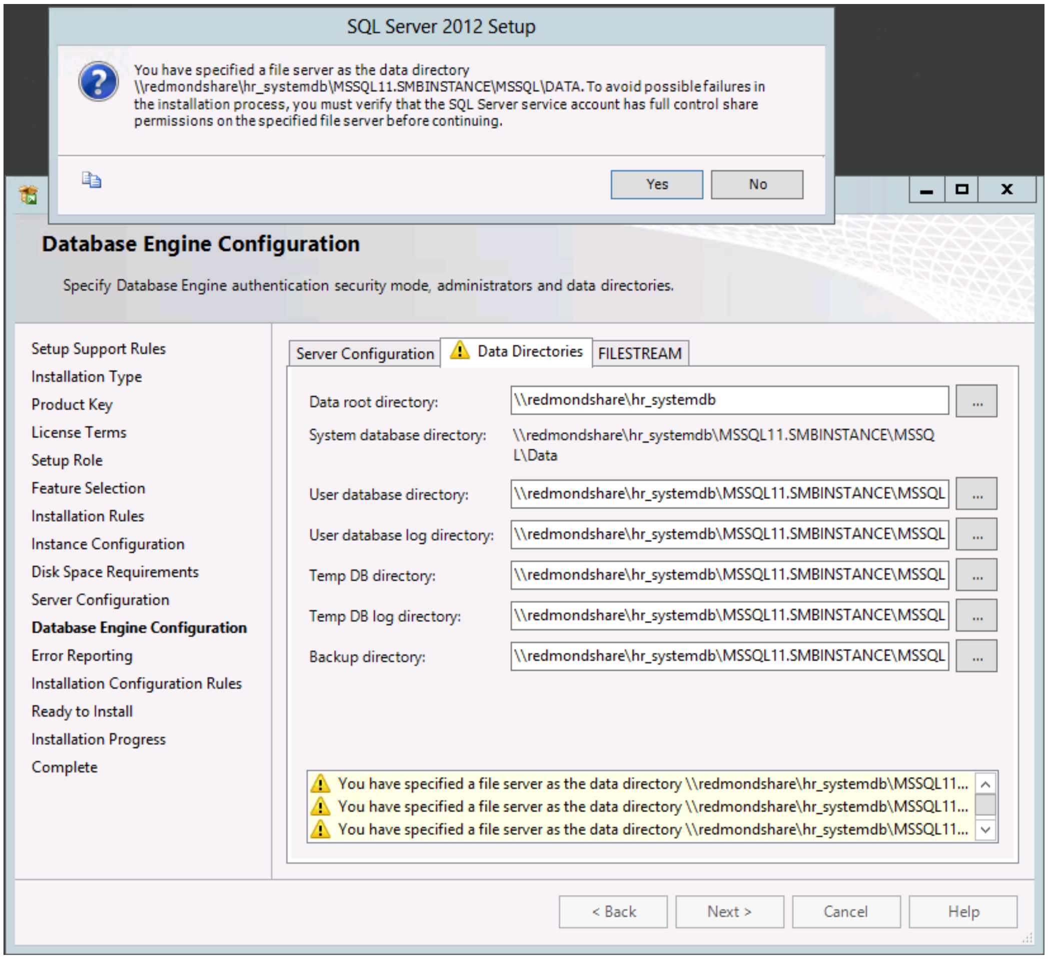

SQL Server 2012 Setup

You have specified a file server as the data directory \\redmondshare\hr_systemdb\MSSQL11.SMBINSTANCE\MSSQL\DATA. To avoid possible failures in the installation process, you must verify that the SQL Server service account has full control share permissions on the specified file server before continuing.

Database Engine Configuration

Specify Database Engine authentication security mode, administrators and data directories.

Data root directory: \\redmondshare\hr_systemdb

System database directory: \\redmondshare\hr_systemdb\MSSQL11.SMBINSTANCE\MSSQL\Data

User database directory: \\redmondshare\hr_systemdb\MSSQL11.SMBINSTANCE\MSSQL

User database log directory: \\redmondshare\hr_systemdb\MSSQL11.SMBINSTANCE\MSSQL

Temp DB directory: \\redmondshare\hr_systemdb\MSSQL11.SMBINSTANCE\MSSQL

Temp DB log directory: \\redmondshare\hr_systemdb\MSSQL11.SMBINSTANCE\MSSQL

Backup directory: \\redmondshare\hr_systemdb\MSSQL11.SMBINSTANCE\MSSQL

⚠ You have specified a file server as the data directory \\redmondshare\hr_systemdb\MSSQL11...
⚠ You have specified a file server as the data directory \\redmondshare\hr_systemdb\MSSQL11...
⚠ You have specified a file server as the data directory \\redmondshare\hr_systemdb\MSSQL11...

7 Creating and Migrating User Databases to SMB Shares

7.1 Creating Databases over SMB Shares

With SQL Server 2012, you can create user databases that have database and log files residing on SMB shares in the same manner as you would do with block storage. The following `T-SQL` command allows you to create a database that has files residing on an SMB share.

```
CREATE DATABASE [Adventureworks]
 CONTAINMENT = NONE
 ON PRIMARY
( NAME = N'Adventureworks', FILENAME = N'\\redmondshare\hr_userdb_data\Adventureworks.mdf' , SIZE
= 4096KB , FILEGROWTH = 1024KB )
 LOG ON
( NAME = N'Adventureworks_log', FILENAME = N'\\redmondshare\hr_userdb_log\Adventureworks_log.ldf'
, SIZE = 1024KB , FILEGROWTH = 10%)
GO
USE [Adventureworks]
GO
IF NOT EXISTS (SELECT name FROM sys.filegroups WHERE is_default=1 AND name = N'PRIMARY')
ALTER DATABASE [Adventureworks] MODIFY FILEGROUP [PRIMARY] DEFAULT
GO
```

7.2 Migrating Databases to SMB Shares

You can also migrate databases from a local drive to SMB shares using SnapManager for SQL Server (SMSQL). In order to migrate a database using SMSQL, SnapDrive for Windows (SDW) 7 must be installed and transport protocol must be configured.

Following are the steps to configure transport protocol.

1. Start SnapDrive.
2. Select your server.
3. Right-click the selected server and select Transport Protocol Settings. The Transport Protocol Settings are displayed.

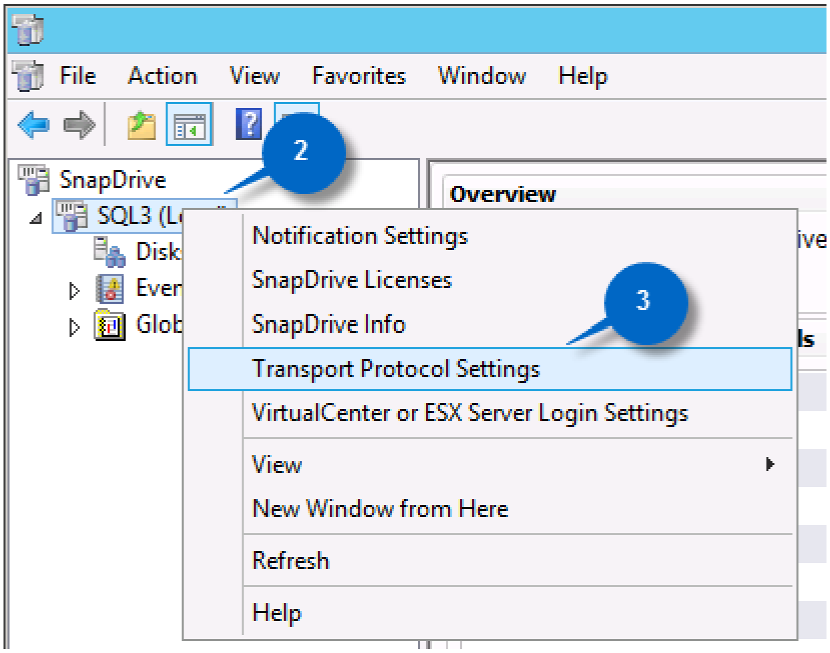

4. Click Add.
5. Type the IP address in the Storage System field and type in the User Name and Password.
6. Click OK.

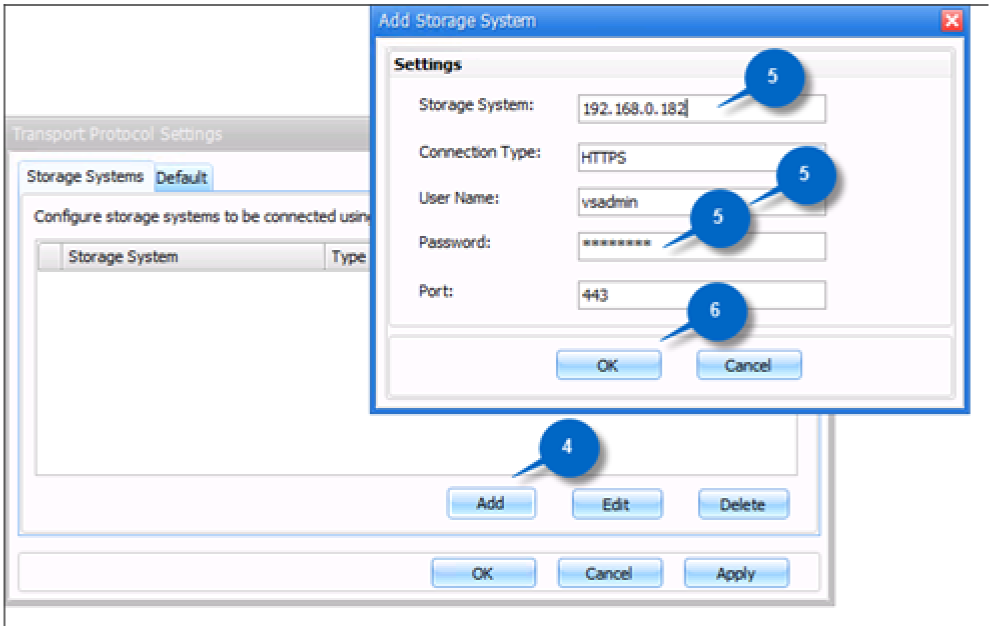

Configure storage systems to be connected usin

Add Storage System

Settings

Storage System: 192.168.0.182

Connection Type: HTTPS

User Name: vsadmin

Password: ********

Port: 443

After SDW has been installed and configured, you can migrate the database using SMSQL 7 in the same manner as you would do with SAN storage.

Following is the screenshot of the database selection step for the configuration wizard:

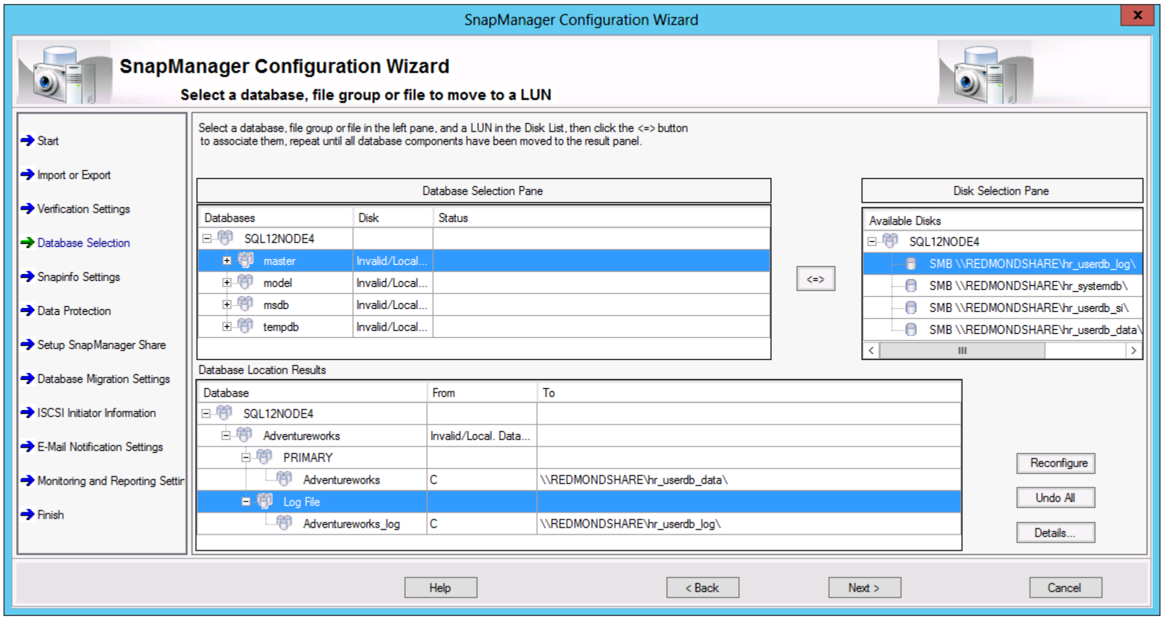

SnapManager Configuration Wizard

Select a database, file group or file to move to a LUN

Select a database, file group or file in the left pane, and a LUN in the Disk List, then click the <=> button to associate them, repeat until all database components have been moved to the result panel.

Database Selection Pane

Databases	Disk	Status
SQL12NODE4		
master	Invalid/Local...	
model	Invalid/Local...	
msdb	Invalid/Local...	
tempdb	Invalid/Local...	

Disk Selection Pane

Available Disks
SQL12NODE4
SMB \\REDMONDSHARE\hr_userdb_log\
SMB \\REDMONDSHARE\hr_systemdb\
SMB \\REDMONDSHARE\hr_userdb_si\
SMB \\REDMONDSHARE\hr_userdb_data\

Database Location Results

Database	From	To
SQL12NODE4		
Adventureworks	Invalid/Local. Data...	
PRIMARY		
Adventureworks	C	\\REDMONDSHARE\hr_userdb_data\
Log File		
Adventureworks_log	C	\\REDMONDSHARE\hr_userdb_log\

Following is the screenshot explaining how to set up a snapinfo directory for all the databases.

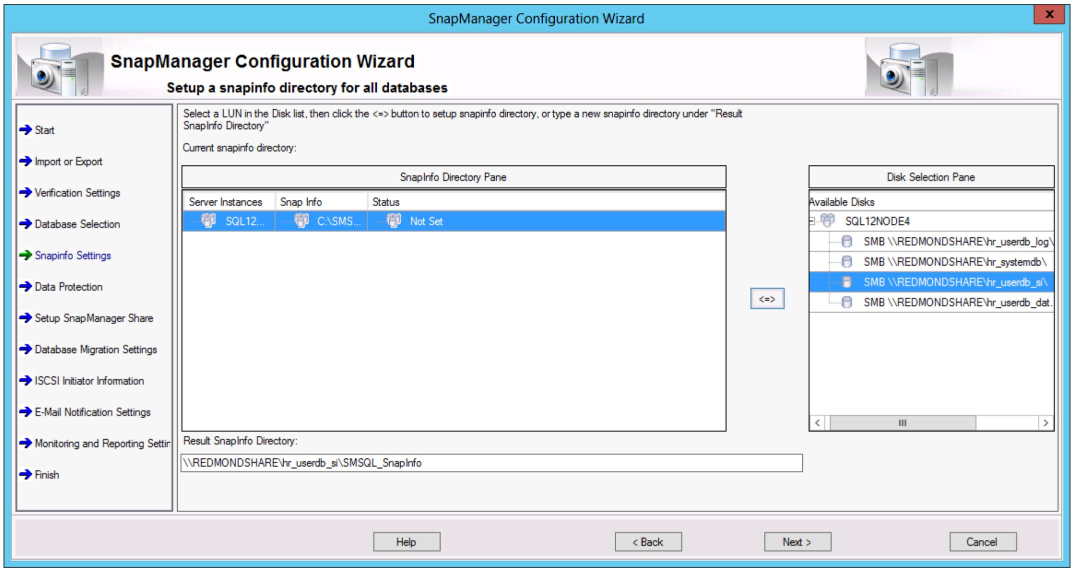
SnapManager Configuration Wizard
SnapManager Configuration Wizard
Setup a snapinfo directory for all databases
Start
Import or Export
Verification Settings
Database Selection
Snapinfo Settings
Data Protection
Setup SnapManager Share
Database Migration Settings
ISCSI Initiator Information
E-Mail Notification Settings
Monitoring and Reporting Settir
Finish
SnapInfo Directory Pane
Disk Selection Pane
<=>
Help
< Back
Next >
Cancel

Select a LUN in the Disk list, then click the <=> button to setup snapinfo directory, or type a new snapinfo directory under "Result SnapInfo Directory"

Current snapinfo directory:

Server Instances	Snap Info	Status
SQL12...	C:\SMS...	Not Set

Available Disks
- SQL12NODE4
  - SMB \\REDMONDSHARE\hr_userdb_log\
  - SMB \\REDMONDSHARE\hr_systemdb\
  - SMB \\REDMONDSHARE\hr_userdb_si\
  - SMB \\REDMONDSHARE\hr_userdb_dat.

Result SnapInfo Directory:

\\REDMONDSHARE\hr_userdb_si\SMSQL_SnapInfo

8 Backing Up Database over SMB Shares Using SnapManager for SQL Server

After the database has been configured using SMSQL, you can use SMSQL to back up in the same manner as you would do in the traditional SAN storage. You can use the SMSQL GUI or a PowerShell script. For more information on how to create database backups, refer to the Best Practice Guide for Microsoft SQL Server and SnapManager 7.0 for SQL Server with Clustered Data ONTAP.

The following describes how to back up a database with a PowerShell script.

```
powershell.exe -ExecutionPolicy Unrestricted -NoLogo -NonInteractive
#Add PowerShell Snapin for SMSQL
add-pssnapin NetApp.SnapManager.SQL.PS.Admin -ErrorAction silentlycontinue

new-backup -Server SQL3 -Database SQL3, ‘1’, ‘Adventureworks’
```

9 Database Clone Lifecycle Management

The database that has the data files residing on SMB shares can be cloned similar to the way the database that has the data files residing on SAN storage can be cloned. The cloning can be done using the GUI or PowerShell. For more information on cloning databases, refer to the Best Practice Guide for Microsoft SQL Server and SnapManager 7.0 for SQL Server with Clustered Data ONTAP.

The following describes how to clone the database using a PowerShell script from the target server.

```
powershell.exe -ExecutionPolicy Unrestricted -NoLogo -NonInteractive
#Import sqlps so we can connect to SQL Server instance
Import-Module sqlps -DisableNameChecking
#Add PowerShell Snapin for SMSQL
add-pssnapin NetApp.SnapManager.SQL.PS.Admin -ErrorAction silentlycontinue

#Define variable
$TargetServer = get-content env:computername
Set-Location SQLSERVER:\SQL\$TargetServer
$TargetInstance = Get-Item Default
$SourceInstance = "SQL3"
$dbname = "Adventureworks"
$db = $TargetInstance.databases[$dbname]
$MountPoint = "C:\MSSQL"
```

```
#If database exists then remove clone database
if ($db -ne $null){
       delete-clone -Server $TargetServer -Inst $TargetServer -Database $dbname
}

#Create a new clone database
clone-database -Server $SourceInstance -ServerInstance $SourceInstance -Database $dbname `
 -TargetServerInstance $TargetServer -TargetDatabase $dbname -RetainBackups 1
```

10 Summary

With Microsoft SQL Server 2012, Microsoft introduced full support for Server Message Block (SMB) protocol. This means that you can install SQL Server 2012 with the system and user databases on SMB file shares. This allows you to build end-to-end network-attached storage (NAS) solutions in which SQL Server can leverage the data management capabilities of NAS, such as volume autogrow and the ability to shrink NAS volumes. With clustered Data ONTAP 8.2.1, you can create database files that reside over continuous availability shares. NetApp delivers a storage environment with always-on availability. Zero downtime means that business-critical applications such as Microsoft SQL Server do not go offline for routine upgrades and maintenance, add-ons, and even complete hardware refreshes.

Appendix

Following is the complete PowerShell script to create SVMs, volumes, and shares and to grant appropriate permissions to shares.

```
#Before running this script demo\SQLAdmins group must be created in the DC

powershell.exe -ExecutionPolicy Unrestricted -NoLogo -NonInteractive
if ((Get-Module| select -exp name) -notcontains 'DataOntap'){Import-module DataOntap}

$cserver = "192.168.0.101"
$vserver = "svm_redmond"
$cifsserver = "redmondshare"
$aggr = "aggr1_01"
$dc = "192.168.0.253"
$netmask = "255.255.255.0"
$domain = "demo.netapp.com"
$UserGroup = "Demo\SQLAdmins"
$AdminPwd = "Netapp1!"
$size = "20g"
$vols = @("hr_systemdb", "hr_userdb_data", "hr_userdb_log", "hr_userdb_si")
$node = "cluster1-01"
$datalif =
@{lifname="cifs_data_lif";Policy="mgmt";ip="192.168.0.181";port="e0c";dataprotocol="cifs"}
$mgmtlif =
@{lifname="cifs_mgmt_lif";Policy="mgmt";ip="192.168.0.182";port="e0g";dataprotocol="none"}

$password = ConvertTo-SecureString "Netapp1!" -AsPlainText -Force
$cred = New-Object -TypeName System.Management.Automation.PSCredential -ArgumentList
"admin",$password

#Connect to Cluster
Connect-NcController $cserver -Credential $cred

#Create vserver
New-NcVserver $vserver "$($vserver)_root" $aggr -NameServerSwitch file -RootVolumeSecurityStyle
ntfs -Language C;
Set-NcUserPassword -UserName vsadmin -Password $AdminPwd -VserverContext $vserver
Unlock-NcUser vsadmin -Vserver $vserver

#Create lifs
foreach ($lif in $datalif, $mgmtlif){
New-NcNetInterface -Name $lif.lifname -Vserver $vserver -Role data -Node $node `
-Port $lif.port -Address $lif.ip -Netmask $netmask `
-FirewallPolicy $lif.Policy -AdministrativeStatus up `
-DataProtocols $lif.dataprotocol
}

#Copy DNS configuration from cluster to vServer.
Get-NcNetDns cluster1 | New-NcNetDns -VserverContext $vserver;

#Add A record to DNS for the new vServer instance.
$dns = [WmiClass]"\\$dc\root\MicrosoftDNS:MicrosoftDNS_ResourceRecord";
$dns.CreateInstanceFromTextRepresentation($dc, $domain, "$cifsserver.demo.netapp.com IN A
$($datalif.get_item("ip"))");

#Create Cifs Server
Add-NcCifsServer -Name $cifsserver -Domain $domain -AdminUsername Administrator `
-AdminPassword $AdminPwd -VserverContext $vserver
Set-NcCifsOption -VserverContext $vserver -DefaultUnixUser pcuser

#Add Volume
Foreach ($vol in $vols) {
        New-NcVol -Name $vol -Aggregate $aggr -Size $size -VserverContext $vserver -JunctionPath
/$vol;
        #Create Shares
        Add-NcCifsShare -Name $vol -Path /$vol -VserverContext $vserver `
        -ShareProperties "oplocks","browsable","changenotify","continuously-available" `
        | Add-NcCifsShareAcl -UserorGroup $UserGroup -Permission "full_control";
        #Remove Everyone Account
        Remove-NcCifsShareAcl -VserverContext $vserver $vol Everyone

}
Add-NcCifsPrivilege -Name $installer -vservercontext $vserver -Privilege sesecurityprivilege
```

Version History

Version	Date	Document Version History
Version 1.0	December 2013	Initial release

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