

Deploying Microsoft Exchange Server 2016 on Dell EMC XC Series using Hyper-V

A fully supported, scalable, hyper-converged solution for Microsoft Exchange Server 2016 deployment

Dell EMC Storage Engineering
January 2017

Revisions

Date	Description
November 2016	Initial release
January 2017	Branding update

Acknowledgements

Author: Daniel J. Curran

Special thanks to: Mark Boeser and Chidambara Shashikiran

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2016 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA.

Dell EMC believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Table of contents

Revisions.....	2
Acknowledgements.....	2
Executive summary.....	4
1 Product overview.....	5
1.1 Dell EMC XC Series overview.....	5
1.2 Dell EMC XC Series and Microsoft Exchange.....	6
1.3 Software-defined storage (SDS) with Nutanix.....	6
1.3.1 Server Message Block (SMB).....	6
1.3.2 SMB with Nutanix.....	7
2 Solution infrastructure.....	8
2.1 Physical system configuration.....	8
2.2 Dell EMC XC Series storage and cluster configuration.....	8
2.3 VM layout.....	9
2.4 Nutanix storage Controller VMs (CVMs).....	10
2.4.1 CVM disk performance balancing.....	10
2.5 Hyper-V hypervisor processor performance.....	10
3 Configuration steps.....	11
3.1 XC Series installation and configuration.....	11
3.1.1 Hyper-V and Nutanix pre-installation and configuration.....	11
4 XC Series cluster installation.....	14
5 Networking overview.....	15
6 Storage overview.....	17
6.1.1 Server Message Block (SMB).....	17
6.1.2 SMB Transparent Failover.....	17
6.1.3 Creating and connecting to SMB shares.....	17
6.1.4 Setting the default storage location with Hyper-V Manager.....	18
6.1.5 Windows failover cluster.....	18
7 Windows Server VM installation and configuration.....	20
7.1 Exchange 2016 simulated DAG servers installation and configuration.....	21
8 Conclusion.....	26
A Additional resources.....	27
A.1 Technical support and resources.....	27

A.2	Dell EMC XC Series solutions	27
A.3	Related documentation	27

Executive summary

The Dell EMC™ XC Series Web-Scale Hyper-converged appliance powered by Nutanix™ delivers a highly resilient, converged compute and storage platform that brings benefits of web-scale architecture to business-critical enterprise applications such as Microsoft® Exchange.

Dell EMC XC Series software installs quickly for deployment of multiple virtualized workloads. The Acropolis Distributed Storage Fabric (DSF) virtualizes local storage from all nodes into a unified storage pool. A container, which is a Nutanix construct, presents this pool as shared storage resources. On top of a storage pool, an administrator can create one or more containers. Containers mount to Microsoft Hyper-V® hosts in the form of Server Message Block (SMB) based file shares.

This paper details the configuration of a simulated multi-Exchange Server 2016 workload on a representative Dell EMC XC Series cluster using SMB-based file shares in Hyper-V.

1 Product overview

1.1 Dell EMC XC Series overview

The Dell EMC XC Series is a hyper-converged solution that combines storage, compute, networking, and virtualization into an industry-proven x86 Dell PowerEdge™ server running Nutanix web-scale software. By combining the hardware resources from each server node into a shared-everything model for simplified operations, improved agility, and greater flexibility, Dell EMC and Nutanix can deliver simple, cost-effective solutions for enterprise workloads. DSF delivers a unified pool of storage from all nodes across the cluster, using techniques including striping, replication, autotiering, error detection, failover, and automatic recovery. A container, which is a Nutanix construct, presents this pool as shared storage resources. Containers mount to hypervisor hosts in the form of Network File System (NFS) or, in the case of Hyper-V, SMB-based file shares; virtual machine data is stored on these containers. A Nutanix storage pool can support multiple containers, each with its own settings for options such as compression, deduplication, or resiliency factor (RF).

The Dell EMC XC Series infrastructure is a scale-out cluster of high-performance nodes, or servers, each running a standard hypervisor and containing processors, memory, and local storage (consisting of SSD flash and high-capacity SATA disk drives). Each node runs virtual machines just like a standard hypervisor host as displayed in Figure 1.

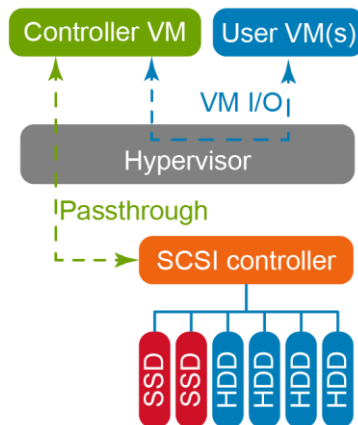


Figure 1 Nutanix node architecture

In addition, the DSF virtualizes local storage from all nodes into a unified pool. DSF uses local SSDs and disks from all nodes to store virtual machine data. Virtual machines running on the cluster write data to DSF as if they were writing to shared storage.

1.2 Dell EMC XC Series and Microsoft Exchange

Acropolis v4.6 supports SMB-based file shares with Hyper-V hypervisors running Microsoft Windows Server® 2012 R2. This operating system was used to validate the Microsoft Exchange Server 2016 solution on the XC Series appliance.

The support for SMB 3.0 allows Microsoft Exchange Server solutions using these SMB share volumes for databases to be fully supported by Microsoft product support. This makes the Dell EMC XC Series an excellent choice for scalable, virtualized Exchange Server solutions using the power of Dell servers with the flexible features of Nutanix and SDS.

1.3 Software-defined storage (SDS) with Nutanix

Nutanix software provides a hyper-converged platform that uses DSF to share and present local storage to all the virtual machines in the cluster. DSF presents the storage in Microsoft Hyper-V environments using SMB-based file shares. Nutanix uses a construct called a container which is presented as an SMB share and mounted to the Microsoft Hyper-V servers as a datastore. VMs store their virtual disks within the datastore. The general architecture is shown in Figure 2.

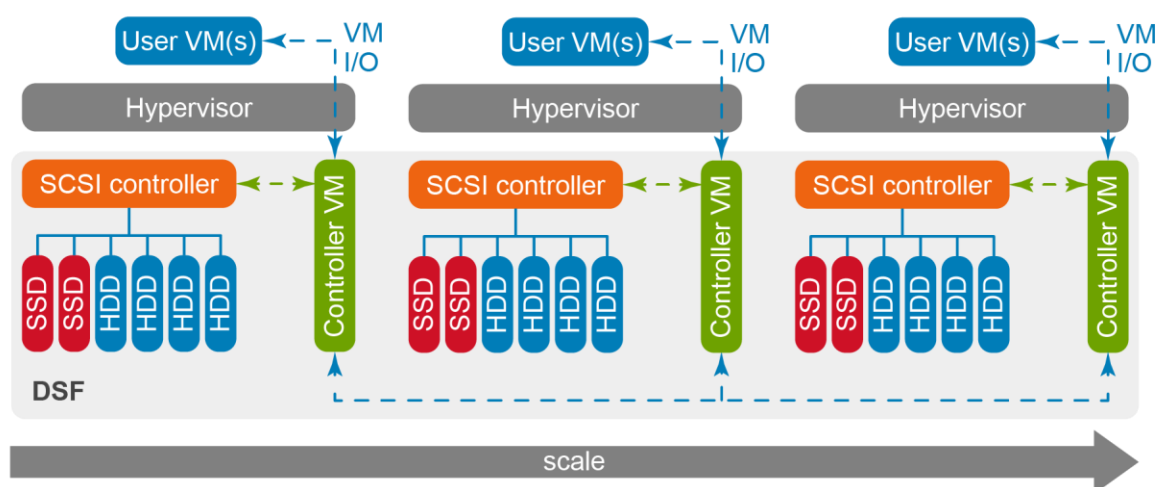


Figure 2 Nutanix architecture using SMB-mounted virtual disks

This paper focuses on using SMB-based file shares as the storage presentation protocol that connects to the Nutanix containers and DSF storage resources.

1.3.1 Server Message Block (SMB)

SMB is an application-layer network protocol typically used for providing access to shared files within a network. With Windows Server 2012, Microsoft introduced support for version 3.0 of the SMB protocol. Windows Server 2012 R2 was updated to support the 3.0.2 version. SMB 3.0 introduced several core features, including SMB Transparent Failover, SMB Direct, and SMB Multichannel to improve the availability and performance of the protocol.

1.3.2 SMB with Nutanix

The Nutanix platform uses SMB 3.0 to present all shared storage in Hyper-V environments. The Nutanix SMB share, exported globally by each Nutanix Controller VM (CVM), is presented directly to the Hyper-V server to host virtual machine files.

The unique architecture of Nutanix enables the processing of all SMB 3.0 traffic locally to the Hyper-V server hosting the virtual machine that is performing the input/output (I/O) operation. This SMB storage I/O traffic occurs over the internal virtual NIC and internal virtual switch. Because of the localized, in-memory SMB traffic, implementing certain features of the SMB 3.0 protocol, such as SMB Direct or SMB Multichannel, are not required on the Nutanix platform.

2 Solution infrastructure

The configuration and solution components are described in this section.

2.1 Physical system configuration

The physical configuration for this environment starts with the basic, three-node Dell EMC XC Series cluster shown in Figure 3.

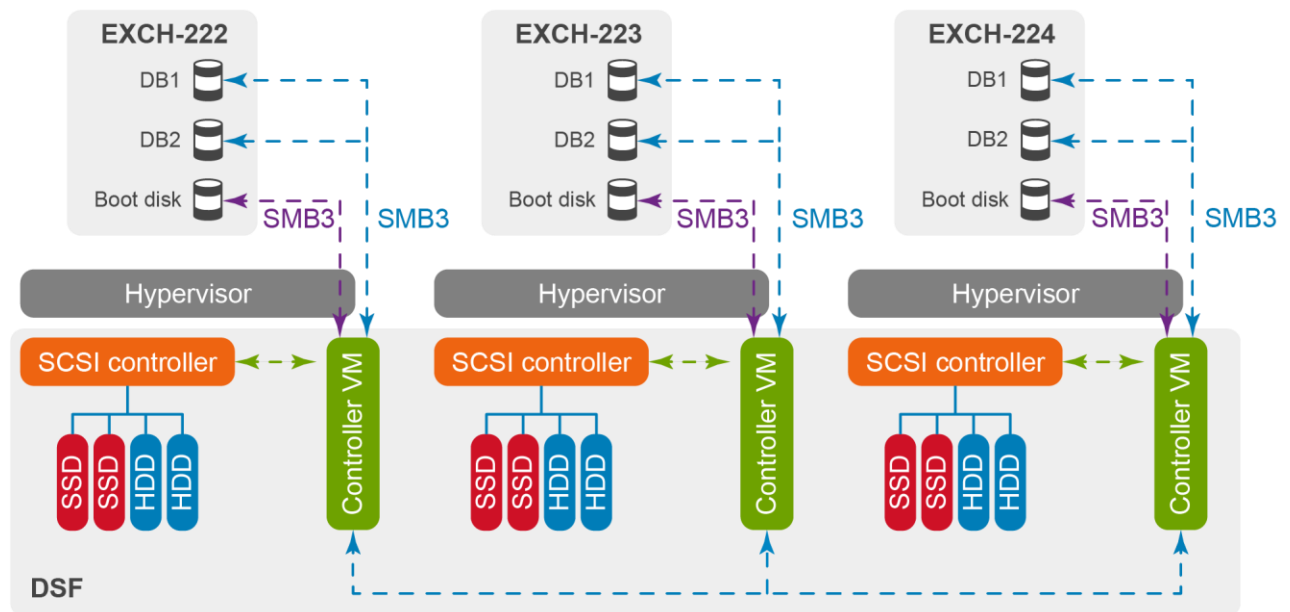


Figure 3 Three-node Dell EMC XC Series cluster representation

2.2 Dell EMC XC Series storage and cluster configuration

Each Dell EMC XC430 Series node used in this configuration is comprised of the following hardware components:

- Two 400GB SATA SSDs
- Two 2TB 7.2K RPM SATA hard disks
- Two 16-core Intel® Xeon® E5-2630 v3 2.40GHz processors
- Twelve 16GB DDR-4 QR 2133MHz RAM modules (192GB total)

The local storage controller on each host ensures that storage performance as well as storage capacity increases when additional nodes are added to the Dell EMC XC Series appliance. Each CVM is directly connected to the local storage controller and its associated disks. By using local storage controllers on each Hyper-V host, access to data through DSF is localized. It does not require read data to be transferred over the network, thereby improving latency. DSF ensures that writes are replicated, distributing data within the platform for resiliency.

The minimum number of Dell EMC XC Series nodes in a cluster is three. When clustered together, the storage across three nodes are virtualized together to create a single storage pool and one or more containers can be created on top of the storage pool. The storage container is presented to all nodes as shared storage within the cluster.

Every attempt is made by the cluster to keep virtual machines and their associated storage on the same cluster node for performance consistency. However, each cluster node is connected to, and communicates with, the other nodes on a 10Gb network. This communication allows virtual machines and their associated storage to reside on different cluster nodes.

2.3 VM layout

The baseline for this study was executed using three XC430 Series cluster nodes (the minimum number of nodes in an XC cluster). The workload applied to the cluster was in the form of three virtual machines running Microsoft Windows Server 2012 R2 and Microsoft Exchange 2016 with one virtual machine per node.

Each virtual machine consisted of two virtual processors, 32GB RAM, and thin-provisioned disks in the layout described in Table 1.

Table 1 Microsoft Exchange disk layout

Size	Disk	Description
100GB	Drive C:	Windows OS disk (SMB)
500GB	Volume: E:\DB1	Database disk (SMB)
500GB	Volume: F:\DB2	Database disk (SMB)

Note: Processor and memory sizing is based on the minimum requirements for testing and not for a production environment. For more sizing information, see the [Exchange Server Role Requirements Calculator](#).

2.4 Nutanix storage Controller VMs (CVMs)

The Nutanix Enterprise Cloud Platform is a scale-out cluster of high-performance nodes, or servers, each running a standard hypervisor and containing processors, memory, and local storage consisting of SSD flash and high-capacity SATA disk drives. Along with the hypervisor, every node runs a Nutanix CVM to which the local storage resources for each server are presented. Virtual machines running on each node can access their storage through the CVM.

2.4.1 CVM disk performance balancing

IOPs generated by the CVM storage controllers were monitored and showed that the disk IOPs were distributed among the three XC Series CVMs.

One aspect of the SMB 3.0 protocol implemented on the Nutanix platform is Transparent Failover (also called continuous availability). If a CVM becomes unavailable on a given Hyper-V server, the routing table for that server updates to reference the SMB server instance running within a CVM on another node in the cluster. Nutanix refers to this CVM storage controller redirection as data path redundancy. Data path redundancy allows virtual machines running on a particular node where the local CVM has become unavailable to continue to access the SMB share without significant interruption. To do this, the file handles associated with the virtual machines need to be maintained across the CVM instances within the cluster. SMB Transparent Failover allows this file handle redirection across CVM instances. This process allows for minimal interruption to virtual machine storage traffic during CVM failure or upgrade scenarios.

2.5 Hyper-V hypervisor processor performance

During the study, the processor utilization on each hypervisor was consistent with the workload placed on both the EXCH server VM and the CVMs supported by that host.

3 Configuration steps

3.1 XC Series installation and configuration

The following section details the system configuration and settings required for the Exchange workload testing.

3.1.1 Hyper-V and Nutanix pre-installation and configuration

Starting with the release of Windows Server 2012 R2, Nutanix supports Hyper-V. Windows Server 2012 R2 is installed on a given node during the imaging process performed by the Nutanix Foundation software or directly during the node imaging process. Windows is installed on a dedicated storage device called the SATA DOM on each node. Nutanix currently supports the installation of the server core mode in combination with the following Windows editions:

- Windows Server 2012 R2 Datacenter
- Windows Server 2012 R2 Standard
- Hyper-V Server 2012 R2

The following steps were completed:

1. Microsoft Hyper-V preinstalled on the SATA DOMs of three Dell EMC XC430-4 Series nodes
2. Nutanix Acropolis 4.6 preinstalled on XC430-4 nodes
3. Windows VMs stored through SMB 3.0 shares hosted on DSF
4. 1TB infrastructure container created in DSF storage to house Hyper-V datastore infrastructure to save all EXCH Windows VMs

Along with the Windows installation, Foundation installs the Nutanix CVM and can be configured to create the Nutanix cluster. Additionally, Foundation performs the following steps for configuring Windows 2012 R2 and Hyper-V:

1. Install and configure Windows server roles including:
 - Hyper-V
 - Failover Clustering
5. Install drivers as required by specific server models.
6. Install Nutanix services to enable communication with Windows and Hyper-V for monitoring and integration purposes.
7. Modify firewall settings to allow WMI (Windows Management Instrumentation), file and print sharing, and WinRM, as required by System Center Virtual Machine Manager (SCVMM) and Nutanix communication between the CVM and host.
8. Configure Windows networking as outlined in the Networking overview section.

Upon completion of the initial imaging process, each host has Hyper-V installed with a CVM configured as a standalone virtual machine as shown in Figure 4. The default CVM naming convention starts with "NTNX", followed by the block serial number, then the node position in the block (as shown in the example figure), and ends with "CVM".




Figure 4 Nutanix CVM

As the Figure 5 shows, the CVM has all storage outside the SATA-DOM presented to it as physical drives, sometimes referred to as pass-through disks. The system marks these drives as “offline” to the Hyper-V host as they are used to create the Nutanix storage pool.

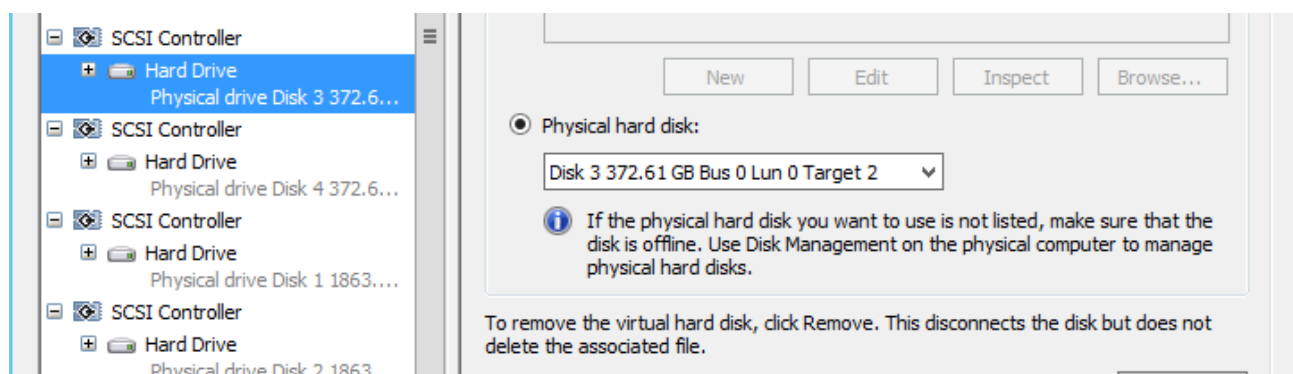


Figure 5 Physical hard disks presented to CVM

At this point, the hosts are ready for final configuration. This can include storage pool and container creation. Additionally, the Windows servers are online and ready for configuration for Windows failover clustering.

To assist with these final configuration steps, Nutanix offers a Hyper-V setup script named **setup_hyperv.py** to automate the process. Starting with release 4.5 of the Acropolis-based software, Nutanix has integrated most of the operations previously performed by the setup_hyperv.py script into the GUI-based Prism interface. Upon logging into Prism, the two steps shown in Figure 6 appear.




Figure 6 Prism Hyper-V setup

The first step joins the Nutanix cluster and the Hyper-V hosts to the specified domain. This step includes configuring the Nutanix cluster:

- Creating a domain object for the Nutanix cluster
- Securely creating an **A** record in DNS for the Nutanix cluster
- Setting the domain controller(s) in the specified domain as the NTP servers for the Nutanix cluster

The second step creates the Windows Server Failover Cluster (WSFC), which includes adding all nodes to the cluster and configuring the default external network for client and live migration use. This step includes configuring the Hyper-V hosts:

- Renaming the Hyper-V hosts and joining them to the specified domain
- Configuring the network adaptor load balancing and failover (LBFO) team
- Configuring DNS

In addition to the two setup steps, there are additional features built into Prism that are specific to the Hyper-V environments, including:

- Enabling Kerberos support for SMB authentication
- Setting a specific Nutanix container as the default SMB share on some or all of the Hyper-V hosts as shown in Figure 7

Create Container ? X

Enter a name for your container and select a storage pool for it. You can provision the container for all hosts, or select individual hosts.

NAME

STORAGE POOL

default-storage-pool-2946274070687002963 +

MAX CAPACITY

4.32 TiB (Physical) Based on storage pool free unreserved capacity

MAKE THIS CONTAINER THE DEFAULT STORE FOR VMS ON HYPER V HOSTS

☐ None

☒ Make default on all Hyper V hosts

☐ Make default on particular Hyper V hosts

Advanced Settings Cancel Save

Figure 7 Container as the default SMB share

4 XC Series cluster installation

A three-node Dell EMC XC430-4 Series cluster was installed with the following specifications.

Feature	Specification	
RAM (per node)	12 x 16GB DIMMS = 192GB	
Processors (2)	Intel Xeon processor E5-2630 v3 @ 2.40GHz, 8 cores Intel Xeon processor E5-2630 v3 @ 2.40GHz, 8 cores	
Network(s)	Management: Broadcom® Gigabit Ethernet BCM5720 embedded NIC (1 of 4) iSCSI: Intel Ethernet 10G 2P X520 adapter x 2	
Shared storage (managed by DSF)	Container: EXCHDATA	Container: Infrastructure
	Datastore: EXCHDATA Free space (logical): 12.28TiB Used: 1.76TiB Max capacity: 14.04TiB Replication factor: 2	Hyper-V datastore: Infrastructure VMs: 3 Free space (logical): 889.69GiB Used: 110.31GiB Max capacity: 0.98TiB Explicit reserved: 0.98TiB Thick provisioned: 100GiB DSF replication factor: 2

5 Networking overview

During the imaging process with Hyper-V, two virtual switches are created on every host. One switch manages network communication between the Nutanix CVM and the Hyper-V host. This switch is called **InternalSwitch**, and it is configured as an internal virtual switch type. The second switch manages external communication between the virtual machines and between the Hyper-V hosts. Its networking traffic also includes CVM replication traffic for the purposes of maintaining resilience factor (RF), as well as any DSF traffic that cannot be processed locally. This second switch is configured as an external virtual switch and is called **ExternalSwitch**. The external switch is assigned a network adapter team as the means to provide connectivity outside of the host. Assuming the 10Gb NICs are used, Figure 8 details the default virtual switch configuration.

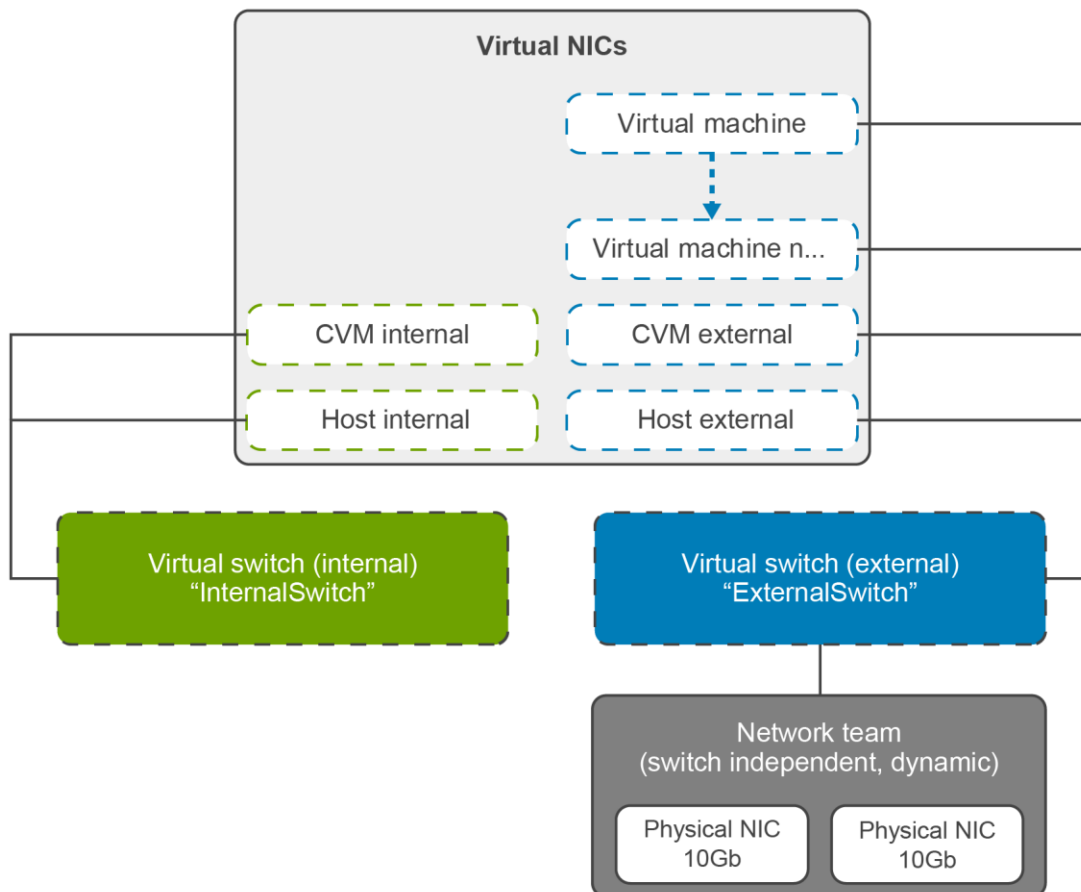


Figure 8 Default Hyper-V networking with Nutanix

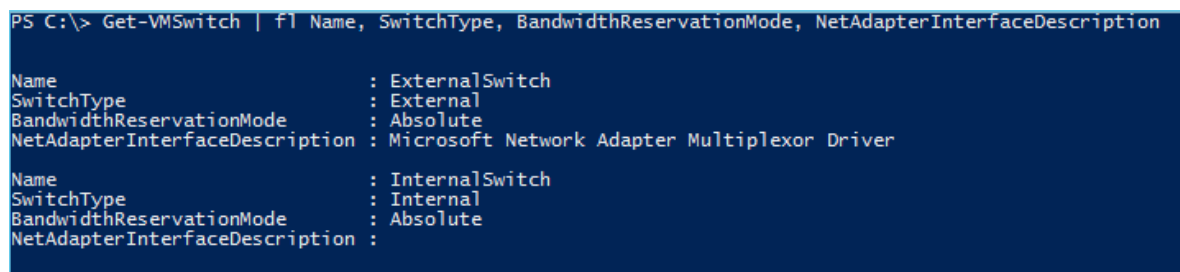
The Windows host and the SMB file share supporting the virtual disks communicate through virtual NICs associated with the internal network. This allows virtual machines that are local to a given node to access their virtual disks through the local CVM.

Virtual machine network traffic, Windows traffic between cluster nodes, and CVM traffic between cluster nodes use the virtual NICs associated with the external network.

You can view the default networking configuration using standard Windows tools, including Hyper-V Manager and PowerShell. The following PowerShell commands manage the virtual switches:

Add | Get | New | Remove | Rename | Set-VMSwitch

The PowerShell example in Figure 9 shows how to view the default configuration, including the switch type and the default bandwidth reservation mode.



```
PS C:\> Get-VMSwitch | fl Name, SwitchType, BandwidthReservationMode, NetAdapterInterfaceDescription

Name                : ExternalSwitch
SwitchType           : External
BandwidthReservationMode : Absolute
NetAdapterInterfaceDescription : Microsoft Network Adapter Multiplexor Driver

Name                : InternalSwitch
SwitchType           : Internal
BandwidthReservationMode : Absolute
NetAdapterInterfaceDescription :
```

Figure 9 Default virtual switch configuration

6 Storage overview

6.1.1 Server Message Block (SMB)

The Nutanix platform uses SMB 3.0 to present all shared storage in Hyper-V environments. SMB 3.0 introduced several core features, including SMB Transparent Failover, SMB Direct, and SMB Multichannel, to improve the availability and performance of the protocol. The SMB share, exported globally by each CVM, is presented directly to the Hyper-V server to host virtual machine files.

The unique architecture of Nutanix enables the processing of all SMB 3.0 traffic locally to the Hyper-V server hosting the virtual machine that is performing the input/output (I/O) operation. This SMB storage I/O traffic occurs over the internal virtual NIC and internal virtual switch. Because of the localized, in-memory SMB traffic, implementing certain features of the SMB 3.0 protocol, such as SMB Direct or SMB Multichannel, are not required on the Nutanix platform.

6.1.2 SMB Transparent Failover

One aspect of the SMB 3.0 protocol implemented on the Nutanix platform is Transparent Failover. Considered a requirement for Hyper-V environments, the continuously available setting (which controls transparent failover) is enabled by default with Nutanix, as shown in Figure 10.

```
PS C:\> Get-SmbConnection | where ShareName -eq ExchData | ft ServerName, ShareName, Dialect, ContinuouslyAvailable -AutoSize
```

ServerName	ShareName	Dialect	ContinuouslyAvailable
EXCHHV.xclab.local	ExchData	3.00	True
EXCHHV.xclab.local	ExchData	3.00	True

Figure 10 Continuously available setting

If a CVM becomes unavailable on a Hyper-V server, the routing table for that server updates to reference the SMB server instance running within a CVM on another node in the cluster.

6.1.3 Creating and connecting to SMB shares

As discussed in the Dell EMC XC Series overview section, storage is presented to the Hyper-V hosts through the use of containers. Containers offer the ability to carve out the resources available within a storage pool and apply specific settings to that unit of storage.

In Hyper-V environments, containers map to specific SMB shares. Each container represents a unique Universal Naming Convention (UNC) path in the form of `\\<NutanixClusterName>\<ContainerName>`. The Nutanix cluster name represents the SMB server name and is registered within Active Directory during the Hyper-V setup process. The Nutanix cluster name is also registered in DNS using an external IP address. This external IP address is not used for Hyper-V host connectivity to the share. As discussed earlier, communication with the Nutanix cluster occurs over the internal network. Presenting the containers to the Hyper-V hosts is as simple as supplying the UNC of the share. You can use PowerShell, Hyper-V manager, or SCVMM to direct virtual machine files onto specific shares. In all cases you should use the cluster name and not an IP address when supplying the share path for virtual machine storage.

6.1.4 Setting the default storage location with Hyper-V Manager

The following is an example of setting up a default storage location with Hyper-V Manager. To set up a default storage location:

1. Log in to Hyper-V Manager.
2. Right-click **Hyper-V Manager** and select **Connect to Server**.
3. Type the name of the host you want to add and click **OK**.
4. Right-click the host and select **Hyper-V Settings**.
5. Select **Virtual Hard Disks** and specify the default folder (see Figure 11) to store virtual hard disk files. As an example, use: `\\<NutanixClusterName>\<ContainerName>`

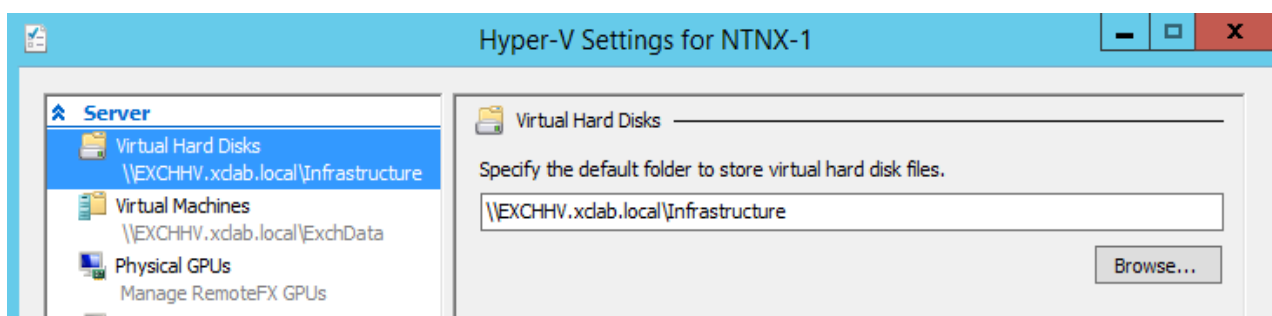


Figure 11 Default folder for virtual hard disk files

6. Click **Virtual Machines** and specify the default folder to store virtual machine configuration files.

6.1.5 Windows failover cluster

In Hyper-V environments, a Nutanix cluster is formed by combining the storage associated with a group of CVMs and exporting that storage in the form of an SMB share. The Nutanix cluster is independent from the Windows failover cluster that is formed to support Hyper-V. Although it is possible to have standalone hosts, the best practice calls for forming Windows failover resources with the underlying Nutanix file shares acting as the shared storage repository. As clustered resources with the support of DSF, the virtual machines can be actively moved between hosts for load balancing, planned maintenance, or restarts during unplanned events.

Virtual machines created with Hyper-V Manager or with the **New-VM** PowerShell command are configured as standalone virtual machines. You can run a PowerShell command to check a virtual machines clustered state by looking at the **IsClustered** attribute of the virtual machines running on a host as shown in Figure 12.




Figure 12 Clustered state of virtual machines

After creating a standalone virtual machine, add it to the cluster by running the **Add-ClusterVirtualMachineRole** PowerShell command. Failover Cluster Manager can also convert an existing virtual machine through the Configure Role high-availability wizard, or it can create a new clustered virtual machine. The Failover Cluster Manager options are shown in Figure 13.

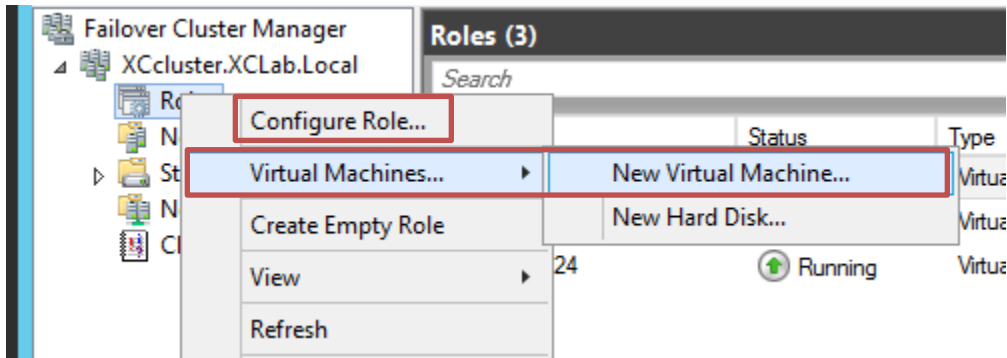


Figure 13 Failover Cluster Manager virtual machine options

Once a virtual machine is clustered, any changes to settings or manipulation to the power state of the virtual machine should be performed from the Failover Cluster Manager or SCVMM and not from Hyper-V Manager.

7 Windows Server VM installation and configuration

Perform the following steps to install and configure a Windows Server VM.

1. From the **New Virtual Machine Wizard**, install **Windows Server 2012 R2** and choose a disk type of **VHDX, fixed**.
2. When the new VM is created, from **Hyper-V Manager > Virtual Machines** panel, select the new VM. Under the **Actions** panel, select **Settings**.
3. Under the **Hardware** panel, select **Add Hardware** and select **SCSI Controller**. Under the **Hardware** panel, click **Add**.
4. Under the **Hardware** panel, select the newly added SCSI controller. Under the **SCSI Controller** panel, click **Hard Drive** and click **Add**.
5. Under the **Hard Drive** panel > **Virtual Hard Disk**, click **New** to open the **New Virtual Hard Disk Wizard** window.
6. Select **Next**, select **Fixed size**, and click **Next** to open the **Specify Name and Location** window.
7. Provide a descriptive name (222-DB1.VHDX in this example). Under **Location**, select **Browse**, select the SMB share (\\EXCHHV.xclab.local\\ExchData in this example), and click **Next**.
8. Under **Configure Disk**, select **Create a New Blank Virtual Hard Disk** (500GB in this example).
9. Click **Next** and review the **Summary** page. Verify the information and click **Finish** to create.
10. Repeat steps 2 through 9 to create another SCSI controller and another 500GB virtual hard disk from the same SMB share. Name this new virtual hard disk **222-DB2.VHDX**.
9. Start the VM from **Hyper-V Manager > Virtual Machines** panel, start the VM, and under the **Actions** panel, select **Connect**.
10. Perform Windows updates.
11. Use **Windows Disk Management** to initialize, mount, and format database disk volumes (DB1, DB2) using 64K NTFS.

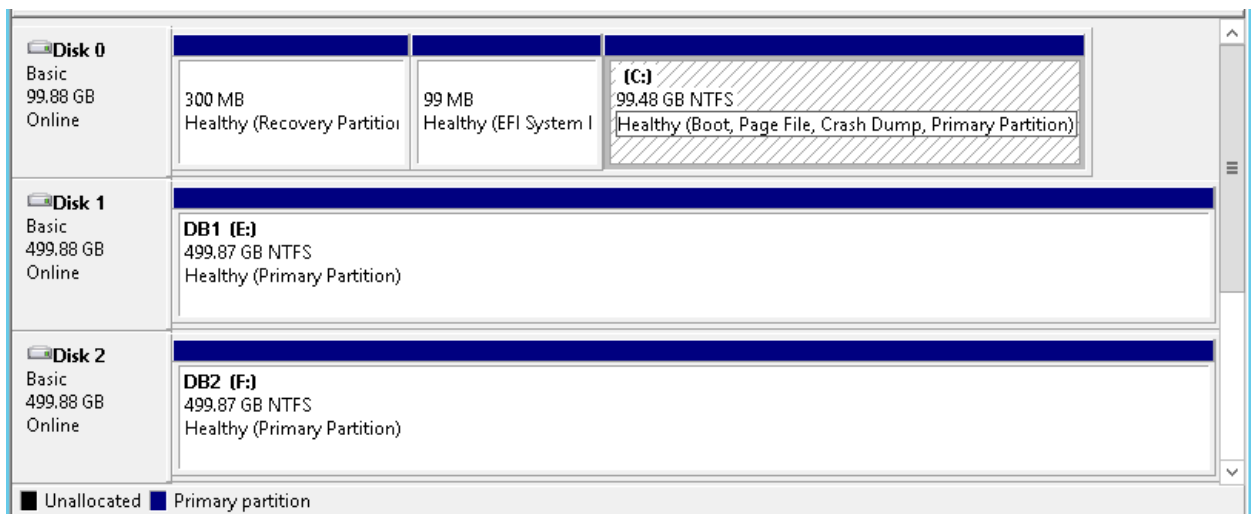


Figure 14 Windows Disk Management

7.1 Exchange 2016 simulated DAG servers installation and configuration

Microsoft supports Exchange Server 2016 storage workloads that are presented using SMB shares to the virtual machine, as stated in the Microsoft TechNet best practices article, [Exchange 2016 virtualization](#).

1. To calculate proper storage configuration for testing, use the [Microsoft Exchange Server role calculator worksheet](#) for Exchange configuration. The following Exchange configuration was used to simulate a real-world Exchange workload using the Jetstress tool:
 - 3 servers
 - 2 database copies
 - 900 simulated mailboxes
 - .06 mailbox profile
 - 2 GB mailbox size quota

Note: The configuration listed in step 1 was not designed to push the limits of the system and achieve maximum performance. See appendix A.3 for links to other performance-focused collateral.

12. Install the latest Jetstress version (2013) on the Exchange Server VMs, copying the Exchange Server 2016 **ese.dll** and **esepperf.*** files to the folder, **C:\Program Files\Exchange Jetstress**. Exchange Server 2016 storage DLLs were used with this latest version of Jetstress 2013 per Microsoft published guidance.
13. Create two Exchange Jetstress databases per Exchange VM server:
 - EXCH-222
 - > 222-DB1
 - > 222-DB2
 - EXCH-223
 - > 223-DB1
 - > 223-DB2
 - EXCH-224
 - > 224-DB1
 - > 224-DB2

Next, configure Jetstress for the selected mailbox profile and quota and prepare the databases for testing.

1. Create a new test configuration.

The screenshot shows the 'Open Configuration' window in the Microsoft Exchange Jetstress 2013 application. The left sidebar contains a navigation menu with the following items: Welcome, Open Configuration (selected), Define Test Scenario, Define Database Configuration, Select Database Source, Review & Execute Test, and a 'See also' section with links to Jetstress Help and About Jetstress. The main content area has the title 'Open Configuration' and two radio button options: 'Create a new test configuration' (selected) and 'Open an existing configuration file'. Both options have a text input field containing the path 'C:\Program Files\Exchange Jetstress\JetstressConfig.xml'. At the bottom of the main area are 'Back' and 'Next' buttons.

14. Test an Exchange mailbox profile. To show an initially small workload for this model, a profile with 300 mailboxes and .06 IOPS per mailbox with a 2 GB quota was chosen. This can easily scale up as demands require.

The screenshot shows the 'Define Test Scenario' window in the Microsoft Exchange Jetstress 2013 application. The left sidebar is identical to the previous screenshot, with 'Define Test Scenario' selected. The main content area has the title 'Define Test Scenario' and a 'Select Test Category' section with two radio button options: 'Test disk subsystem throughput' and 'Test an Exchange mailbox profile' (selected). Below this is a 'Describe test scenario' section with a text input field containing the text: '300 Mailboxes', '.06 profile', and '2GB Mailboxes'. At the bottom of the main area are 'Back' and 'Next' buttons.

15. Set a small workload of three threads to generate the target I/O.

The screenshot shows the 'Exchange Mailbox Profile' configuration window in the Microsoft Exchange Jetstress 2013 application. The left sidebar contains a navigation menu with the following items: Welcome, Open Configuration, Define Test Scenario (highlighted), Define Database Configuration, Select Database Source, Review & Execute Test, and a 'See also' section with links to Jetstress Help and About Jetstress. The main content area is titled 'Exchange Mailbox Profile' and contains the following settings:

Number of mailboxes	300
IOPS/Mailbox	0.06
Mailbox size (MB)	2048
<input checked="" type="checkbox"/> Suppress tuning and use thread count (global)	3

At the bottom of the main area, there are two green buttons: 'Back' and 'Next'.

16. Select the performance test with multi-host and background database maintenance.

The screenshot shows the 'Select Test Type' configuration window in the Microsoft Exchange Jetstress 2013 application. The left sidebar is identical to the previous screenshot, with 'Define Test Scenario' highlighted. The main content area is titled 'Select Test Type' and contains the following options:

- ☒ Performance
- ☐ Database backup
- ☐ Soft recovery
- ☒ Multi-host test
- ☒ Run background database maintenance
- ☐ Continue the test run despite encountering disk errors

At the bottom of the main area, there are two green buttons: 'Back' and 'Next'.

17. Define a test run of two hours.

The screenshot shows the 'Define Test Run' window in Microsoft Exchange Jetstress 2013. The left sidebar contains a navigation menu with the following items: Welcome, Open Configuration, Define Test Scenario (highlighted), Define Database Configuration, Select Database Source, Review & Execute Test, and a 'See also' section with links to Jetstress Help and About Jetstress. The main area is titled 'Define Test Run' and contains the following fields: 'Output path for test results' with a text box containing 'C:\Program Files\Exchange Jetstress' and a 'Browse' button; 'Test duration (hours)' with a dropdown menu set to '2'; and 'Back' and 'Next' buttons at the bottom.

18. Define the database configuration with two database volumes.

The screenshot shows the 'Define Database Configuration' window in Microsoft Exchange Jetstress 2013. The left sidebar is identical to the previous screenshot, with 'Define Database Configuration' highlighted. The main area is titled 'Define Database Configuration' and contains the following fields: 'Number of databases' with a dropdown menu set to '2'; 'Number of copies per database' with a dropdown menu set to '2'; and two tables. The first table, 'Database', has columns 'Name', 'Database Path', and an expand/collapse icon. It contains two rows: 'Database1' with path 'E:\DB1' and 'Database2' with path 'F:\DB2'. The second table, 'Log', has columns 'Name', 'Log Path', and an expand/collapse icon. It contains two rows: 'Log1' with path 'E:\DB1' and 'Log2' with path 'F:\DB2'.

	Name	Database Path	...
▶	Database1	E:\DB1	
	Database2	F:\DB2	

	Name	Log Path	...
▶	Log1	E:\DB1	
	Log2	F:\DB2	

19. Create new databases.

The screenshot shows the 'Select Database Source' dialog box in the Microsoft Exchange Jetstress 2013 application. The left sidebar contains a list of steps: Welcome, Open Configuration, Define Test Scenario, Define Database Configuration, Select Database Source (highlighted), Review & Execute Test, Jetstress Help, and About Jetstress. The main area has three radio buttons: 'Create new databases' (selected), 'Attach existing databases', and 'Restore backup database'. Below these is a text field for 'Backup path (for create and restore)' with a 'Browse...' button. At the bottom are 'Back' and 'Next' buttons.

20. Save the test to save the configuration details to the XML file and then prepare the test to create the databases.

The screenshot shows the 'Review & Execute Test' dialog box in the Microsoft Exchange Jetstress 2013 application. The left sidebar is identical to the previous screenshot, with 'Review & Execute Test' highlighted. The main area displays a 'Test Scenario and Exchange Profile Summary' with the following details: Test Scenario: Exchange Mailbox Profile Test, Test type: Performance, Run Database Maintenance: True, Test duration: 02:00:00, Mailbox Profile: mailbox count 300, mailbox quota 2048, mailbox lops 0.06, Suppress tuning: True, ThreadCount: 3, Output path: C:\Program Files\Exchange Jetstress, Database source: CreateBrandNewDatabases, Number of copies per database: 2, Database paths: E:\DB1, F:\DB2, and Log paths: E:\DB1, F:\DB2. At the bottom are 'Prepare test', 'Run test', and 'Save test' buttons.

8 Conclusion

The tests results detailed in this paper show that Dell EMC XC Series appliances provide fast performance, easy scalability, simple management, and high availability for Microsoft Exchange environments. This paper also demonstrated how SMB 3.0 shares can be used to present storage disks directly to Windows guest VMs.

The Dell EMC XC Series solution supporting the Microsoft Exchange environment described in this document is:

Fully supported by Dell EMC: Dell EMC fully supports the XC Series platform hardware as a scalable, converged-application platform.

Fully supported by Nutanix: Nutanix fully supports the Nutanix platform in conjunction with SMB 3.0.

Fully supported by Microsoft: SMB 3.0 share-connected block storage for Exchange Server databases and log volumes is fully support by Microsoft Product support.

IT administrators deploying this solution can be confident about the performance, reliability, and support for this Microsoft Exchange solution.

Note: The tests results showcase the benefits of XC Nutanix hyper-converged infrastructure for enterprise applications like Microsoft Exchange. A set of Jetstress tests was executed to simulate real-world Exchange workload and validate the architecture. The tests were not designed to push the limits and achieve maximum performance. See appendix A.3 for links to Nutanix ESRP and best practices.

A Additional resources

A.1 Technical support and resources

Dell.com/Support is focused on meeting customer needs with proven services and support.

[Dell TechCenter](#) is an online technical community where IT professionals have access to numerous resources for Dell software, hardware and services.

[Storage Solutions Technical Documents](#) on Dell TechCenter provide expertise that helps to ensure customer success on Dell Storage platforms.

A.2 Dell EMC XC Series solutions

For XC Series technical content, visit [Dell XC Series Hyper-converged Appliances powered by Nutanix - Documentation](#).

A.3 Related documentation

See the following referenced or recommended resources related to this document:

- Mike McGhee (Nutanix) blog post: [How to Configure Nutanix Volume Groups to Support Windows Failover Clustering](#) for specifics on iSCSI setup with Nutanix
- Microsoft Best Practices: [Exchange 2016 virtualization](#)
- [Deploying Microsoft Exchange Server 2016 on Dell XC Series using iSCSI](#)
- [Nutanix 24000-Mailbox Virtualized Exchange ESRP](#)
- [Best Practices Guide: Hyper-V Windows Server 2012 R2 - Nutanix](#)