

Deploying Microsoft Exchange Server 2016 on Dell EMC XC Series using iSCSI

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

Dell EMC Storage Applications Engineering
January 2017

Revisions

Date	Description
February 2016	Initial release
November 2016	Terminology and branding updates
January 2017	Branding updates

Acknowledgements

Authors: Mark Boeser, 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 [2/6/2017] [Best Practice Guide] [3109-BP-SDS]

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 and Microsoft Exchange	5
1.2 Software-defined storage (SDS) with Nutanix.....	5
1.2.1 Presentation of iSCSI and NFS	6
1.2.2 Windows MPIO importance with in-guest iSCSI	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
3 Configuration steps	10
3.1 XC Series installation and configuration.....	10
3.1.1 VMware and Nutanix pre-installation and configuration	10
3.1.2 XC Series cluster installation.....	10
3.1.3 iSCSI vgroup and vdisk creation	11
3.1.4 Setting Jumbo Frames for iSCSI traffic	12
3.1.5 Setting Jumbo Frames for iSCSI traffic on VMware vSwitch	14
3.2 Windows Server VM installation and configuration	15
3.2.1 Windows MPIO with iSCSI installation and configuration	15
3.2.2 Setting Jumbo Frames for iSCSI traffic on Windows Server NIC	21
3.3 Exchange 2016 simulated DAG servers installation and configuration.....	21
4 Conclusion.....	27
A Additional resources.....	28
A.1 Technical support and customer service	28
A.2 Dell EMC XC Series solutions	28
A.3 Related documentation.....	28

Executive summary

The Dell EMC™ XC Series Web-Scale Hyper-converged appliance powered by Nutanix™ delivers a highly resilient, hyper-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. Previously, storage from the Nutanix Acropolis Distributed Storage Fabric (ADSF or DSF) could only be presented directly to virtual machines through NFS. Now, Nutanix software supports volume group vdisks to be connected to VMs through iSCSI with in-guest iSCSI initiators. This change is significant for Microsoft Exchange Server because NFS-connected virtual disks, although fully supported by Nutanix, are not supported by Microsoft. However, in-guest iSCSI is supported by Microsoft because the storage is presented as block-based storage to the virtual machines.

This paper details the configuration of this iSCSI support feature using a simulated multi-Exchange Server 2016 workload on a representative Dell EMC XC Series cluster.

1 Product 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 virtualized enterprise workloads.

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.




Figure 1 Nutanix node architecture

In addition, the DSF virtualizes local storage from all nodes into a unified pool. In effect, DSF acts like an advanced NAS that 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.

1.1 Dell EMC XC Series and Microsoft Exchange

Acropolis v4.5 supports in-guest iSCSI volumes or LUNs in virtual machines using the Microsoft iSCSI initiator. This operating system was used to validate the Microsoft Exchange Server solution on the XC Series appliance running the VMware® vSphere® 6 hypervisor.

The support for in-guest iSCSI allows Microsoft Exchange Server solutions using these iSCSI 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 software-defined storage (SDS).

1.2 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. Traditionally, DSF presented the storage in VMware ESXi® environments using the Network File System (NFS) protocol. Nutanix uses a construct called a container which is presented

as an NFS share and mounted to the VMware ESXi™ servers as a datastore. VMs store their virtual disks within the datastore. The general architecture is shown in Figure 2.

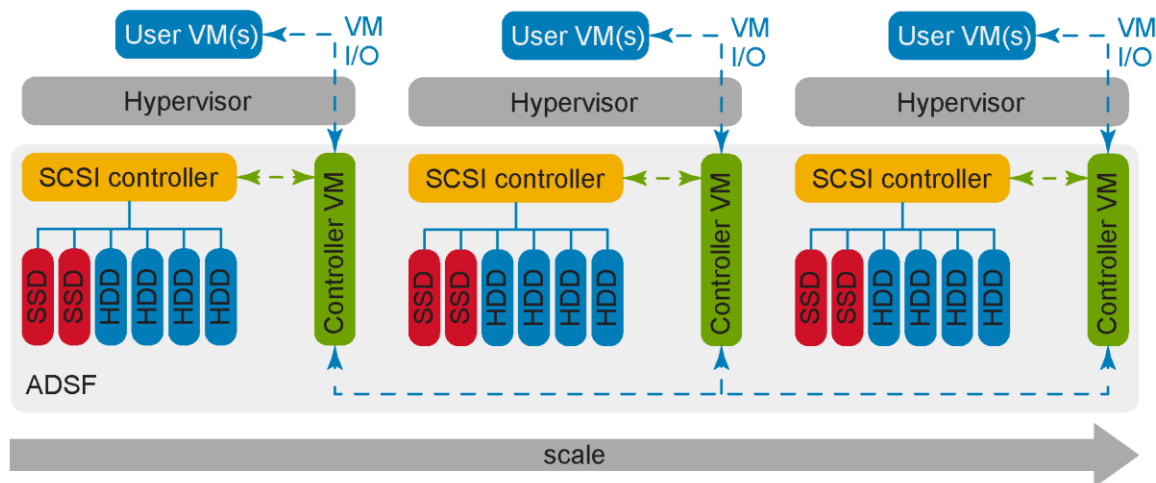


Figure 2 Standard Nutanix architecture using NFS-mounted virtual disks

For specific use cases, such as in-guest, shared-storage clustering or Microsoft Exchange, Nutanix also supports storage presentation using iSCSI. This allows Nutanix to present virtual disk resources directly to VMs to support these use cases.

This paper focuses on using iSCSI as the storage presentation protocol that connects through Microsoft in-guest iSCSI initiators and MPIO to the Nutanix containers and DSF storage resources.

1.2.1 Presentation of iSCSI and NFS

With the iSCSI connection model presented in this document, the Microsoft VM initiator connects directly to each Controller VM (CVM) as an iSCSI target portal*. Windows Server® MPIO, with the built-in DSM for iSCSI, performs the failover and load balancing of the I/O workload to the connected storage vgroups or vdisks. Multiple connections are made from each Windows VM to each of the unique Nutanix cluster CVMs.

***NOTE:** Nutanix has added a more simplified iSCSI connection model since this paper was originally written. Starting with AOS v4.7, Acropolis Block Services (ABS) is the new preferred method for iSCSI block storage connectivity running on Dell EMC XC Series. See the online Nutanix overview guidance for ABS [here](#) and enablement procedure [here](#). According to the Nutanix guidance, it is also still valid to use the MPIO connection method described in this paper.

At least two paths per disk are recommended to allow for path redundancy during Nutanix non-disruptive software upgrades, or during host-failure events. Three paths are used in this example shown in Figure 3.

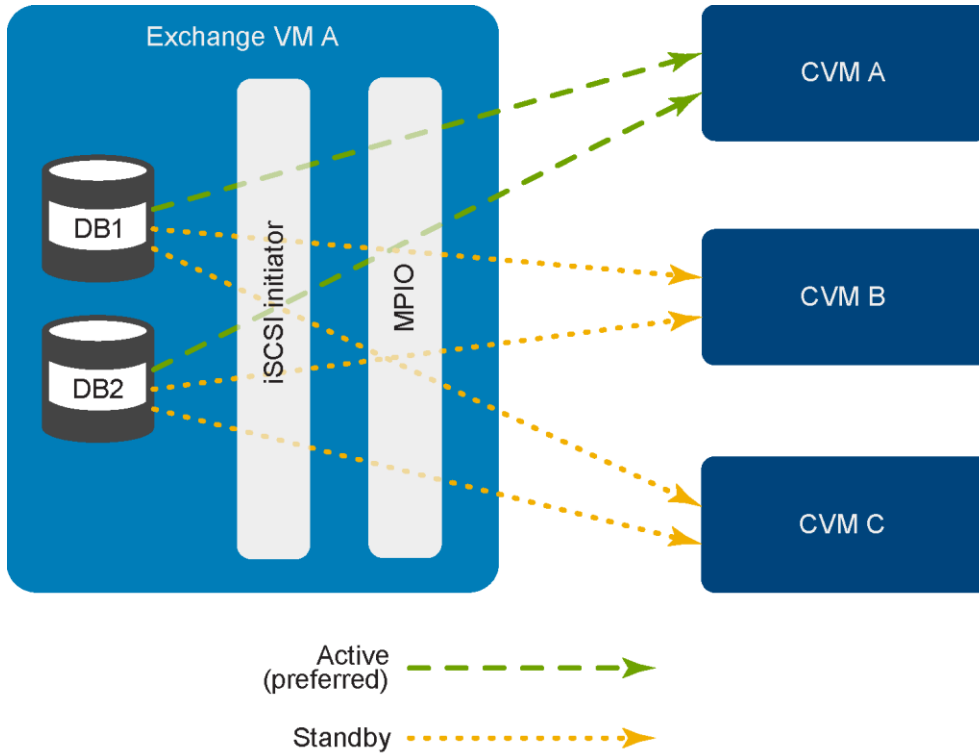


Figure 3 In-guest iSCSI initiator with Failover Only MPIO active and standby connections

1.2.2 Windows MPIO importance with in-guest iSCSI

In this configuration, the storage is being connected as block storage through iSCSI instead of a managed .vmdk file as it would be using NFS. This makes the proper iSCSI initiator and MPIO configuration very important to the overall loading and redundancy of the solution.

It was found that preparing the Exchange Server VM multipathed iSCSI connections was best accomplished using Windows PowerShell® commands to establish these connections and their preferred active and standby paths. Windows PowerShell in the Windows Server 2012 R2 VMs worked well in making consistent connections to the proper CVM targets on each Exchange Server Windows VM. This is detailed in the configuration section **Error! Reference source not found.**

The minimum number of Dell EMC XC Series nodes in a cluster is three. When clustered together, they create a single storage pool and a single storage container within it. The storage container is presented to all nodes within the cluster and is spread across all of the node's disks.

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. In such a case, the cluster will begin to move the data on that virtual machine to the local cluster node using the same 10Gb network.

2.3 VM layout

The baseline for this configuration 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 and Microsoft Exchange 2016 with one virtual machine per node.

Each virtual machine consisted of two virtual processors, 48GB RAM, and thin-provisioned disks in the layout described in Table 1.

Table 1 Microsoft Exchange disk layout

100GB	Drive C:	Windows OS disk (NFS)
500GB	Mount point: C:\DB\DB1	Database disk (iSCSI)
500GB	Mount point: C:\DB\DB2	Database disk (iSCSI)

3 Configuration steps

3.1 XC Series installation and configuration

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

3.1.1 VMware and Nutanix pre-installation and configuration

The following steps were completed:

- VMware vSphere ESXi 6.0 pre-installed on the SATADOMs of three Dell EMC XC430-4 Series nodes
- Nutanix Acropolis 4.5.1 pre-installed on XC430-4 nodes (pre-installation automatically configures CVMs on each hypervisor host)
- Windows VMs stored in traditional VMFS datastore hosted on Nutanix DSF
- 1TB infrastructure container created in Acropolis DSF storage to house VMware datastore infrastructure to save all EXCH Windows VMs

3.1.2 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 Nutanix ADSF)	Container: EXCHDATA	Container: Infrastructure
	Datastore: EXCHDATA Free space (logical): 12.28TiB Used: 1.76TiB Max capacity: 14.04TiB Replication factor: 2	VMware 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

3.1.3 iSCSI vgroup and vdisk creation

This section describes creating iSCSI vgroup or vdisk targets for each Exchange database through Nutanix DSF using the CVM console.

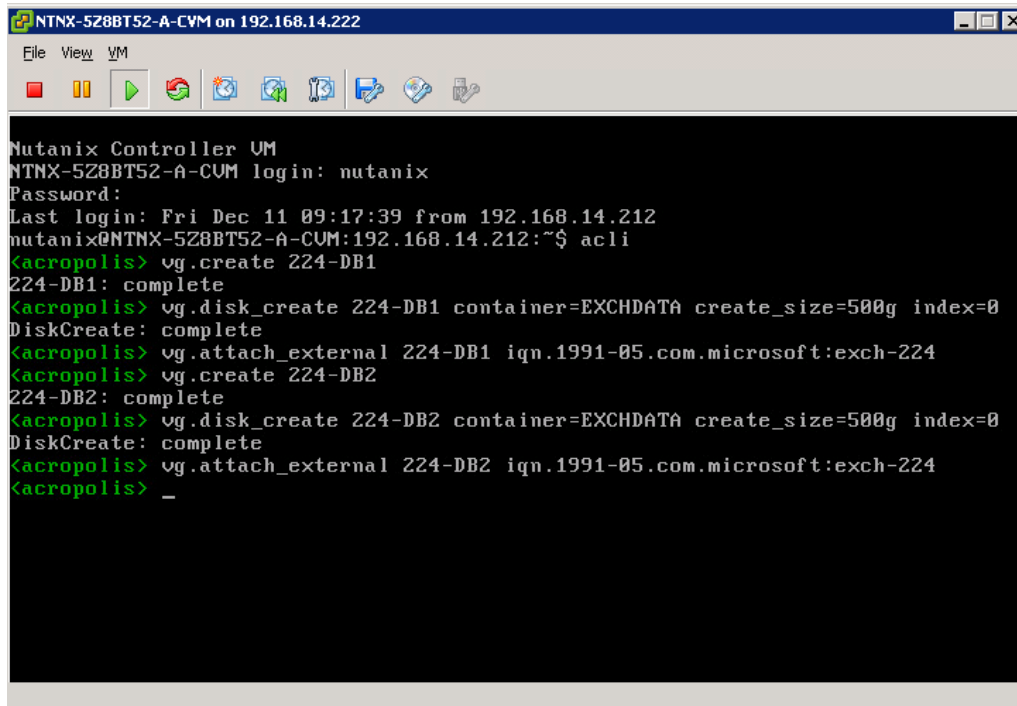
Issue the following commands by acli command or by going directly into the acli:

1. SSH into a CVM, or use the VMware console of CVM
2. After successful logon to a CVM console, at the prompt type acli and press <Enter>:

```
nutanix@NTNX-5Z8CT52-A-CVM:192.168.14.213:~$ acli
```

The acli prompt appears in green: `<acropolis>`

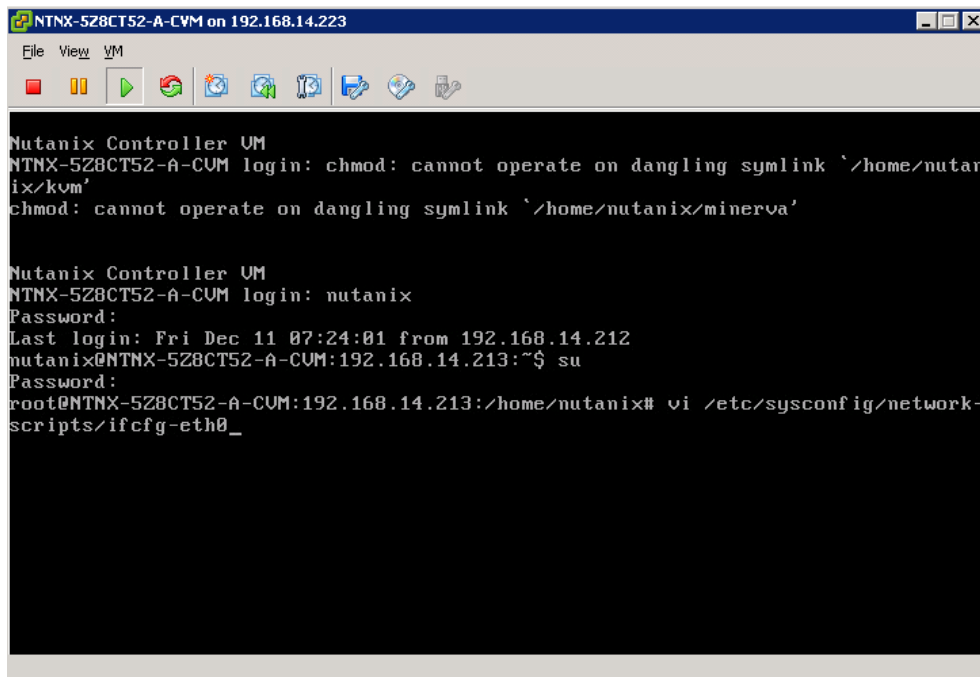
3. Use these three main acli commands to create volumes accessible by iSCSI:
 - `vg.create volume group name`
 - `vg.disk_create vg name create_size=xg container=your container index=lun id`
 - `vg.attach_external <disk name> <iqn string of Windows iSCSI initiator>`
4. The following example shows these commands used to create and attach two vgroups each with one vdisk of 500GB to store the two databases on the EXCH server. Creating EXCH-224 databases with acli (acropolis) shell using VMware console NTNX-5Z8BT52-A-CVM.



```
NTNX-5Z8BT52-A-CVM on 192.168.14.222
File View VM
Nutanix Controller VM
NTNX-5Z8BT52-A-CVM login: nutanix
Password:
Last login: Fri Dec 11 09:17:39 from 192.168.14.212
nutanix@NTNX-5Z8BT52-A-CVM:192.168.14.212:~$ acli
<acropolis> vg.create 224-DB1
224-DB1: complete
<acropolis> vg.disk_create 224-DB1 container=EXCHDATA create_size=500g index=0
DiskCreate: complete
<acropolis> vg.attach_external 224-DB1 iqn.1991-05.com.microsoft:exch-224
<acropolis> vg.create 224-DB2
224-DB2: complete
<acropolis> vg.disk_create 224-DB2 container=EXCHDATA create_size=500g index=0
DiskCreate: complete
<acropolis> vg.attach_external 224-DB2 iqn.1991-05.com.microsoft:exch-224
<acropolis> _
```

3.1.4 Setting Jumbo Frames for iSCSI traffic

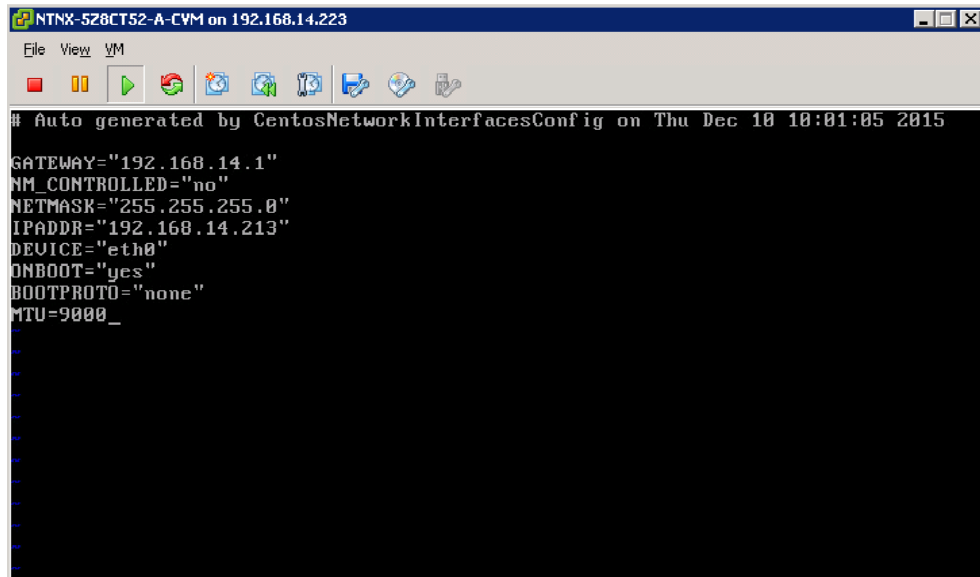
1. Connect to each CVM console using the VMware console or SSH and log on as **su** (super user).
2. Use the vi editor to edit `/etc/sysconfig/network-scripts/ifcfg-eth0`.



```
NTNX-5Z8CT52-A-CVM on 192.168.14.223
File View VM
Nutanix Controller UM
NTNX-5Z8CT52-A-CVM login: chmod: cannot operate on dangling symlink `/home/nutanix/kvm'
chmod: cannot operate on dangling symlink `/home/nutanix/minerva'

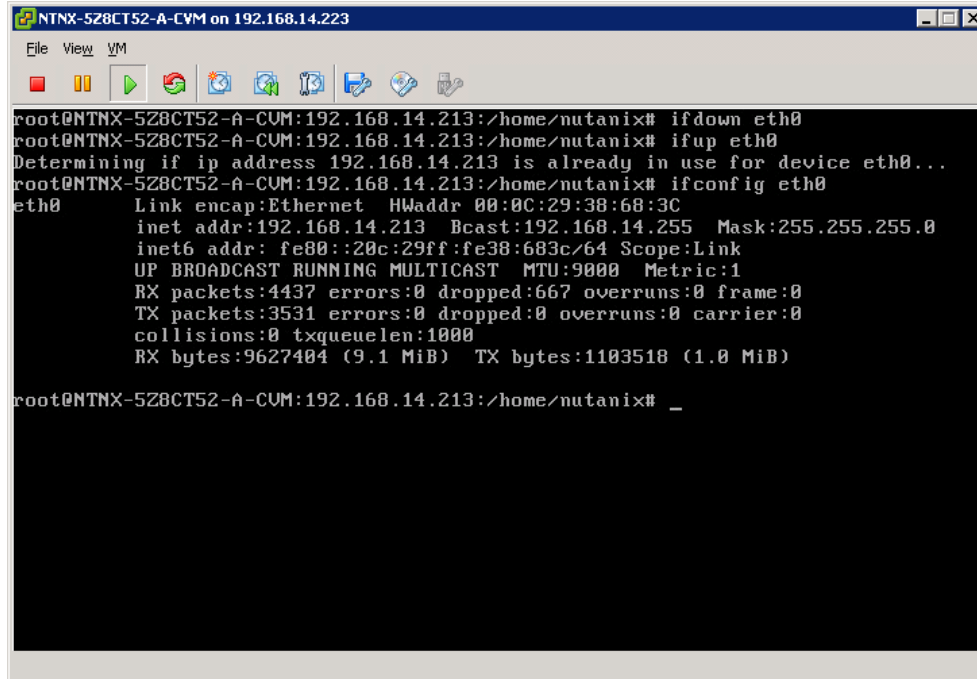
Nutanix Controller UM
NTNX-5Z8CT52-A-CVM login: nutanix
Password:
Last login: Fri Dec 11 07:24:01 from 192.168.14.212
nutanix@NTNX-5Z8CT52-A-CVM:192.168.14.213:~$ su
Password:
root@NTNX-5Z8CT52-A-CVM:192.168.14.213:/home/nutanix# vi /etc/sysconfig/network-scripts/ifcfg-eth0_
```

3. Add or change line `MTU=9000`, then save and quit the vi editor.



```
NTNX-5Z8CT52-A-CVM on 192.168.14.223
File View VM
# Auto generated by CentosNetworkInterfacesConfig on Thu Dec 10 10:01:05 2015
GATEWAY="192.168.14.1"
NM_CONTROLLED="no"
NETMASK="255.255.255.0"
IPADDR="192.168.14.213"
DEVICE="eth0"
ONBOOT="yes"
BOOTPROTO="none"
MTU=9000_
```

4. Reset the **eth0** interface using `ifdown eth0` and `ifup eth0`. Then use `ifconfig` to show the MTU:9000 setting as being set on the interface.

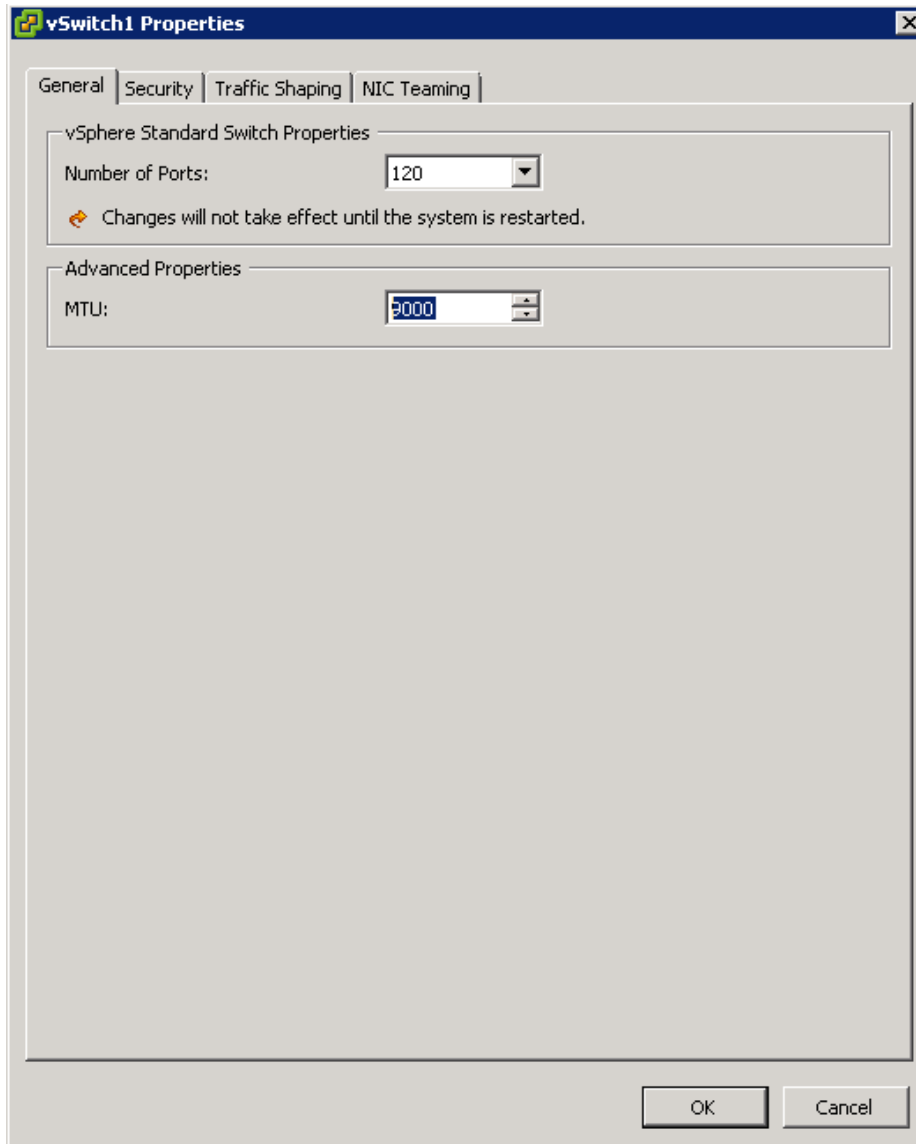


```
NTNX-5Z8CT52-A-CVM on 192.168.14.223
File View VM
root@NTNX-5Z8CT52-A-CVM:192.168.14.213:/home/nutanix# ifdown eth0
root@NTNX-5Z8CT52-A-CVM:192.168.14.213:/home/nutanix# ifup eth0
Determining if ip address 192.168.14.213 is already in use for device eth0...
root@NTNX-5Z8CT52-A-CVM:192.168.14.213:/home/nutanix# ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:0C:29:38:68:3C
          inet addr:192.168.14.213  Bcast:192.168.14.255  Mask:255.255.255.0
          inet6 addr: fe80::20c:29ff:fe38:683c/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:9000  Metric:1
          RX packets:4437 errors:0 dropped:667 overruns:0 frame:0
          TX packets:3531 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:9627404 (9.1 MiB)  TX bytes:1103518 (1.0 MiB)

root@NTNX-5Z8CT52-A-CVM:192.168.14.213:/home/nutanix# _
```

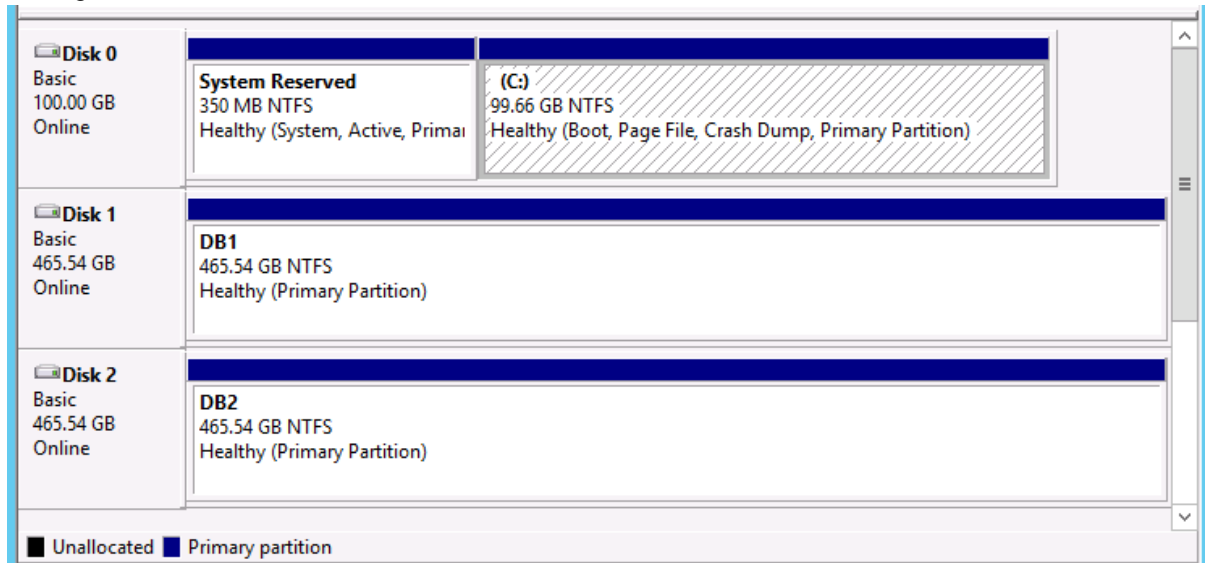
3.1.5 Setting Jumbo Frames for iSCSI traffic on VMware vSwitch

Set the **MTU** to **9000** on the vSwitch that will connect the iSCSI adapters using the VMware vSphere client.



3.2 Windows Server VM installation and configuration


1. Install Windows Server 2012 R2 from the standard VM template and perform Windows updates.
2. Start the in-guest iSCSI initiator GUI through the Windows control panel.
3. Connect redundant iSCSI connections from each Exchange server to each of the three controllers on each 10Gb NIC, creating six iSCSI connections per database volume.
4. Initialize, mount, and format database disk volumes (DB1, DB2) using 64K NTFS with Windows Disk Management.



3.2.1 Windows MPIO with iSCSI installation and configuration

1. Install Multipath I/O to Windows Server.
2. Use the MPIO Control Panel Microsoft DSM and configure to use iSCSI.
3. Use Failover Only (recommended) instead of Round Robin. Failover Only is recommended for the chosen workload in this scenario. See this link: [How to Configure Nutanix Volume Groups to Support Windows Failover Clustering.](#)

- Point the Exchange database disk MPIO Failover Only target to the CVM that is on the local host with the Exchange VM. For example, choose **NTNX-5Z8BT52-A-CVM** iSCSI targets for **EXCH-222** which has IP address of **192.168.14.212** on the 10Gb LAN.



- Use the `Connect-IscsiTarget` Windows PowerShell® command to connect the multiple iSCSI targets per DB volume for each 10Gb NIC initiator IP address to each target portal IP address:

```
# ** EXCH-222
# *** DB1
Connect-IscsiTarget -NodeAddress iqn.2010-06.com.nutanix:222-db1 -
InitiatorPortalAddress 192.168.14.225 -IsMultipathEnabled $true -
IsPersistent $true -TargetPortalAddress 192.168.14.212
Connect-IscsiTarget -NodeAddress iqn.2010-06.com.nutanix:222-db1 -
InitiatorPortalAddress 192.168.14.226 -IsMultipathEnabled $true -
IsPersistent $true -TargetPortalAddress 192.168.14.212
Connect-IscsiTarget -NodeAddress iqn.2010-06.com.nutanix:222-db1 -
InitiatorPortalAddress 192.168.14.225 -IsMultipathEnabled $true -
IsPersistent $true -TargetPortalAddress 192.168.14.213
Connect-IscsiTarget -NodeAddress iqn.2010-06.com.nutanix:222-db1 -
InitiatorPortalAddress 192.168.14.226 -IsMultipathEnabled $true -
IsPersistent $true -TargetPortalAddress 192.168.14.213
Connect-IscsiTarget -NodeAddress iqn.2010-06.com.nutanix:222-db1 -
InitiatorPortalAddress 192.168.14.225 -IsMultipathEnabled $true -
IsPersistent $true -TargetPortalAddress 192.168.14.214
Connect-IscsiTarget -NodeAddress iqn.2010-06.com.nutanix:222-db1 -
InitiatorPortalAddress 192.168.14.226 -IsMultipathEnabled $true -
IsPersistent $true -TargetPortalAddress 192.168.14.214
# *** DB2
```



```

Connect-IscsiTarget -NodeAddress iqn.2010-06.com.nutanix:222-db2 -
InitiatorPortalAddress 192.168.14.225 -IsMultipathEnabled $true -
IsPersistent $true -TargetPortalAddress 192.168.14.212
Connect-IscsiTarget -NodeAddress iqn.2010-06.com.nutanix:222-db2 -
InitiatorPortalAddress 192.168.14.226 -IsMultipathEnabled $true -
IsPersistent $true -TargetPortalAddress 192.168.14.212
Connect-IscsiTarget -NodeAddress iqn.2010-06.com.nutanix:222-db2 -
InitiatorPortalAddress 192.168.14.225 -IsMultipathEnabled $true -
IsPersistent $true -TargetPortalAddress 192.168.14.213
Connect-IscsiTarget -NodeAddress iqn.2010-06.com.nutanix:222-db2 -
InitiatorPortalAddress 192.168.14.226 -IsMultipathEnabled $true -
IsPersistent $true -TargetPortalAddress 192.168.14.213
Connect-IscsiTarget -NodeAddress iqn.2010-06.com.nutanix:222-db2 -
InitiatorPortalAddress 192.168.14.225 -IsMultipathEnabled $true -
IsPersistent $true -TargetPortalAddress 192.168.14.214
Connect-IscsiTarget -NodeAddress iqn.2010-06.com.nutanix:222-db2 -
InitiatorPortalAddress 192.168.14.226 -IsMultipathEnabled $true -
IsPersistent $true -TargetPortalAddress 192.168.14.214

```

6. Use these commands in PowerShell to set the global policy and preferred targets for each database disk volume on each EXCH server as in the following example:

```

PS C:\Users\Administrator> Set-MSDSMGlobalDefaultLoadBalancePolicy -Policy
FOO
PS C:\Users\Administrator> mpclaim -s -d 0
MPIO Disk2: 06 Paths, Fail Over Only, ALUA Not Supported
Controlling DSM: Microsoft DSM
SN: iqn.2010-06.com.nutanix:222-db1,L,0x981efd0f2dad197d
Supported Load Balance Policies: FOO RR RRWS LQD WP LB

```

Path ID	State	SCSI Address	Weight
0000000077030005	Standby	003 000 005 000	0
0000000077030004	Standby	003 000 004 000	0
0000000077030003	Standby	003 000 003 000	0
0000000077030002	Standby	003 000 002 000	0
0000000077030001	Standby	003 000 001 000	0
0000000077030000	Active/Optimized	003 000 000 000	0

PS C:\Users\Administrator> mpclaim -s -d 1

MPIO Disk3: 06 Paths, Fail Over Only, ALUA Not Supported
 Controlling DSM: Microsoft DSM
 SN: iqn.2010-06.com.nutanix:222-db2,L,0xe908f6e6098b2f50
 Supported Load Balance Policies: FOO RR RRWS LQD WP LB

Path ID	State	SCSI Address	Weight
000000007703000b	Standby	003 000 011 000	0
000000007703000a	Standby	003 000 010 000	0
0000000077030009	Standby	003 000 009 000	0
0000000077030008	Standby	003 000 008 000	0
0000000077030007	Standby	003 000 007 000	0
0000000077030006	Active/Optimized	003 000 006 000	0

7. To determine which Path ID is preferred, use the `iscsicli` PowerShell command in the following example. The Target Id # matches the last digit in the Path ID (for example, Target ID 6 matches 0000000077030006 in the previous `mpclaim` command). This path was chosen because it matched the IP of the iSCSI target portal for the local CVM and was the first virtual NIC.

```
PS C:\Users\Administrator> iscsicli sessionlist | select-string "Reported Mappings", "Target Portal"
```

```
Target Portal      : 192.168.14.212/3260
Reported Mappings  : Port 3, Bus 0, Target Id 0, LUN 0
Target Portal      : 192.168.14.212/3260
Reported Mappings  : Port 3, Bus 0, Target Id 1, LUN 0
Target Portal      : 192.168.14.213/3260
Reported Mappings  : Port 3, Bus 0, Target Id 2, LUN 0
Target Portal      : 192.168.14.213/3260
Reported Mappings  : Port 3, Bus 0, Target Id 3, LUN 0
Target Portal      : 192.168.14.214/3260
Reported Mappings  : Port 3, Bus 0, Target Id 4, LUN 0
Target Portal      : 192.168.14.214/3260
Reported Mappings  : Port 3, Bus 0, Target Id 5, LUN 0
Target Portal      : 192.168.14.212/3260
Reported Mappings  : Port 3, Bus 0, Target Id 6, LUN 0
Target Portal      : 192.168.14.212/3260
Reported Mappings  : Port 3, Bus 0, Target Id 7, LUN 0
Target Portal      : 192.168.14.213/3260
Reported Mappings  : Port 3, Bus 0, Target Id 8, LUN 0
Target Portal      : 192.168.14.213/3260
Reported Mappings  : Port 3, Bus 0, Target Id 9, LUN 0
Target Portal      : 192.168.14.214/3260
Reported Mappings  : Port 3, Bus 0, Target Id 10, LUN 0
Target Portal      : 192.168.14.214/3260
Reported Mappings  : Port 3, Bus 0, Target Id 11, LUN 0
```

8. Use the `mpclaim` command to choose the preferred Target Id based on which is local to this VM:

```
PS C:\Users\Administrator> mpclaim -l -d 0 1 0000000077030000 0 0 1
Successfully set the new load balance policy settings.
```

```
PS C:\Users\Administrator> mpclaim -l -d 1 1 0000000077030006 0 0 1
Successfully set the new load balance policy settings.
```

9. Use `mpclaim` to see that the preferred paths are marked with an asterisk (*),

```
PS C:\Users\Administrator> mpclaim -s -d 0
MPIO Disk2: 06 Paths, Fail Over Only, ALUA Not Supported
Controlling DSM: Microsoft DSM
SN: iqn.2010-06.com.nutanix:222-db1,L,0x981efd0f2dad197d
Supported Load Balance Policies: FOO RR RRWS LQD WP LB
```

Path ID	State	SCSI Address	Weight
0000000077030005	Standby	003 000 005 000	0
0000000077030004	Standby	003 000 004 000	0
0000000077030003	Standby	003 000 003 000	0
0000000077030002	Standby	003 000 002 000	0
0000000077030001	Standby	003 000 001 000	0
* 0000000077030000	Active/Optimized	003 000 000 000	0

PS C:\Users\Administrator> mpclaim -s -d 1

MPIO Disk3: 06 Paths, Fail Over Only, ALUA Not Supported
 Controlling DSM: Microsoft DSM
 SN: iqn.2010-06.com.nutanix:222-db2,L,0xe908f6e6098b2f50
 Supported Load Balance Policies: FOO RR RRWS LQD WP LB

Path ID	State	SCSI Address	Weight
000000007703000b	Standby	003 000 011 000	0
000000007703000a	Standby	003 000 010 000	0
0000000077030009	Standby	003 000 009 000	0
0000000077030008	Standby	003 000 008 000	0
0000000077030007	Standby	003 000 007 000	0
* 0000000077030006	Active/Optimized	003 000 006 000	0

3.2.2 Setting Jumbo Frames for iSCSI traffic on Windows Server NIC

On each virtual NIC that will be used for iSCSI traffic, set Jumbo Packet to **Jumbo 9000**.

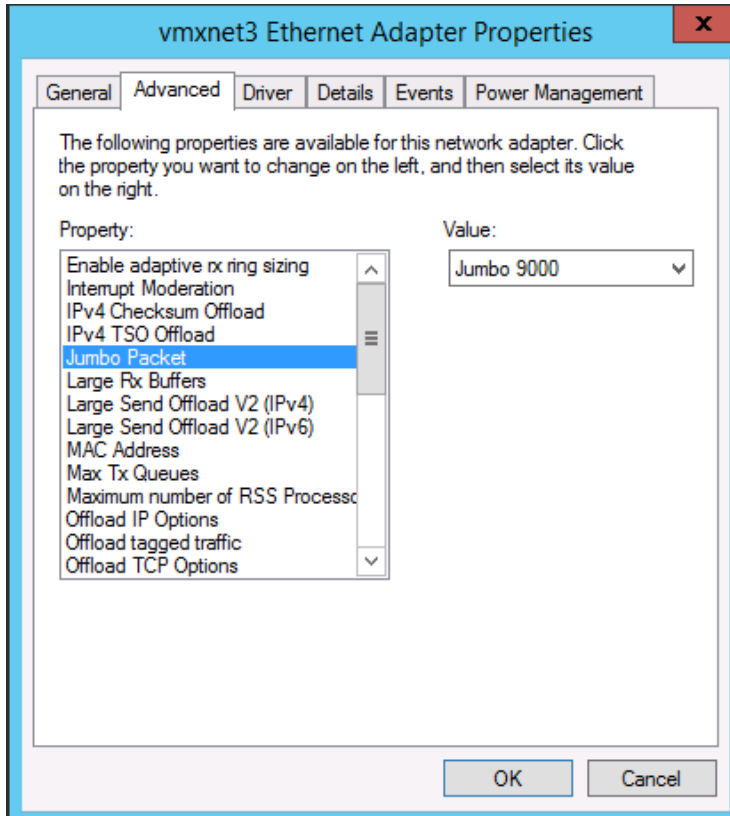


Figure 5 Jumbo Packet setting

3.3 Exchange 2016 simulated DAG servers installation and configuration

Microsoft supports Exchange Server 2016 storage workloads that are presented using an iSCSI initiator inside the guest virtual machine, as stated in the Microsoft article, [Microsoft Best Practices - Exchange 2016 virtualization](#).

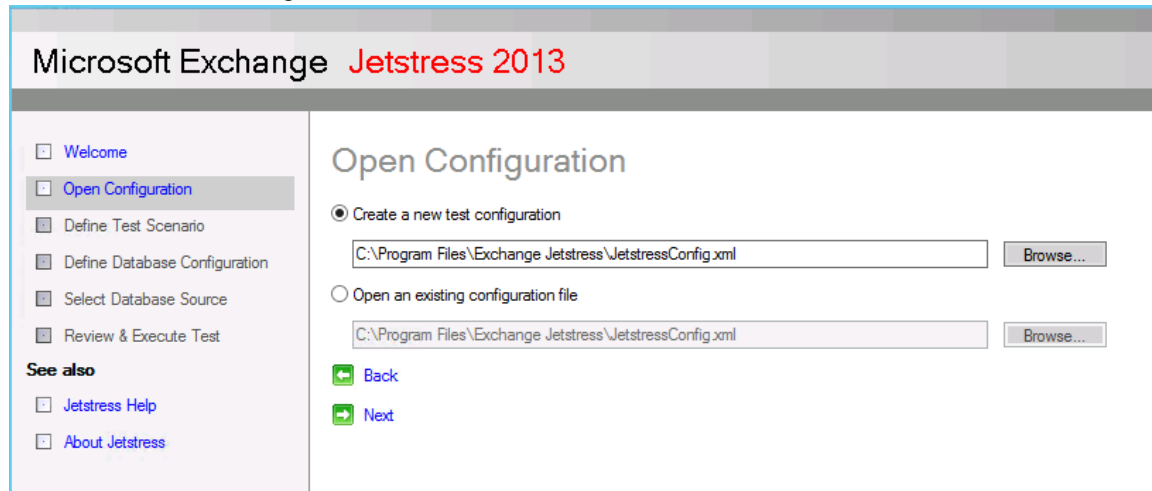
1. To calculate proper storage configuration for testing, use the [Microsoft Exchange Server role calculator worksheet](#) for Exchange configuration. The following test configuration was used to simulate real-world Exchange workload using Jetstress tool:
 - 3 servers
 - 3 database copies
 - 900 simulated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

The above configuration was not designed to push the limits of the system and achieve maximum performance. Please see section A.3 for links to other performance focused collaterals.

2. 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.
3. 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

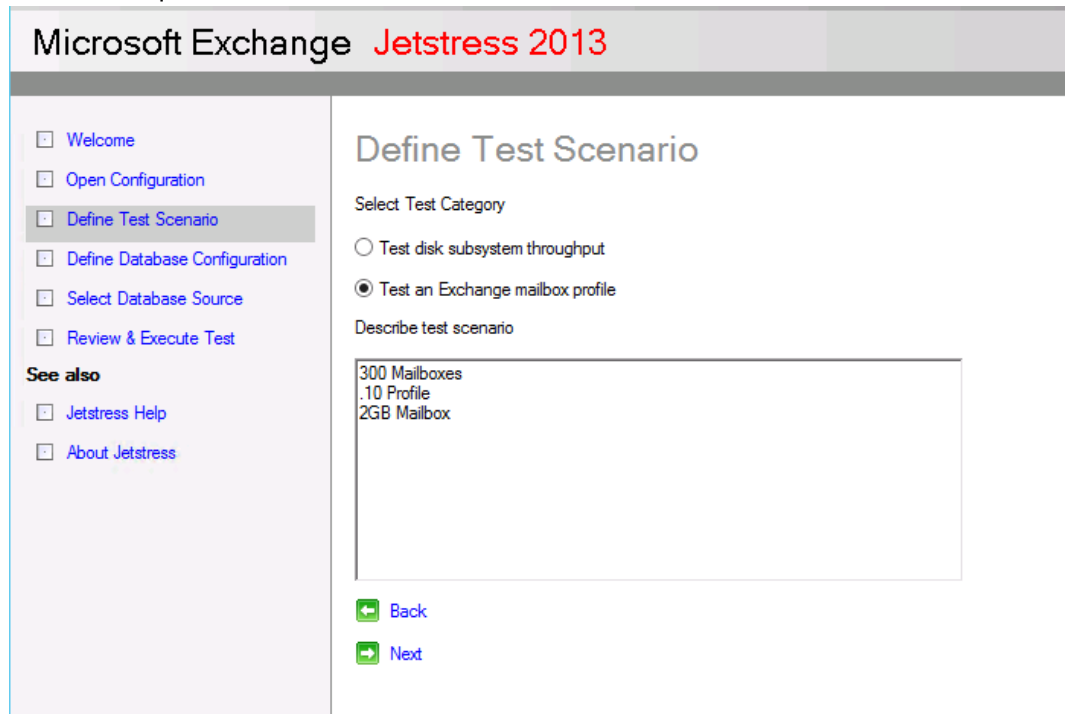
Next, configure Jetstress for the selected mailbox profile and quota and prepare the databases for testing.

1. Create a new test configuration.



The screenshot shows the 'Open Configuration' screen in the Microsoft Exchange Jetstress 2013 application. The left sidebar contains a navigation menu with 'Open Configuration' selected. The main area has two radio button options: 'Create a new test configuration' (selected) and 'Open an existing configuration file'. Both options have a text input field containing 'C:\Program Files\Exchange Jetstress\JetstressConfig.xml' and a 'Browse...' button. At the bottom, there are 'Back' and 'Next' navigation buttons.

2. Test an Exchange mailbox profile. To show an initially small workload for this model, a profile with 300 mailboxes and .10 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' screen in the Microsoft Exchange Jetstress 2013 application. The left sidebar has 'Define Test Scenario' selected. The main area has 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 area containing the text: '300 Mailboxes', '.10 Profile', and '2GB Mailbox'. At the bottom, there are 'Back' and 'Next' navigation buttons.

3. This small initial workload requires just one thread to generate the target I/O.

Microsoft Exchange **Jetstress 2013**

- Welcome
- Open Configuration
- Define Test Scenario**
- Define Database Configuration
- Select Database Source
- Review & Execute Test

See also

- Jetstress Help
- About Jetstress

Exchange Mailbox Profile

Number of mailboxes

IOPS/Mailbox

Mailbox size (MB)

Suppress tuning and use thread count (global)

[← Back](#)

[→ Next](#)

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

Microsoft Exchange **Jetstress 2013**

- Welcome
- Open Configuration
- Define Test Scenario**
- Define Database Configuration
- Select Database Source
- Review & Execute Test

See also

- Jetstress Help
- About Jetstress

Select Test Type

Performance

Database backup

Soft recovery

Multi-host test

Run background database maintenance

Continue the test run despite encountering disk errors

[← Back](#)

[→ Next](#)

- Define a test run of two hours.

Microsoft Exchange **Jetstress 2013**

- [Welcome](#)
- [Open Configuration](#)
- [Define Test Scenario](#)
- [Define Database Configuration](#)
- [Select Database Source](#)
- [Review & Execute Test](#)

See also

- [Jetstress Help](#)
- [About Jetstress](#)

Define Test Run

Output path for test results

Test duration (hours)

- Define the database configuration with two databases with mount points to the mounted iSCSI database volumes.

Microsoft Exchange **Jetstress 2013** Microsoft Exchange Server

- [Welcome](#)
- [Open Configuration](#)
- [Define Test Scenario](#)
- [Define Database Configuration](#)
- [Select Database Source](#)
- [Review & Execute Test](#)

See also

- [Jetstress Help](#)
- [About Jetstress](#)

Define Database Configuration

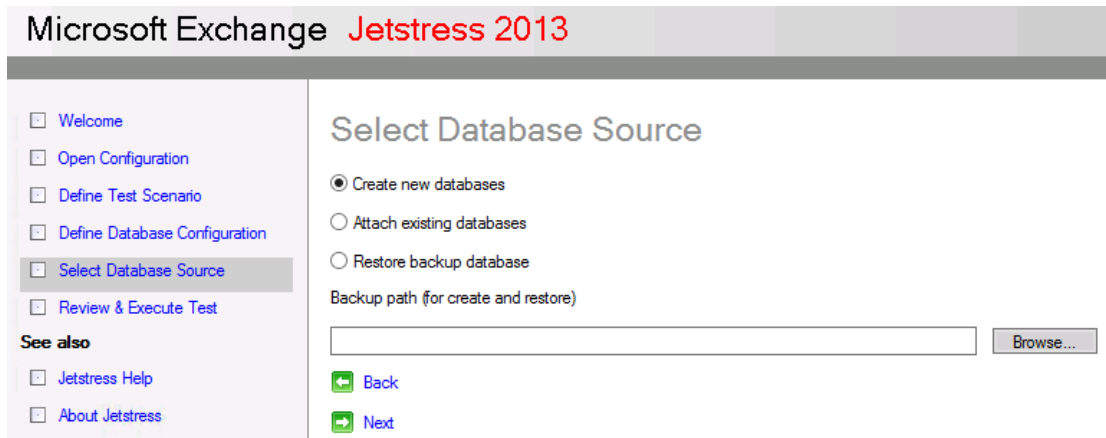
Number of databases

Number of copies per database

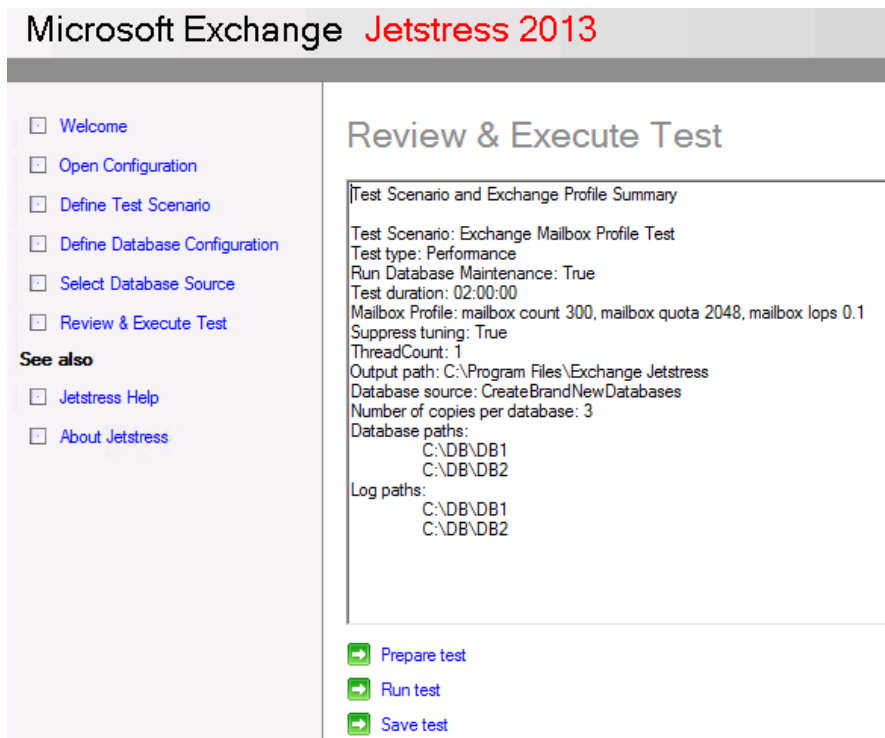
	Name	Database Path	
▶	Database1	C:\DB\DB1	<input type="button" value="..."/>
	Database2	C:\DB\DB2	<input type="button" value="..."/>

	Name	Log Path	
▶	Log1	C:\DB\DB1	<input type="button" value="..."/>
	Log2	C:\DB\DB2	<input type="button" value="..."/>

7. Create new databases.



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



4 Conclusion

Dell EMC XC Series appliances provide fast performance, easy scalability, simple management, and high availability for Microsoft Exchange environments. This paper demonstrated how in-guest iSCSI initiators 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, hyper-converged-application platform.

Fully supported by Nutanix: Nutanix fully supports the Nutanix platform in conjunction with in-guest iSCSI since version 4.3.

Fully supported by VMware: Microsoft Exchange that is deployed using in-guest iSCSI initiators on Windows VMs is fully supported by VMware.

Fully supported by Microsoft: In-guest iSCSI-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.

A Additional resources

A.1 Technical support and customer service

Offering online and telephone-based support and service options, Dell EMC support service can answer your questions about XC Series arrays, groups, volumes, array software, and host software. Availability varies by country and product, and some services might not be available in your area.

Visit Dell.com/support or call 800-945-3355 (United States and Canada).

Note: If you do not have access to an Internet connection, contact information is printed on your invoice, packing slip, bill, or Dell product catalog.

A.2 Dell EMC XC Series solutions

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

More Dell EMC Storage technical content can be found at [Storage Solutions Technical Documents](#).

A.3 Related documentation

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

- Nutanix best practices: [Virtualizing Microsoft Exchange on Nutanix: VMware vSphere Edition](#)
- Mike McGhee (Nutanix) blog post (archived): [How to Configure Nutanix Volume Groups to Support Windows Failover Clustering](#) for specifics on iSCSI setup with Nutanix
- Microsoft Best Practices: [Exchange 2016 virtualization](#)
- Virtualizing Exchange 2013: [Microsoft Exchange 2013 on VMware Best Practices Guide](#)