



Technical Report

# **VMware Configuration Guide for E-Series SANtricity iSCSI Integration with ESXi 6.x and 7.x Solution Design**

Darshan Hosad, Joe McCormick, Kelly Kemnitz, NetApp  
November 2020 | TR-4789

## **Abstract**

For full information about supported iSCSI host ports on a particular NetApp® E-Series system, see the [NetApp Hardware Universe](#).

## TABLE OF CONTENTS

<b>Overview of E-Series in VMware Environments .....</b>	<b>3</b>
What This Document Covers.....	3
<b>E-Series and VMware iSCSI Architecture.....</b>	<b>3</b>
Environments Using iSCSI Host Interfaces .....	3
iSCSI HA Architecture .....	3
Path Management .....	6
VMware ESXi 6.X and 7.Xwith Volume Groups and Dynamic Disk Pool Configuration.....	6
VMware Network and iSCSI Storage Adapter Configuration Details .....	7
Configuring One vSwitch Configuration with Two iSCSI Ports on Each Controller.....	7
iSCSI Initiator/Target Configuration on ESXi Hosts .....	12
Create and Configure Hosts and Clusters in SANtricity.....	15
Tuning VMware Settings to Maximize Performance .....	21
I/O Operation Limit—Performance Implications.....	22
Jumbo Frames.....	23
Performance Degradation with Data-Assurance-Enabled Volumes and iSCSI .....	29
VMware Port Binding.....	32
<b>Conclusion .....</b>	<b>32</b>
<b>Appendix A: Changing Jumbo Frame Settings from a VMware vSphere Web Client .....</b>	<b>33</b>
Change the MTU on a Virtual Switch from a VMware vSphere Web Client.....	33
Change the MTU on VMkernel Adapters.....	34
<b>Appendix B: Configuring iSCSI CHAP Authentication .....</b>	<b>35</b>
VMware vSphere Web Client View.....	35
Related Resources .....	37
<b>Where to Find Additional Information .....</b>	<b>37</b>
<b>Version History .....</b>	<b>37</b>

## LIST OF FIGURES

Figure 1) VMware HA architecture with E-Series storage systems—a single-vSwitch configuration with four iSCSI HIC ports per controller. ....	4
Figure 2) VMware HA architecture with E-Series storage systems—a single-vSwitch configuration with two iSCSI base ports per controller.....	5

# Overview of E-Series in VMware Environments

NetApp® E-Series storage systems integrate seamlessly with existing or new VMware environments. The flexible host interfaces and easy-to-integrate, understand, and manage storage configuration features make E-Series systems a natural choice for storage administrators and IT directors. Customers who need to balance the total cost of ownership with superior performance and features will enjoy the flexibility delivered by the range of E-Series products.

Using the NetApp SANtricity® System Manager software, storage administrators can quickly deploy E-Series systems in most configurations with little guidance or training. The intuitive E-Series SANtricity Storage Manager interface provides the tools needed to perform the following functions:

- Discover and name the storage system
- Manage software
- Complete systemwide implementation settings such as storage-system alerts and NetApp AutoSupport®
- Monitor and maintain the platform hardware over time

NetApp SANtricity System Manager can be used to create new VMware hosts, create and map volumes (LUNs), control E-Series copy service functions, and monitor the system for faults.

With ease of integration, system reliability, and service flexibility, NetApp E-Series storage systems offer cost-effective storage for customers who use VMware tool sets to manage the day-to-day complexities of their data centers.

## What This Document Covers

This technical report describes the steps needed to configure iSCSI integration with VMware. For VMware Express configuration, see [NetApp E-series and SANtricity 11 Documentation Center, SANtricity Software Express Configuration for VMware.](#)

This document does not cover VLANs, virtual machine (VM)/iSCSI pass through, or distributed vSwitches. For information about these topics, see [VMware Storage and Availability Technical Documents.](#)

## E-Series and VMware iSCSI Architecture

NetApp E-Series storage systems support up to four 25Gb optical iSCSI ports on each controller that interface with servers running the VMware vSphere ESXi OS. The VMware native multipathing (NMP) feature provides multipath management without adding the complexity associated with other OS-based multipath drivers used in bare-metal server implementations. The path policy defaults to round robin and can be tuned to force alternate path selections on a smaller number of I/O requests.

## Environments Using iSCSI Host Interfaces

VMware environments often use the iSCSI protocol to connect ESXi hosts to a multivendor storage platform in the data center. Unfortunately, the vast tuning and configuration options available with iSCSI implementations can make this protocol choice very complicated. Careful planning is required to properly lay out the iSCSI network for a given implementation so that all target-to-initiator paths are strictly layer 2. Layer 3 routing of I/O between ESXi host initiators and E-Series storage targets is not supported.

## iSCSI HA Architecture

E-Series storage systems offer full redundancy when the paths from a single ESXi host are spread across the A-side and B-side controllers on each storage system. This configuration is indicated by the blue

(controller A) and red (controller B) paths in Figure 1 and Figure 2. The only difference between the two configurations shown is the number of iSCSI ports on the controller.

Figure 1 has four iSCSI HIC ports per controller and thus has four VMkernel ports on each ESXi host. Figure 2 has two iSCSI ports per controller and thus has two VMkernel port on each ESXi host.

For both architectural configurations, all VMkernel ports reside in the same vSwitch and can share the physical NICs for basic link redundancy within the vSwitch. Under link-fault conditions using the default VMware ESXi settings, the configurations have the same failover behaviors. The use of one configuration rather than the other should be based on the number of paths between host and storage array.

For more information about supported ports and speed on E-Series hardware, see [Introduction to NetApp E-Series E2800 Arrays](#) and [Introduction to NetApp E-Series E5700 Arrays](#).

Figure 1) VMware HA architecture with E-Series storage systems—a single-vSwitch configuration with four iSCSI HIC ports per controller.

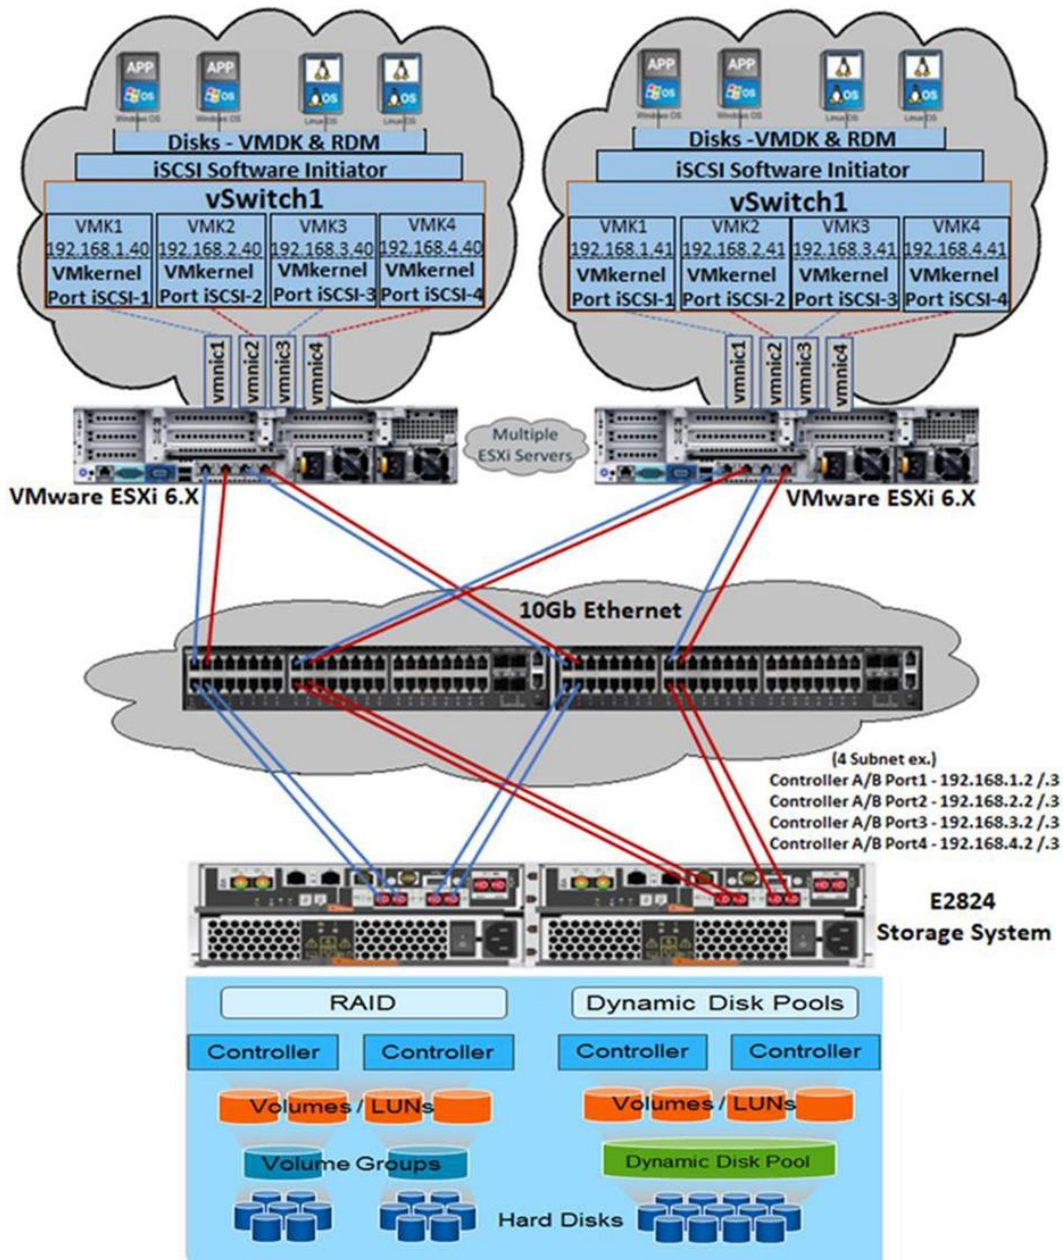
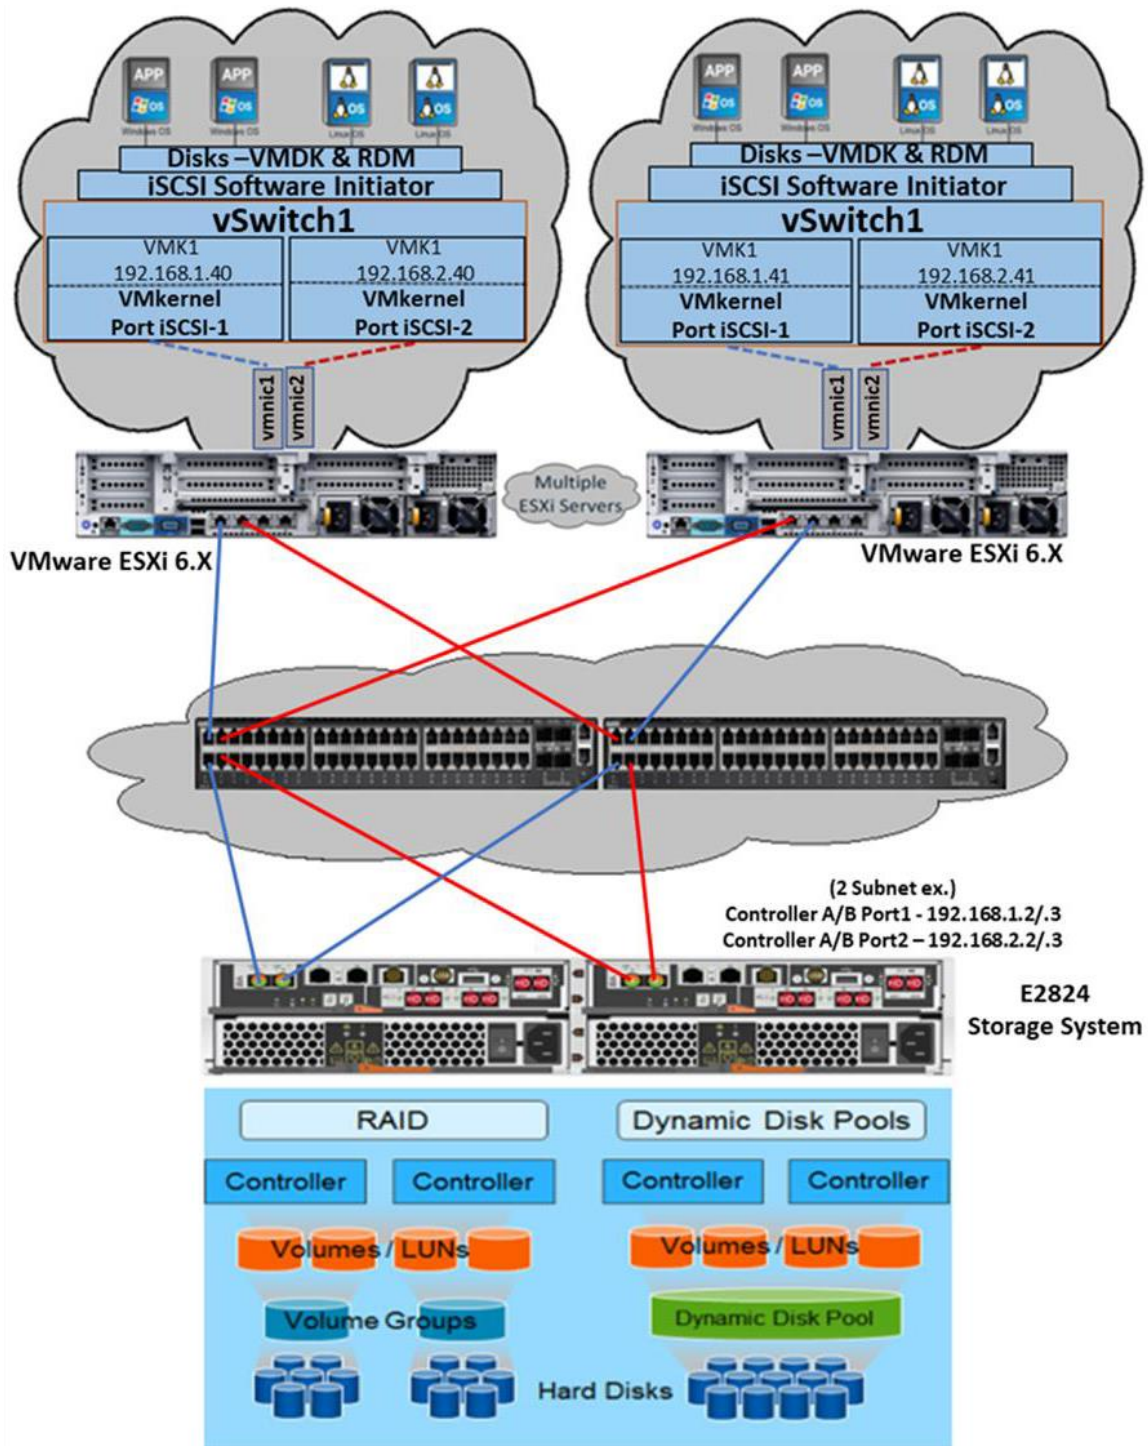


Figure 2) VMware HA architecture with E-Series storage systems—a single-vSwitch configuration with two iSCSI base ports per controller.



**Note:** The VMware ESXi 6.x and 7.x documentation states that up to eight paths from an ESXi host to a single LUN are supported. As a result, each controller host port must be in a different IP subnet. Failure to put the port pairs (that is, Controller A Port 1 and Controller B Port 1, Controller A Port 2 and Controller B Port 2, and so on) in individual subnets can result in the host discovering more than eight paths to each LUN or potentially not discovering all of the intended eight paths to each LUN.



The Figure 1 configuration uses a single vSwitch and four iSCSI HIC ports per controller. Each ESXi host can establish eight physical paths to each storage LUN, four active-optimized paths through the controller with LUN ownership, and four active, nonoptimized paths through the alternate controller in the storage system.

The Figure 2 configuration uses a single vSwitch and two iSCSI base ports per controller. Each ESXi can establish four physical paths to each storage LUN, two active-optimized paths through the controller with LUN ownership, and two active nonoptimized paths through the alternate controller in the storage system.

### Best Practice

Place each controller host port pair in a different IP subnet or VLAN.

## Path Management

By default, ESXi 6.x and 7.x contains storage claim rules associated with the paths from VMware devices to NetApp E-Series storage systems. A path policy defined in the ESXi claim rules specifies round robin for all NetApp E-Series devices. Specifically, the path failover for the one-vSwitch architecture is handled in the vSwitch.

The physical storage LUNs from the E-Series storage system are assigned to the ESXi host by using the E-Series SANtricity System Manager. Each HIC port is configured with an IP address in a local subnet to a specific NIC port on the ESXi host, as shown in Figure 1. This method divides traffic by using the subnets. However, both controllers should have access to all subnets so that the VMware multipath policy on each host manages all available paths to the storage system correctly.

### Best Practice

All ESXi hosts that are connected to a single storage system should use the same vSwitch and multipath settings to avoid inconsistent load balancing behaviors on the storage system host interface ports.

## VMware ESXi 6.x and 7.x with Volume Groups and Dynamic Disk Pool Configuration

Options for using E-Series volume groups or NetApp Dynamic Disk Pools (DDPs) for the storage configuration supporting VMware are shown in Figure 1 and Figure 2. VMware ESXi 6.x and 7.x software writes a variable segment size of up to 128KB. Therefore, standard RAID-based volume groups that are tuned to match specific segment sizes or DDP volumes that have a default nontunable 128KB segment size are well suited for VMware workloads. As a result, either E-Series storage configuration can be used to meet the requirements for individual storage implementations. In VMware, E-Series volumes are commonly used as VMFS datastores, but they can also be used for raw device mappings (RDMs).

All the possible storage and LUN mapping options can deliver low-latency I/O at various levels of IOPS and throughput for random I/O. However, volume group configurations that use the VMware RDM option are best suited for large sequential I/O.

### Best Practice

For random workloads, DDPs match the performance of and in some cases outperform comparable RAID 6 volume group configurations. As a result, when reliability, availability, and serviceability are the overriding considerations and VMware disks greater than 2TB are required, NetApp recommends E-Series DDPs with the VMware RDM feature. For LUNs smaller than 2TB, NetApp recommends E-Series DDPs with VMware virtual disks.

## Host Block Size Requirements

For EF300 and EF600 E-Series systems, a volume can be set to support a 512-byte block size (also called sector size) through the SANtricity System Manager. You must set the correct value during the volume creation process. If possible, the System Manager interface suggest the appropriate default value.

- Before setting the volume block size, read the following limitations and guidelines.
- At this time, VMware require a 512-byte block size and does not support 4KiB.
- The type of drives you select for your pool or volume group also determines what volume block sizes are supported, as follows:
  - If you create a volume group using drives that write to 512-byte blocks, then you can only create volumes with 512-byte blocks.
  - If you create a volume group using drives that write to 4KiB blocks, then you can create volumes with either 512-byte or 4KiB blocks.
- If the array has an iSCSI host interface card, all volumes are limited to 512-byte blocks (regardless of volume group block size). This is due to a specific hardware implementation.
- You cannot change block size once it is set. If you need to change a block size, you must delete the volume and re-create it.

For more information about how to set the host block size for volumes within the SANtricity System Manager, see the [SANtricity System Manager Online Help](#).

## VMware Network and iSCSI Storage Adapter Configuration Details

VMware allows multiple configurations of virtual networks to accommodate redundancy and throughput requirements. In many cases, an ESXi server must drive workflows by using multiple 10Gb or 25Gb links to an E-Series storage system. In that case, care must be taken so that traffic uses all available paths in a balanced manner. Various configurations have been tested so that performance and link-fault characteristics are well documented.

This VMware configuration guide uses a virtual switch configuration in which all VMK ports are associated with a single storage system. In this configuration, each VMK is assigned a unique IP address and subnet that is then associated with an assigned primary vmnic.

Based on the physical network architecture and IP scheme, each VMK port is configured to access two paths for each LUN on the E-Series storage system. One path is through controller A and one path is through controller B. By using the architecture in .

For more information on supported ports and speed on E-Series hardware, see Introduction to NetApp E-Series E2800 Arrays and Introduction to NetApp E-Series E5700 Arrays.

The configurations in Figure 1 support a maximum of eight paths to any LUN on the storage system. By using the Figure 2 architecture, the configuration supports four paths to any LUN on the storage system. The following section describes the configuration of E-Series and VMware connectivity over iSCSI using the second architecture (Figure 2).

Please use the following [video](#) as a visual guide for configuring a vSwitch on an ESXi host. There are also instructions regarding setting up an iSCSI initiator and target as well as instructions for configuring the host on SANtricity.

## Configuring One vSwitch Configuration with Two iSCSI Ports on Each Controller

To configure one vSwitch on an ESXi host, complete the following steps:

1. Create the vSwitch and add uplinks:
  - a. On the ESXi Host on the Navigator tab, select Networking>Virtual Switches.

- b. Click Add Standard Virtual Switch and choose the specific vmnic on the Uplink1 option. Click Add.

**Add standard virtual switch - vSwitch1**

**Add uplink**

vSwitch Name	vSwitch1
MTU	1500
Uplink 1	vmnic5
▶ Link discovery	Click to expand
▶ Security	Click to expand

**Add** **Cancel**

- c. You can add only one uplink at a time. To add more uplinks, select the virtual switch you already created, click Add Uplink, and select the specific vmnic on the Uplink 2 options. Then click Save. Each vmnic should be connected to a different physical switch to eliminate a single point of failure on the physical switch.

**Edit standard virtual switch**

**Add uplink**

vSwitch Name	vSwitch1
MTU	1500
Uplink 1	vmnic5
Uplink 2	vmnic1
▶ Link discovery	Click to expand
▶ Security	Click to expand
▶ NIC teaming	Click to expand
▶ Traffic shaping	Click to expand

**Save** **Cancel**

- d. Verify that a switch with two uplinks has been created.



localhost.ict.englab.netapp.com - Networking

Port groups **Virtual switches** Physical NICs VMkernel NICs TCP/IP stacks Firewall rules

Add standard virtual switch Add uplink Edit settings Refresh Actions Search

Name	Port groups	Uplinks	Type
vSwitch0	2	1	Standard vSwitch
vSwitch1	2	2	Standard vSwitch

2 items

2. Add a VMkernel NIC and assign the IP address:
  - a. Go to Networking > VMkernel NICs and click Add VMkernel NIC.
  - b. From the Virtual Switch drop-down menu, select the virtual switch that you created in step 1.
  - c. In the New Port Group field, enter the port group name (for example, iSCSI-1).
  - d. In the IPv4 settings, select Static.
  - e. From the drop-down menu, assign the IP address for the VMkernel NIC.
  - f. Click Create.

Add VMkernel NIC

Port group: New port group

New port group: iSCSI-1

Virtual switch: vSwitch1

VLAN ID: 0

MTU: 1500

IP version: IPv4 only

IPv4 settings

Configuration:  DHCP  Static

Address: 192.168.1.40

Subnet mask: 255.255.255.0

TCP/IP stack: Default TCP/IP stack

Services:  vMotion  Provisioning  Fault tolerance logging

Create Cancel

- g. Repeat step 2 to create additional VMkernel NICs. Click Create.
- h. Verify that the VMkernel NICs have been created with IPv4 address.

localhost.ict.englab.netapp.com - Networking

Port groups   Virtual switches   Physical NICs   **VMkernel NICs**   TCP/IP stacks   Firewall rules

Add VMkernel NIC   Edit settings   Refresh   Actions   Search

Name	Portgroup	TCP/IP stack	Services	IPv4 address	IPv6 addresses
vmk0	Management Network	Default TCP/IP stack	Management	10.113.84.179	None
vmk1	iSCSI-1	Default TCP/IP stack		192.168.1.40	None
vmk2	iSCSI-2	Default TCP/IP stack		192.168.2.40	None

3 items

### 3. Configure port groups:

- a. Go to Networking > Port Groups, select iSCSI-1, and click Edit Settings.

localhost.ict.englab.netapp.com - Networking

**Port groups**   Virtual switches   Physical NICs   VMkernel NICs   TCP/IP stacks   Firewall rules

Add port group   **Edit settings**   Refresh   Actions   Search

Name	Active p...	VLAN ID	Type	vSwitch	VMs
VM Network	1	0	Standard port group	vSwitch0	1
Management Network	1	0	Standard port group	vSwitch0	N/A
iSCSI-2	1	0	Standard port group	vSwitch1	N/A
<b>iSCSI-1</b>	<b>1</b>	<b>0</b>	<b>Standard port group</b>	<b>vSwitch1</b>	<b>N/A</b>

4 items

- b. Click NIC Teaming and then click Yes for Override Failover Order.
- c. Select one active vmnic, with the rest set to Standby for each port group. Click Save.

The choice of which vmnic to activate depends on which subnet the VMkernel NIC equivalent to the port group is on. For example, iSCSI-1 is on subnet 192.168.1.X. vmnic5 is connected to the iSCSI port on the storage on subnet 192.168.1.X, and vmnic1 is connected to an iSCSI port on the storage on subnet 192.168.2.X. Therefore, vmnic5 should be set to Active and vmnic1 should be set to Standby for that port group.


- d. Click Save.










**Edit port group - iSCSI-1**

Name	iSCSI-1									
VLAN ID	0									
Virtual switch	vSwitch1									
Security	Click to expand									
<b>NIC teaming</b>										
Load balancing	Inherit from vSwitch									
Network failover detection	Inherit from vSwitch									
Notify switches	<input type="radio"/> Yes <input type="radio"/> No <input checked="" type="radio"/> Inherit from vSwitch									
Failback	<input type="radio"/> Yes <input type="radio"/> No <input checked="" type="radio"/> Inherit from vSwitch									
Override failover order	<input checked="" type="radio"/> Yes <input type="radio"/> No									
Failover order	<div> <input type="checkbox"/> Mark standby    <input type="checkbox"/> Move up    <input type="checkbox"/> Move down </div> <table border="1"> <thead> <tr> <th>Name</th> <th>Speed</th> <th>Status</th> </tr> </thead> <tbody> <tr> <td>vmnic5</td> <td>10000 Mbps, full duplex</td> <td>Active</td> </tr> <tr> <td>vmnic1</td> <td>10000 Mbps, full duplex</td> <td>Standby</td> </tr> </tbody> </table>	Name	Speed	Status	vmnic5	10000 Mbps, full duplex	Active	vmnic1	10000 Mbps, full duplex	Standby
Name	Speed	Status								
vmnic5	10000 Mbps, full duplex	Active								
vmnic1	10000 Mbps, full duplex	Standby								
Traffic shaping	Click to expand									

e. Override the failover order of vmnics on the iSCSI-2 port group.

 Edit port group - iSCSI-2

Name	ISCSI-2									
VLAN ID	0									
Virtual switch	vSwitch1									
▶ Security	Click to expand									
▼ NIC teaming										
Load balancing	Inherit from vSwitch									
Network failover detection	Inherit from vSwitch									
Notify switches	<input type="radio"/> Yes <input type="radio"/> No <input checked="" type="radio"/> Inherit from vSwitch									
Failback	<input type="radio"/> Yes <input type="radio"/> No <input checked="" type="radio"/> Inherit from vSwitch									
Override failover order	<input checked="" type="radio"/> Yes <input type="radio"/> No									
Failover order	<div style="display: flex; align-items: center; gap: 5px;"> <span> Mark standby</span> <span> Move up</span> <span> Move down</span> </div> <table border="1"> <thead> <tr> <th>Name</th> <th>Speed</th> <th>Status</th> </tr> </thead> <tbody> <tr> <td> vmnic 1</td> <td>10000 Mbps, full duplex</td> <td>Active</td> </tr> <tr> <td> vmnic5</td> <td>10000 Mbps, full duplex</td> <td>Standby</td> </tr> </tbody> </table>	Name	Speed	Status	 vmnic 1	10000 Mbps, full duplex	Active	 vmnic5	10000 Mbps, full duplex	Standby
Name	Speed	Status								
 vmnic 1	10000 Mbps, full duplex	Active								
 vmnic5	10000 Mbps, full duplex	Standby								
▶ Traffic shaping	Click to expand									

4. Configure the vSwitch on the other ESXi host.

## iSCSI Initiator/Target Configuration on ESXi Hosts

To configure the iSCSI initiator/target on ESXi hosts, complete the following step:

1. In SANtricity, go to Settings > System > iSCSI settings > Configure iSCSI Ports.

To configure all iSCSI ports with IPv4 addresses on Controller A and Controller B, complete the following steps:

1. Click Controller A and then click Next.

Configure iSCSI Ports

I want to configure iSCSI ports for...

Controller A

Controller B

Next > Cancel

2. From the drop-down menu, select the port on Controller A and then click Next.

Configure iSCSI Ports

1 Select Port 2 Configure Port 3 Configure Network Settings

I want to configure network settings for the following Controller A port...

Port 0a Link Status: Up

What else do I need to do to configure or diagnose iSCSI?

Cancel Next >

3. In the iSCSI Ports window, enable IPv4 and enable ICMP ping responses. Click Next.

Configure iSCSI Ports

1 Select Port 2 Configure Port 3 Configure Network Settings

I want to configure network settings for the following Controller A port...

Port 0a settings Show more port settings

MAC address: 00:A0:98:BF:86:E9

Enable IPv4:

Enable IPv6:

Enable ICMP PING responses (applies to all iSCSI ports on the storage array)

< Back Cancel Next >

4. Select the Manually Specify Static Configuration option and enter the IP address for the iSCSI port. Click Finish.



Configure iSCSI Ports
✕

1 Select Port
2 Configure Port
3 Configure Network Settings

I want to configure IPv4 for my port...

**Port 0a IPv4 network settings** Show more IPv4 settings

Automatically obtain configuration from DHCP server

**Manually specify static configuration:**

IP address

Subnet mask

Gateway

< Back
Cancel
Finish

5. Go to Settings > System > iSCSI settings and copy the target IQN from SANtricity.

**iSCSI settings**

**Configure iSCSI Ports**  
Configure your iSCSI host connections on the storage array for I/O connectivity.

**Configure Authentication**  
Your iSCSI authentication method is currently set to **no authentication**.

**View/Edit Target Discovery Settings**  
Register your storage array's iSCSI information with an iSNS server and choose whether unnamed iSCSI discovery sessions are allowed.

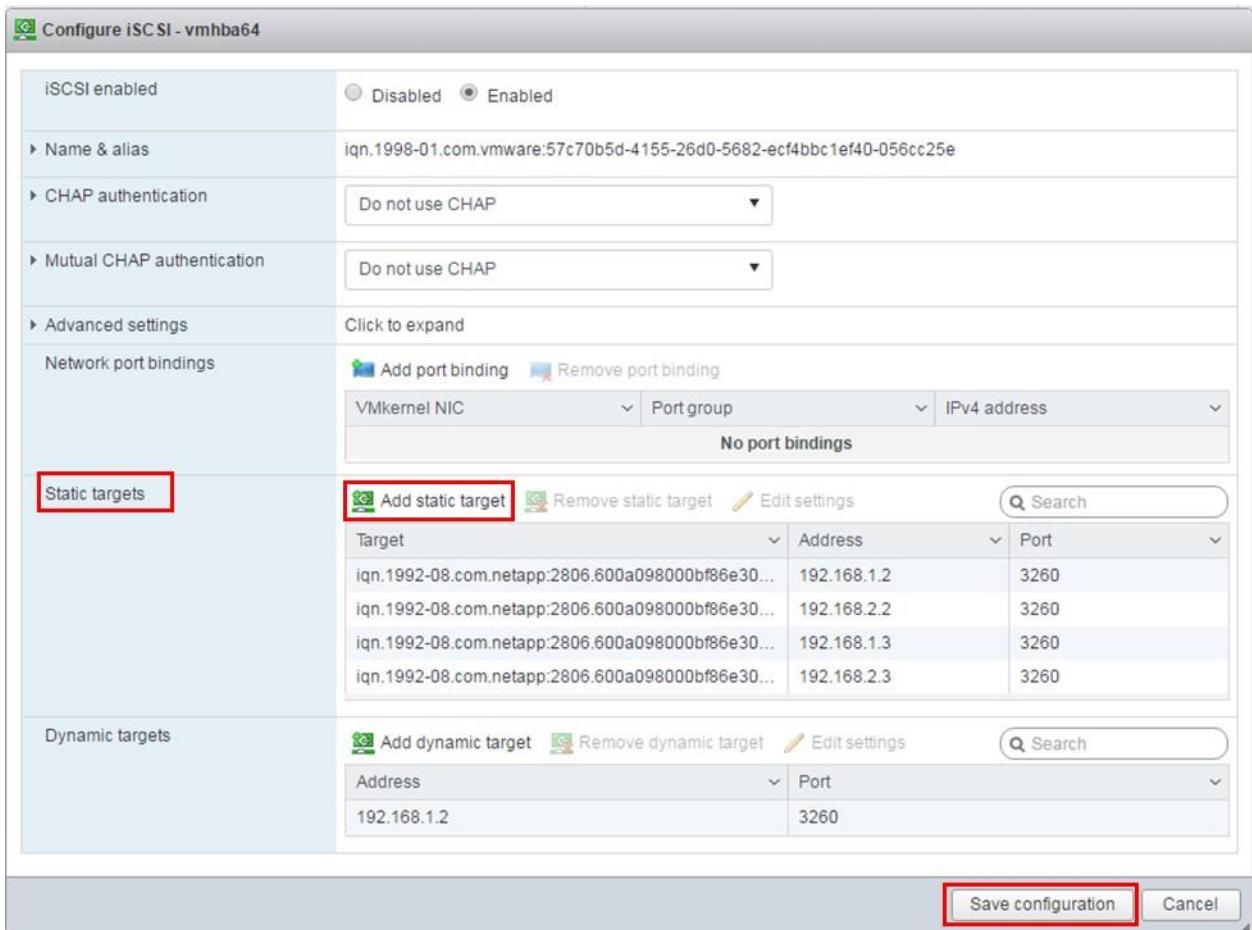
**View iSCSI Statistics**  
View the iSCSI statistics available on your storage array.

**View/End iSCSI Sessions**  
View and/or end iSCSI sessions to force initiators off your storage array.

---

**Target IQN:** iqn.1992-08.com.netapp:2806.600a098000bf86e3000000005c5c3a78

6. On the ESXi host, go to Storage > Adapters > Configure iSCSI:
  - a. In the Static Targets menu, click Add Static Target.
  - b. For the Target option, paste the target IQN that you copied in step 2.
  - c. In the Address field, enter the IP address of the iSCSI port that you configured in step 1. Keep port 3260 as the default.
  - d. Enter all the static targets. The number of static targets is equal to the number of iSCSI ports configured on the storage array.
  - e. In the Dynamic Target menu, add one of the static IP addresses that you have already set. Keep port 3260 as the default. Click Save Configuration.



7. Follow steps 1 through 3 to configure iSCSI targets on other ESXi hosts.

## Create and Configure Hosts and Clusters in SANtricity

To create hosts in SANtricity after creating volumes from the DDP, complete the following steps:

1. Go to Storage > Hosts > Create > Host.
2. Select VMware as the Host Operating System Type.
3. Under Host Ports, specify the IQN of the ESXi host. Click Create.

Create Host
✕

[How do I match the host ports to a host?](#)  
[How do I know which host operating system type is correct?](#)

Name ?

Host operating system type

VMware
▼

Host ports ?

Set CHAP initiator secret ?

Create

Cancel

Repeat step 1 to create additional hosts in SANtricity. The number of hosts depends on the number of ESXi hosts in the environment. The architectures illustrated in .

For more information about supported ports and speed on E-Series hardware, see Introduction to NetApp E- Series E2800 Arrays and Introduction to NetApp E-Series E5700 Arrays.

Figure 1 and Figure 2 have two ESXi hosts.

4. (Optional) Create a cluster.

In a VMware environment, hosts typically need concurrent access to some volumes for HA purposes. For example, volumes used as datastores often need to be accessible by all hosts in the VMware cluster. Volumes that need to be accessed by more than one host must be mapped to a host cluster. Make sure that you have created at least two hosts before creating a cluster.

- a. To create a cluster, go to Storage > Hosts > Create > Host Cluster.
- b. Enter a name for the cluster and select the host to add to the cluster. Click Create.

### Create Host Cluster ✕

Why would I need to create a host cluster?

Name ?

Select hosts to share volume access ?

Create Cancel

3. Assign volumes to a host.

**Note:** This step applies only to volumes that are accessed by a single host, which are typically boot LUNs or standalone ESXi hosts. See step 5 to map volumes to a host cluster.

- a. Select the host and then click Assign Volumes. Select the volumes and then click Assign.

## Assign Volumes ✕

Filter ?

Select volumes to assign to Host **ESXi\_Host1**...

<input type="checkbox"/> Name	Capacity (GiB)	DA Enabled
<input type="checkbox"/> Access	N/A	N/A
<input checked="" type="checkbox"/> Datastore_1	2400.00	No
<input checked="" type="checkbox"/> Datastore_2	2400.00	No
<input checked="" type="checkbox"/> Datastore_3	2400.00	No
<input checked="" type="checkbox"/> Datastore_4	2400.00	No
<input type="checkbox"/> Datastore_5	2400.00	No
<input type="checkbox"/> Datastore_6	2400.00	No
<input type="checkbox"/> Datastore_7	2400.00	No
<input type="checkbox"/> Datastore_8	2400.00	No
<input type="checkbox"/> Datastore_9	1200.00	Yes

Selected rows: 4 of 11

Assign
Cancel

- b. Repeat step 3-a to assign volumes to other hosts.
- c. After assigning volumes, each host shows one additional volume, which is the access volume/LUN.

<span>Create ▾</span> <span>Assign Volumes</span> <span>Unassign Volumes</span> <span>View/Edit Settings</span> <span style="float: right;">Delete</span>						
Name	Type	Associated Objects	Total Assigned Volumes	Reported Capacity (GiB)	Host Type	Edit
Cluster	Cluster	2 Host(s)	0	0.00	VMware	
ESXi_Host1	Host Member	Cluster	5	9600.00	VMware	
ESXi_Host2	Host Member	Cluster	5	9600.00	VMware	

Total rows: 3

4. Unassign the access LUN. Access LUNs are used for in-band array management, which requires running SANtricity on the host. Because this is not possible with ESXi, you can unassign the access LUN using the following steps.
  - a. Select the host and click Unassign Volumes.
  - b. In the Unassign Volumes window, check the Access LUN and enter `unassign`. Click Unassign.



## Unassign Volumes



Volumes currently assigned to Host **ESXi\_Host1**

<input type="checkbox"/> Name	Capacity (GiB)
<input type="checkbox"/> Datastore_1	2400.00
<input type="checkbox"/> Datastore_2	2400.00
<input type="checkbox"/> Datastore_3	2400.00
<input type="checkbox"/> Datastore_4	2400.00
<input checked="" type="checkbox"/> Access	N/A

Selected rows: 1 of 5

If you proceed, you will lose any in-band management capabilities unless you have an access volume assigned to another host.

Type UNASSIGN to confirm that you want to perform this operation.

Unassign

Cancel

### Best Practice

Unassign the access LUN when using ESXi.

5. Assign volumes to a cluster.

**Note:** This step applies only to volumes that are to be shared between ESXi hosts.

- a. Select the cluster and then select the Assign Volumes option. Select the volumes and then click Assign.

## Assign Volumes



Filter



Select volumes to assign to Host Cluster **Cluster...**

<input type="checkbox"/>	Name	Capacity (GiB)	DA Enabled
<input type="checkbox"/>	Access	N/A	N/A
<input checked="" type="checkbox"/>	Datastore_9	1200.00	Yes
<input checked="" type="checkbox"/>	Datastore_10	1200.00	Yes

Selected rows: 2 of 3

Assign

Cancel

5. Verify that the volumes are mounted on the ESXi host.
  - a. Log into both ESXi hosts and verify that volumes are mounted. In the Navigator tab, go to Storage > Devices.
  - b. Click Rescan and Refresh.

localhost.ict.englab.netapp.com - Storage

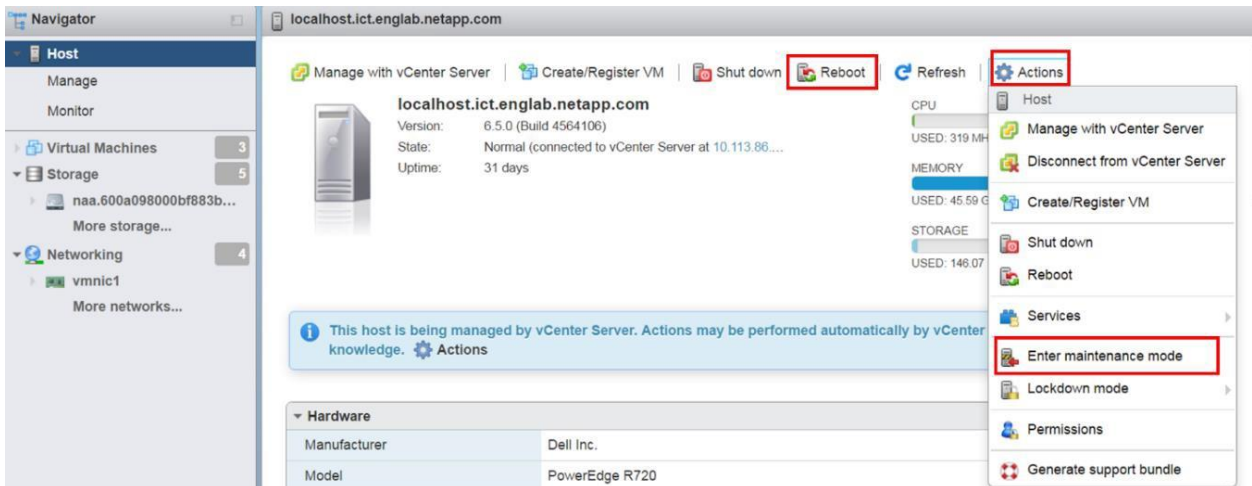
Datstores Adapters **Devices**

New datastore Increase capacity Rescan Refresh Actions

Name	Status	Type	Capacity	Queue Depth	Vendor
Local TSSSTcorp CD-ROM (mpx.vmhba0:C0:T4:L0)	Normal	CDROM	Unknown	N/A	TSSSTcorp
Local ATA Disk (t10.ATA_____WDC_WD5003ABYX2D18WERA0_____)	Normal	Disk	465.76 GB	31	ATA
NETAPP iSCSI Disk (naa.600a098000b883b000028855c5235c)	Normal	Disk	1.17 TB	128	NETAPP
NETAPP iSCSI Disk (naa.600a098000b883b000027905cb4382e)	Normal	Disk	2.34 TB	128	NETAPP
Local ATA Disk (t10.ATA_____WDC_WD5003ABYX2D18WERA0_____)	Normal	Disk	465.76 GB	31	ATA
NETAPP iSCSI Disk (naa.600a098000b883b00002c035cb4379f)	Normal	Disk	2.34 TB	128	NETAPP
NETAPP iSCSI Disk (naa.600a098000b883b00002c015cb43750)	Normal	Disk	1.17 TB	128	NETAPP
NETAPP iSCSI Disk (naa.600a098000b883b00002c015cb43750)	Normal	Disk	2.34 TB	128	NETAPP
NETAPP iSCSI Disk (naa.600a098000b883b0000278f5cb437de)	Normal	Disk	2.34 TB	128	NETAPP

9 items

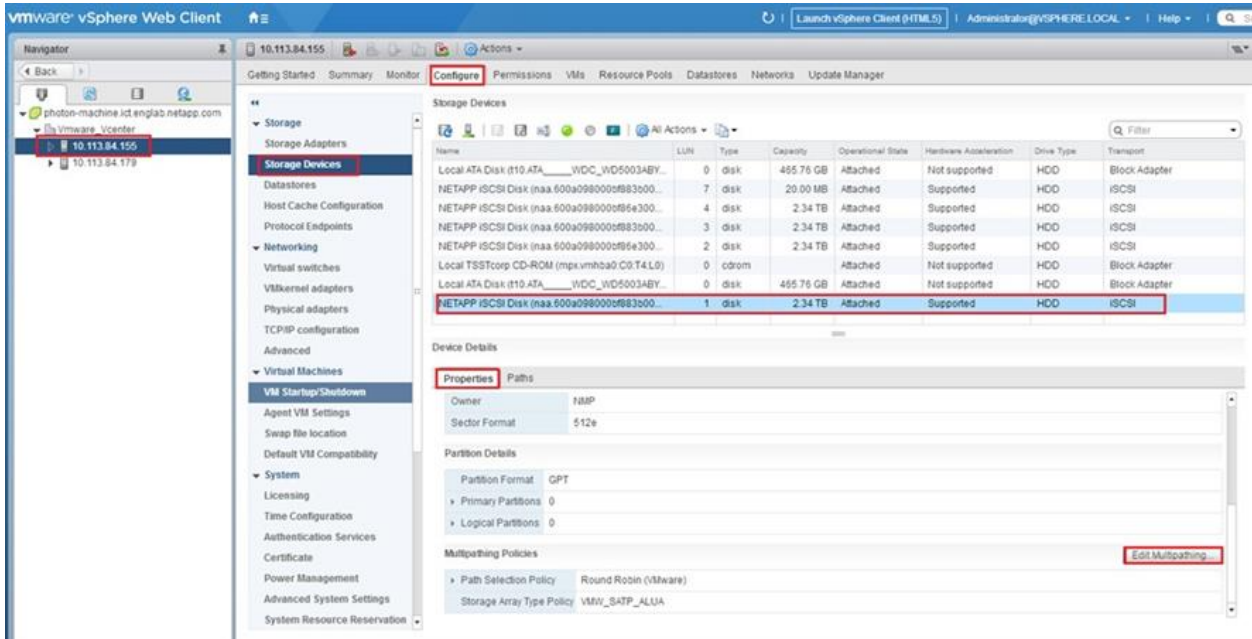
- c. If volumes did not show up after the rescan and refresh, you can try rebooting the host. From the Navigator tab, go to Host > Actions.
- d. From the drop-down menu, Click Enter Maintenance Mode and then click Yes. After the host enters maintenance mode, click Reboot.



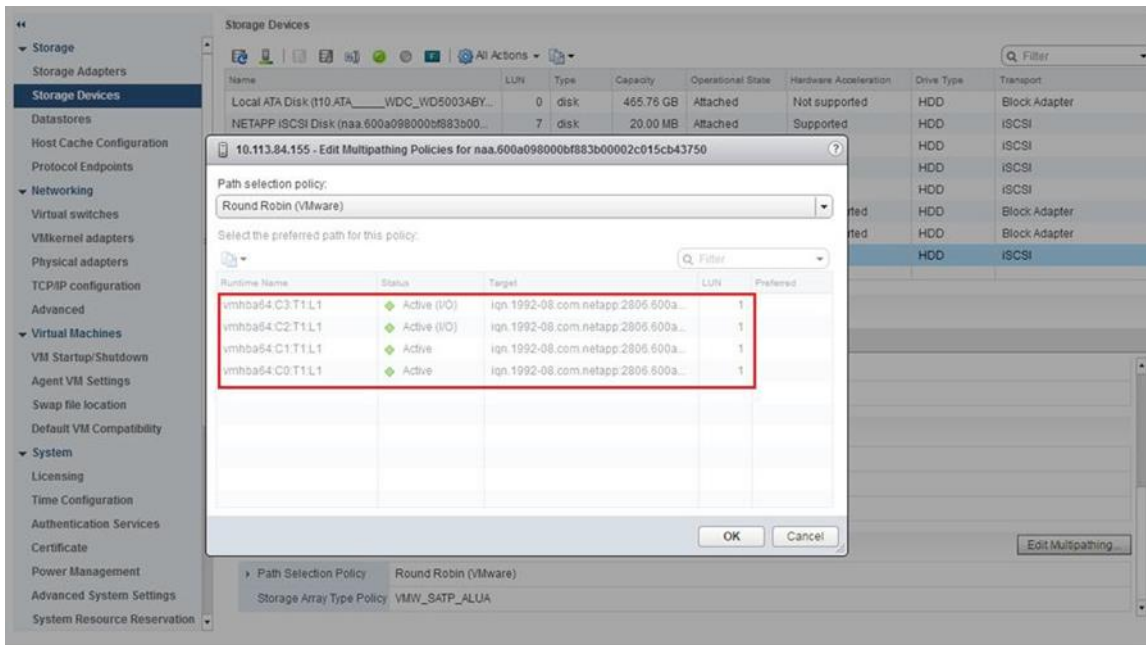
## Tuning VMware Settings to Maximize Performance

VMware ESXi 6.x and 7.x defaults to a round-robin multipath policy to balance I/O for each storage LUN across the available optimized paths to the storage system. After the NetApp E-Series devices (LUNs) have been discovered by an ESXi host, view the Manage Paths window for each E-Series device to verify that the paths are set up as expected.

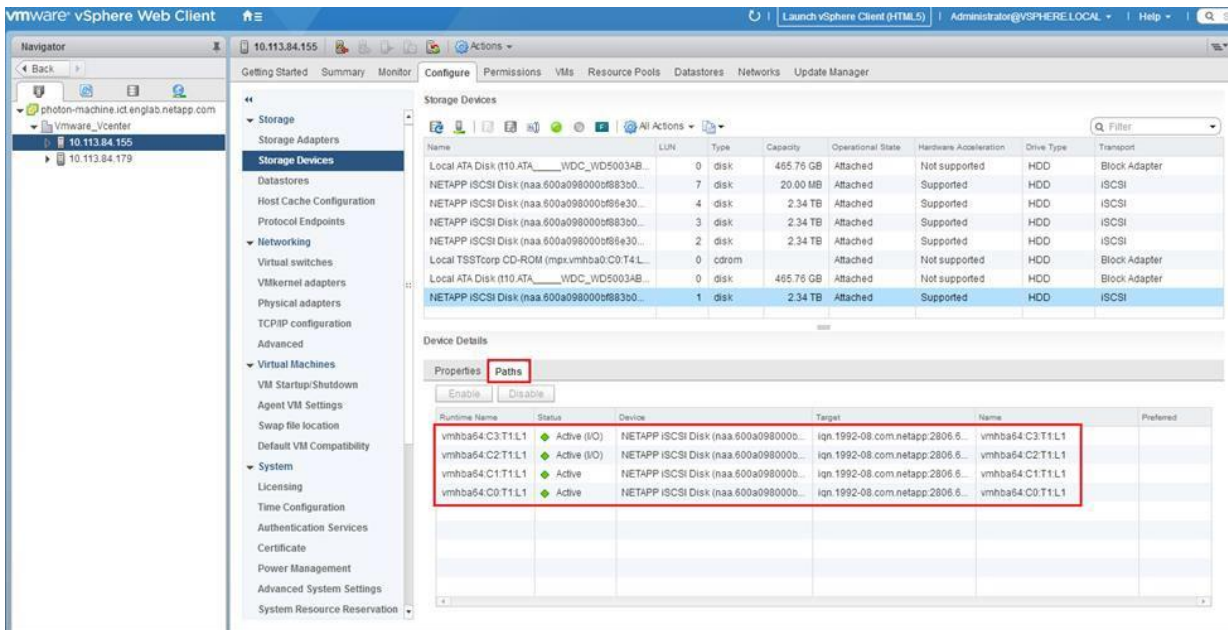
1. On the VMware vSphere web client, select the ESXi host and go to Configure > Storage Devices.
2. Select any iSCSI disk, go to Properties, and select Edit Multipathing.



With a fully configured E-Series storage system (using all four iSCSI ports) connected to two NIC ports on an ESXi host, there should be two active I/O paths and two active (nonoptimized) paths for each device.



3. You can also check the path by clicking Configure > Storage Devices.
4. Select any iSCSI disk and go to Path.



## I/O Operation Limit—Performance Implications

By default, the VMware multipath round-robin policy balances I/O requests across the available active (I/O) paths for each LUN by switching paths for each one thousand I/Os (IOOperations Limit).

Testing in our lab showed that the default IOOperations Limit (1,000) did not maximize load balancing on the host NIC ports. However, when the default I/O limit was adjusted to 250, the I/O load was much more evenly distributed between the two NIC ports on each host. For more information, see [Adjusting Round Robin IOPS limit from default 1000 to 1 \(2069356\)](#).

To view the current IOOperations Limit setting on the ESXi host, run the `esxcli storage nmp psp roundrobin deviceconfig get -d <device ID>` command.

```
~ # esxcli storage nmp psp roundrobin deviceconfig get -d naa.60080e50002935dc00003c7d540f7619
Byte Limit: 10485760
Device: naa.60080e50002937e0000044dd540f7483 IOOperation Limit: 1000
```

The default IOOperations Limit can be adjusted on an existing device as required by running the following command:

```
esxcli storage nmp psp roundrobin deviceconfig set -d <device ID> -t iops -I <1 to 1000>
```

Setting the value to 1 forces the ESXi server to send each I/O through a different path from the previous I/O whenever multiple active (I/O) paths are available. To return the setting to the default value of 1,000, run the `esxcli storage nmp psp roundrobin deviceconfig set -d <device ID> -t iops -I 1000` command.

```
~ # esxcli storage nmp psp roundrobin deviceconfig set -d naa.60080e50002935dc00003c7d540f7619 -t iops -I 1000
~ # esxcli storage nmp psp roundrobin deviceconfig get -d naa.60080e50002935dc00003c7d540f7619
Byte Limit: 10485760
Device: naa.60080e50002935dc00003c7d540f7619 IOOperation Limit: 1000 Limit Type: Iops
```

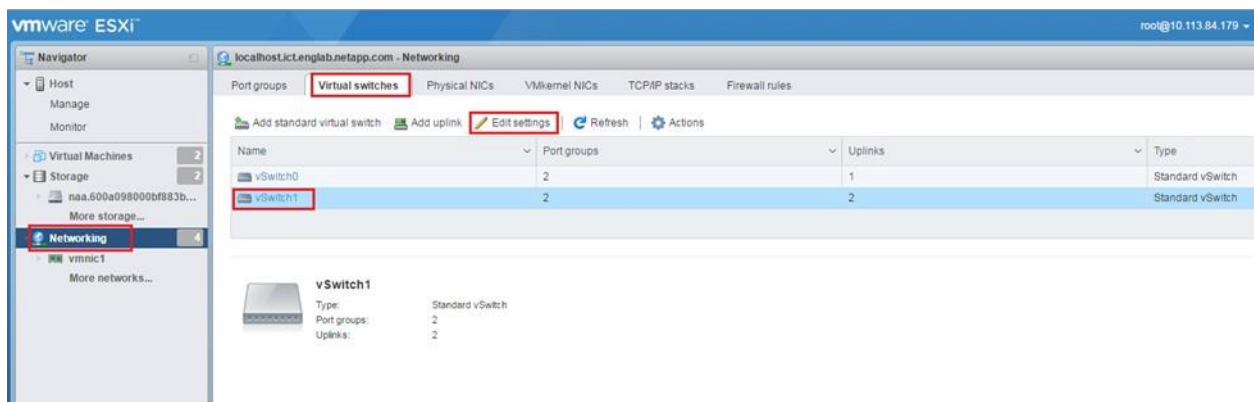
To automatically set the IOOperations limit when a new device is created in the ESXi host, create a claim rule that overrides the ESXi systemwide claim rule for E-Series storage systems by running the `esxcli storage nmp satp rule add -s "VMW_SATP_ALUA" -V "NETAPP" -M "INF-01-00" -P "VMW_PSP_RR" -O "iops=<1 to 1000>"` command. After new devices are created, be sure to confirm that the setting was successful by using the `esxcli storage nmp psp roundrobin deviceconfig get <device ID>` command.

## Jumbo Frames

In addition to setting the round-robin parameters, it is important to change the jumbo frames default setting to an MTU of 9,000 for all network interfaces in the I/O path between the host and the storage. This is not a global setting in the ESXi host and instead must be set in multiple locations, once on the virtual switch and again on each iSCSI VMkernel adapter. This task can be performed through the ESXi host and the VMware vSphere web client. Changing the jumbo frame setting from the VMware vSphere web client is shown in Appendix A.

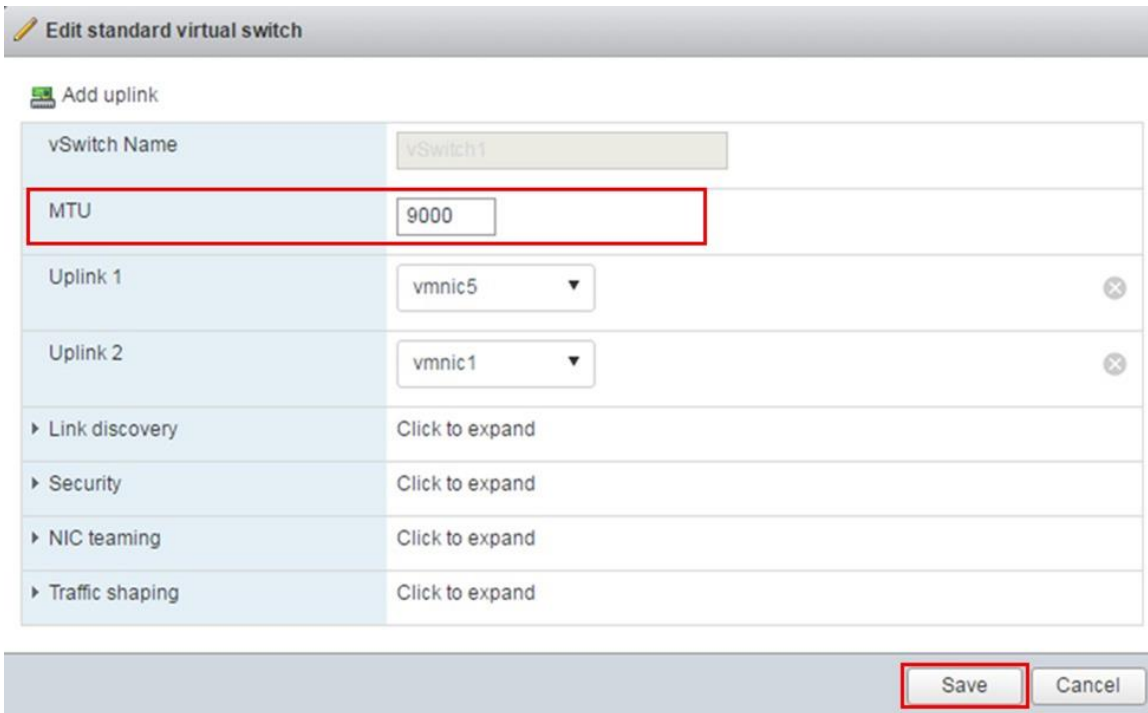
To change the jumbo frames setting using the ESXi interface, complete the following these steps:

1. In the VMware ESXi host view, log in to the ESXi host from the web browser.
2. To change the MTU on a virtual switch from the Navigator tab, go to Networking>Virtual Switches and click the virtual switch. Click Edit Settings.

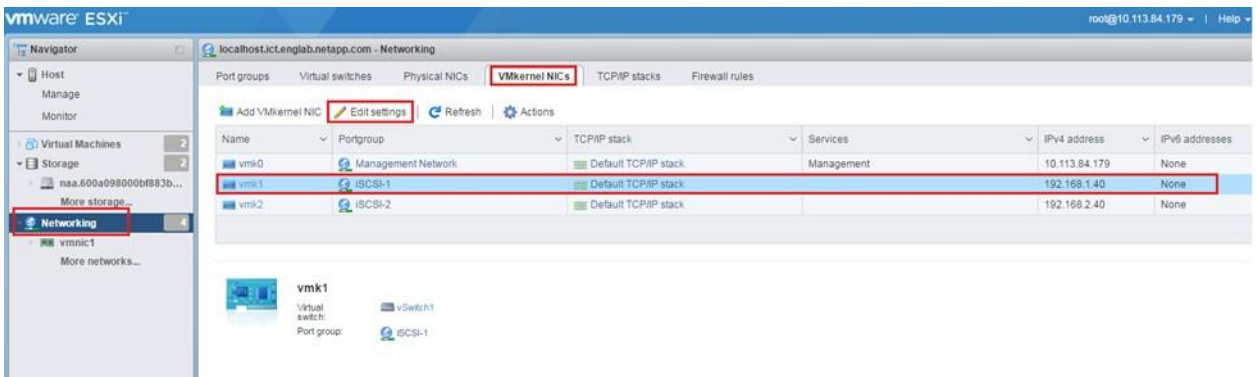





3. In the resulting window, change the MTU to 9000 and click Save.



4. To change the MTU on the VMkernel adapters, complete the following steps:
  - a. From the Navigator tab, go to Networking > VMkernel NICs and click the iSCSI VMkernel NIC. Select Edit Settings.



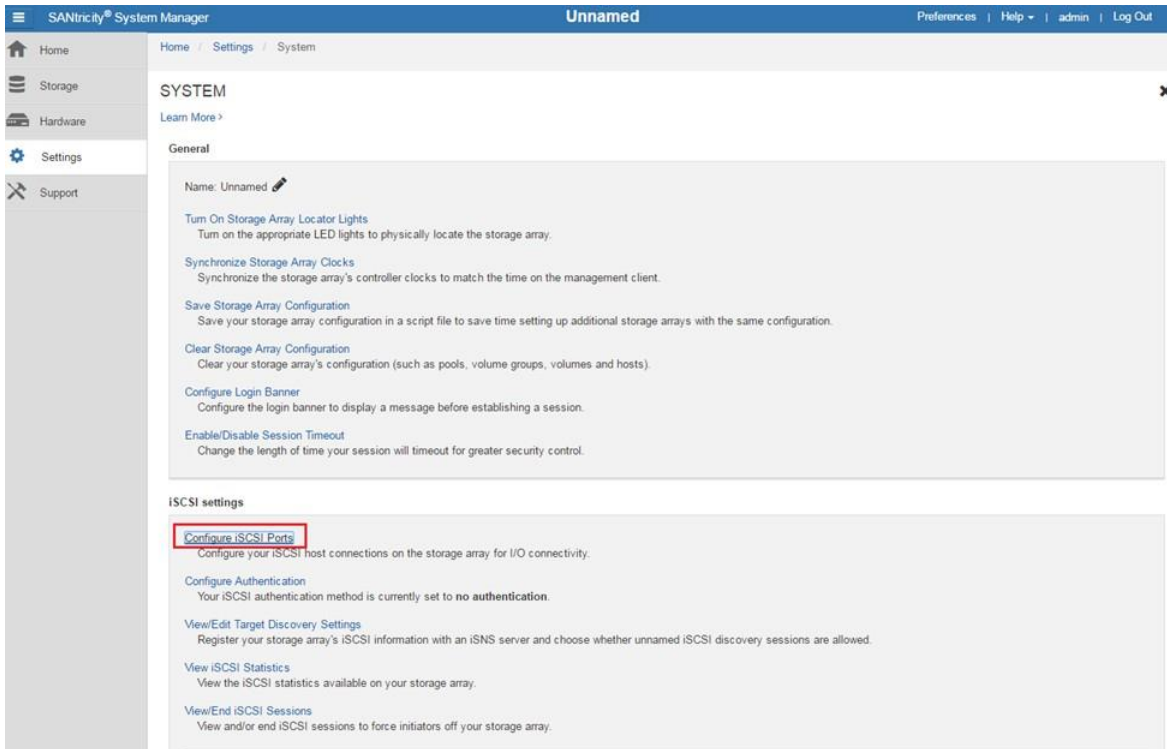
- e. In the resulting window, change the MTU to 9000 and click Save.

 Edit settings - vmk1

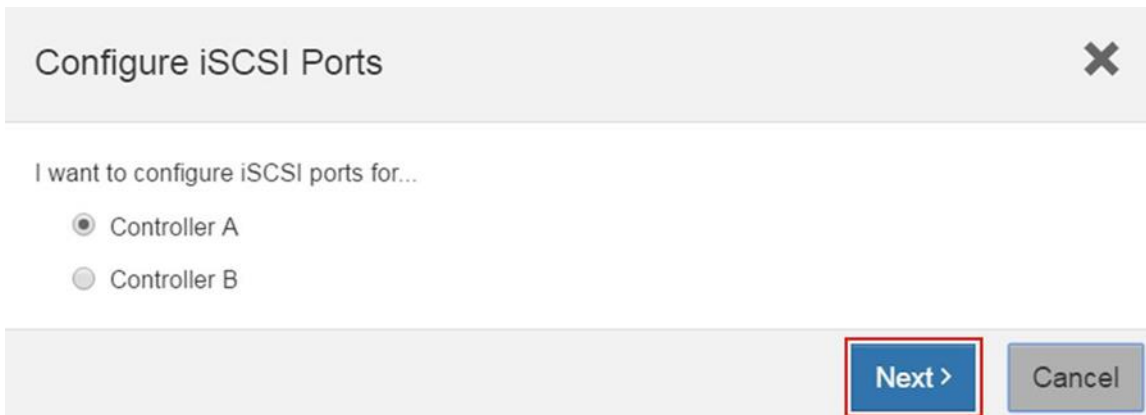
Port group	iSCSI-1
MTU	9000
IP version	IPv4 only
▶ IPv4 settings	<input type="radio"/> DHCP <input checked="" type="radio"/> Static
TCP/IP stack	Default TCP/IP stack
Services	<input type="checkbox"/> vMotion <input type="checkbox"/> Provisioning <input type="checkbox"/> Fault tolerance logging <input type="checkbox"/> Management <input type="checkbox"/> Replication <input type="checkbox"/> NFC replication

- f. Be sure to do this on all iSCSI VMkernel NICs.
5. In addition to the VMware configuration, jumbo frames must be enabled for each HIC port on the E-Series controllers. To change the jumbo frame setting on E-Series controllers, complete the following steps:
  - a. Log in to the E-series array SANtricity System Manager and go to Settings > System.
  - b. In the iSCSI settings, select Configure iSCSI Ports.



c. Select the controller and click Next.



d. Select the HIC port from the drop-down menu and click Next.

Configure iSCSI Ports ✕

**1** Select Port    **2** Configure Port    **3** Configure Network Settings

I want to configure network settings for the following Controller A port...

Port 0a Link Status: Up ▾

[What else do I need to do to configure or diagnose iSCSI?](#)

Cancel    **Next >**

e. Click Show More Port Settings.

Configure iSCSI Ports ✕

**1** Select Port    **2** Configure Port    **3** Configure Network Settings

I want to configure network settings for the following Controller A port...

Port 0a settings **Show more port settings**

MAC address: 00:A0:98:BF:86:E9

Enable IPv4:

Enable IPv6:

Enable ICMP PING responses (applies to all iSCSI ports on the storage array)

< Back    Cancel    **Next >**

f. Change the MTU to 9000 and click Next.

## Configure iSCSI Ports



1 Select Port

2 Configure Port

3 Configure Network Settings

I want to configure network settings for the following Controller A port...

Port 0a settings

[Show fewer port settings](#)

MAC address: 00:A0:98:BF:86:E9

Enable IPv4:

Enable IPv6:

Port 0a TCP listening port

3260

Port 0a MTU size

- 9000 + bytes per frame

**Note:** TCP listening port and MTU size settings apply to both IPv4 and IPv6.

Enable ICMP PING responses (applies to all iSCSI ports on the storage array)

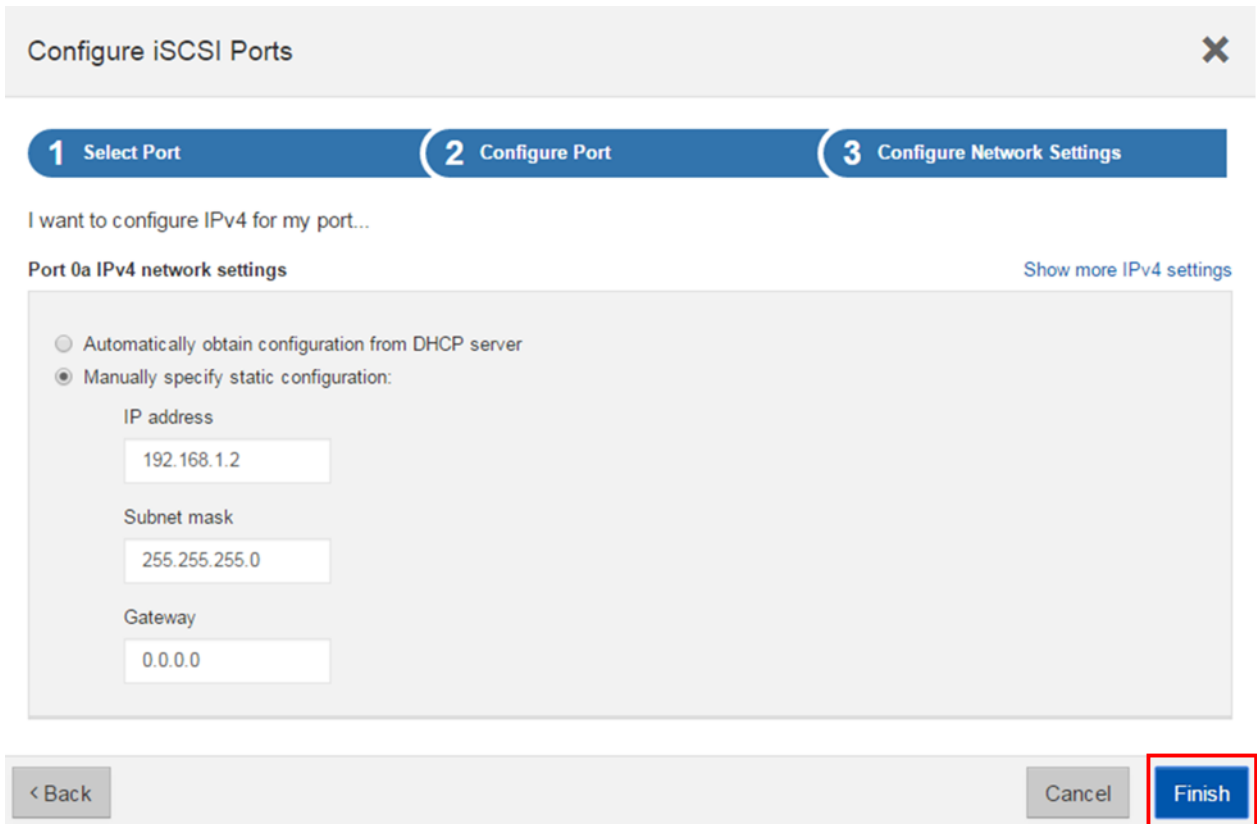
< Back

Cancel

Next >

- g. Make sure that the IP address is correct and click Finish.





6. Change the MTU settings for all HIC ports on Controller A and Controller B. Verify that the jumbo frame settings are set correctly from host to storage.
  - a. Log in to the ESXi host management IP address.
  - b. Run the `vmkping -s 8972 -d <target IP> -I <source VMK port ID>` command for each possible path combination so that all intended paths can pass large packets. For more information, see [Testing VMkernel network connectivity with the vmkping command \(1003728\)](#).

```
[root@localhost:~] vmkping -s 8972 -d 192.168.1.2 -I vmk1 PING 192.168.1.2 (192.168.1.2): 8972
data bytes
8980 bytes from 192.168.1.2: icmp_seq=0 ttl=64 time=0.792 ms 8980 bytes from 192.168.1.2:
icmp_seq=1 ttl=64 time=0.688 ms 8980 bytes from 192.168.1.2: icmp_seq=2 ttl=64 time=0.664 ms
--- 192.168.1.2 ping statistics ---
```

## Performance Degradation with Data-Assurance-Enabled Volumes and iSCSI

When using an iSCSI initiator to issue reads to an iSCSI volume with data assurance (DA) enabled, you might experience read performance degradation compared to a non-DA-enabled iSCSI volume. The degradation is more noticeable if the queue depth equals 1. Extensive performance tests were performed by E-Series engineering and the IOP (Interoperability) group. These tests determined that the main contributor to this performance effect is a TCP feature called Delayed Acknowledgment, which is enabled by default on most common host operating systems.

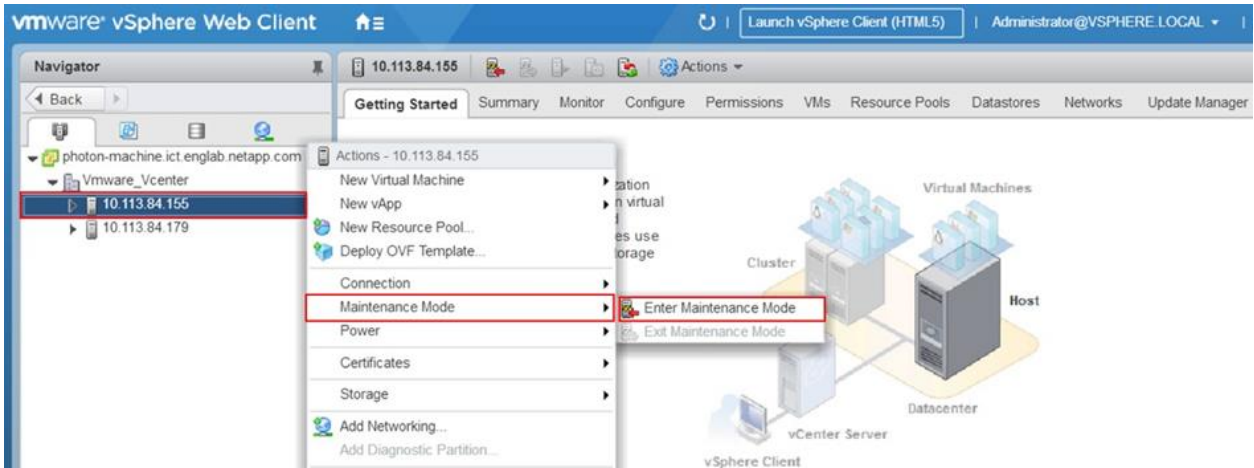
### Best Practice

Disable Delayed Acknowledgment on the host OS.

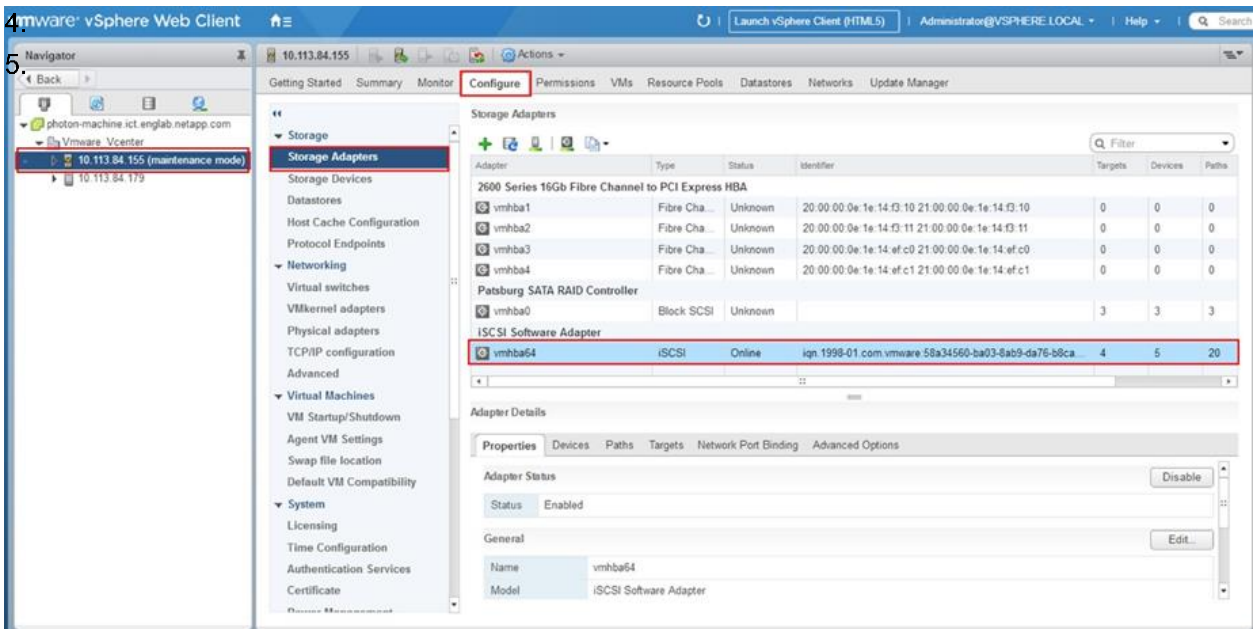
For more information, see [Performance Degradation with Data Assurance enabled Volumes and iSCSI and ESX/ESXi hosts might experience read or write performance issues with certain storage arrays \(1002598\)](#).

To disable Delayed Acknowledgment on the ESXi host, complete the follow steps:

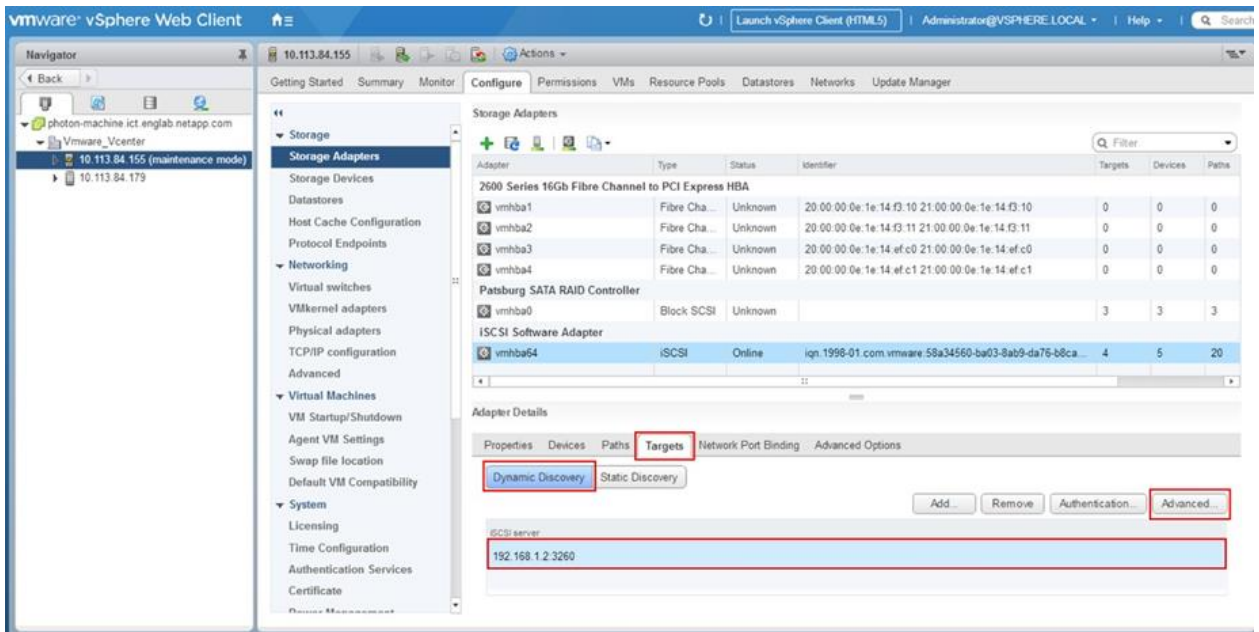
1. Log in to the vSphere Client and select the host.
2. Right-click the host, select Maintenance Mode, and select Enter Maintenance Mode.



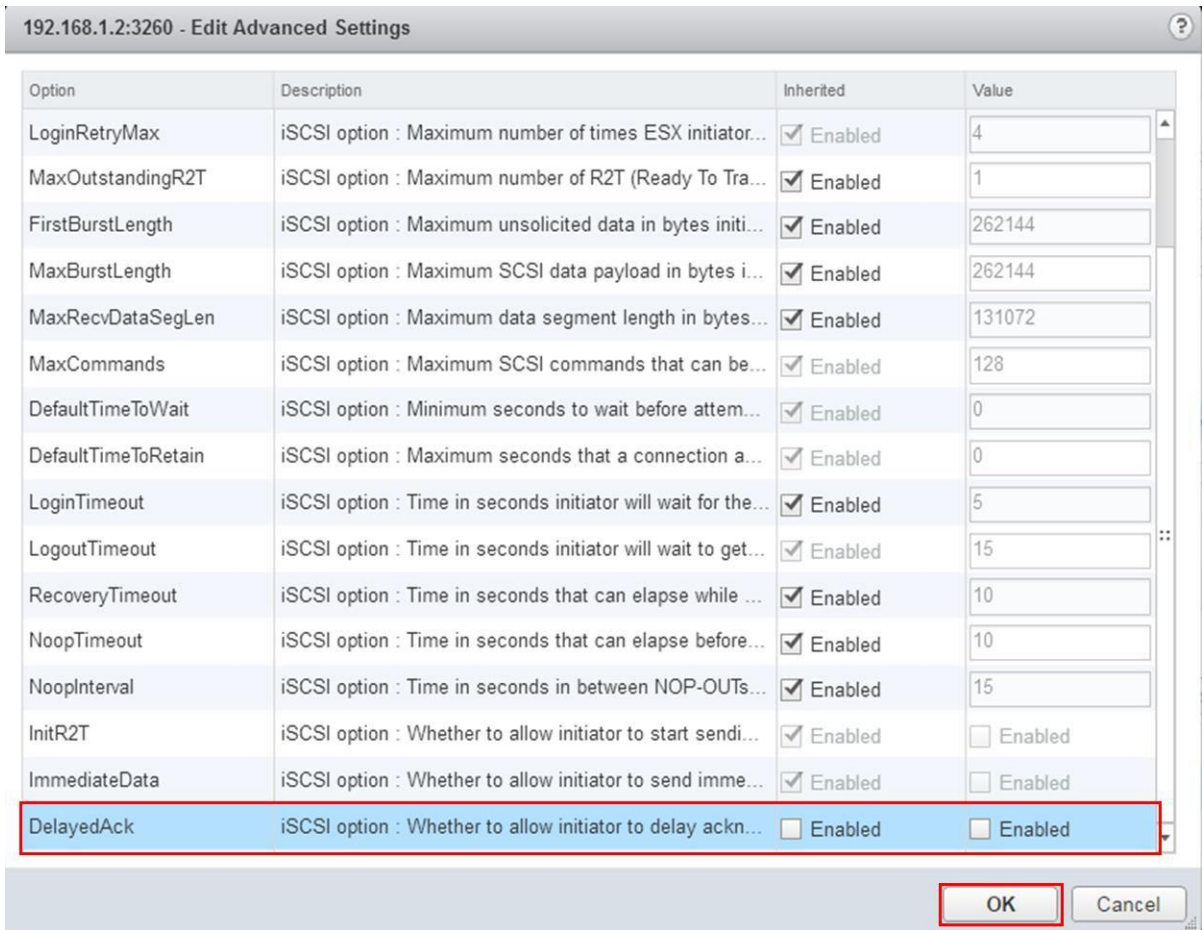
3. Wait for the process to complete. Navigate to the Configuration tab. Click Storage Adapters.



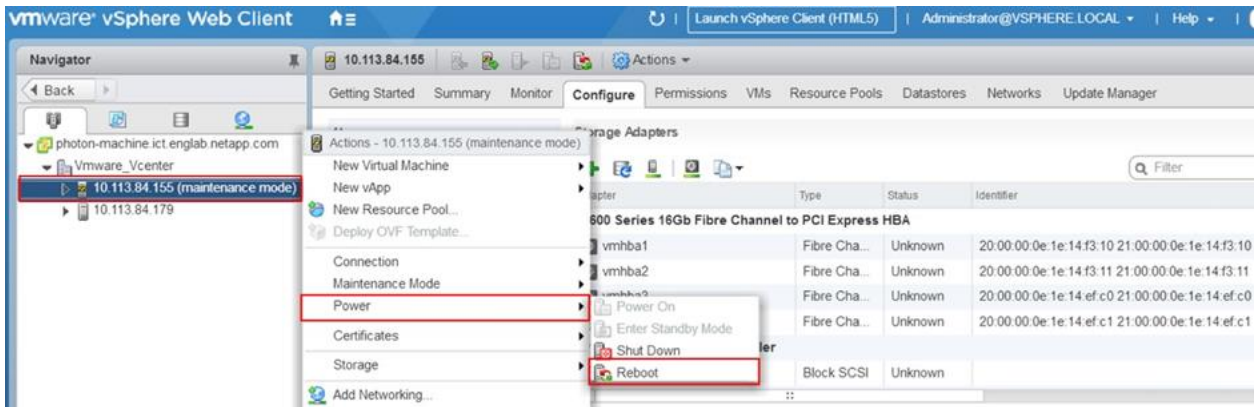
4. Click the iSCSI vmhba that you want to modify.
5. Modify the Delayed Acknowledgement setting on a discovery address.
  - a. Under Adapter Details, click the Targets tab.
  - b. Click Dynamic Discovery.
  - c. Click the Server Address tab.
  - d. Click Advanced.



- e. In the Edit Advanced Settings window, uncheck Inherited and Value for the DelayedAck option and then click OK.



## 6. Reboot the ESXi host.



## VMware Port Binding

By default, the VMware iSCSI initiator makes only a single connection to each target port presented by a storage system. The iSCSI port binding feature forces the iSCSI initiator to make connections from each host-side port to each target port. This feature is meant to be used with storage systems that present only a single IP address for the target.

Without the port-binding feature, regardless of how many host-side ports were configured and capable of connecting to the storage system, the ESXi host would make only a single connection to such a storage system. The remaining connections would never be used.

Conversely, if this feature is used with an E-Series storage system that presents multiple IP addresses, too many connections would be established. For example, there could be four host-side interfaces configured and a total of four target-side interfaces, two per controller. You would end up with a total of 16 connections to the storage system. This exceeds the maximum of eight paths per volume supported by VMware.

For more information, see [Considerations for using software iSCSI port binding in ESX/ESXi \(2038869\)](#). If iSCSI port binding is used when it should not be, you might experience longer rescan times and incorrect path detection.

### Best Practice

Do not use port binding with E-Series storage arrays.

## Conclusion

NetApp E-Series and EF-Series storage systems are well suited for serving workloads in VMware environments. The available host interfaces allow flexible integration, depending on performance and network architecture requirements. Administrators who spend much of their time working with the vSphere and vCenter interfaces will quickly appreciate the intuitive storage system management interface.

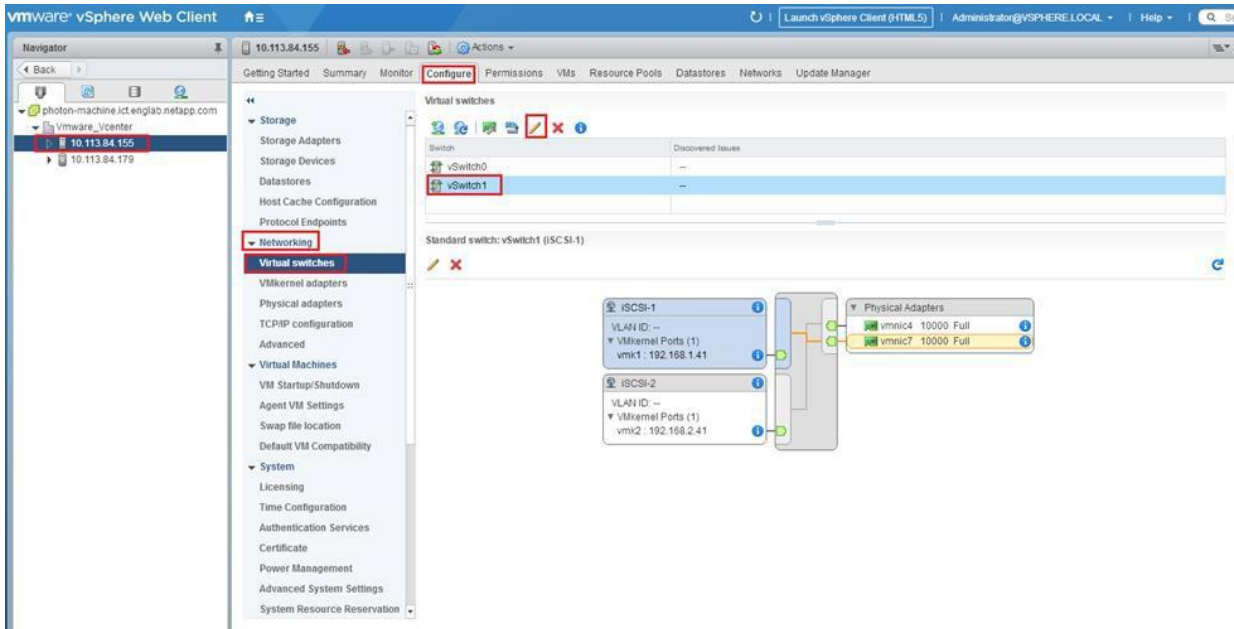
In addition to ease of administration and integration, E-Series storage systems provide dependable and consistent performance with easy-to-implement DDP architectures that are advantageous for random I/O workloads. E-Series systems can also be tuned for specific workloads across a wide range of RAID types, LUN capacities, and drive speeds, including adding SSDs to provide high-performance LUNs for demanding workloads.

# Appendix A: Changing Jumbo Frame Settings from a VMware vSphere Web Client

## Change the MTU on a Virtual Switch from a VMware vSphere Web Client

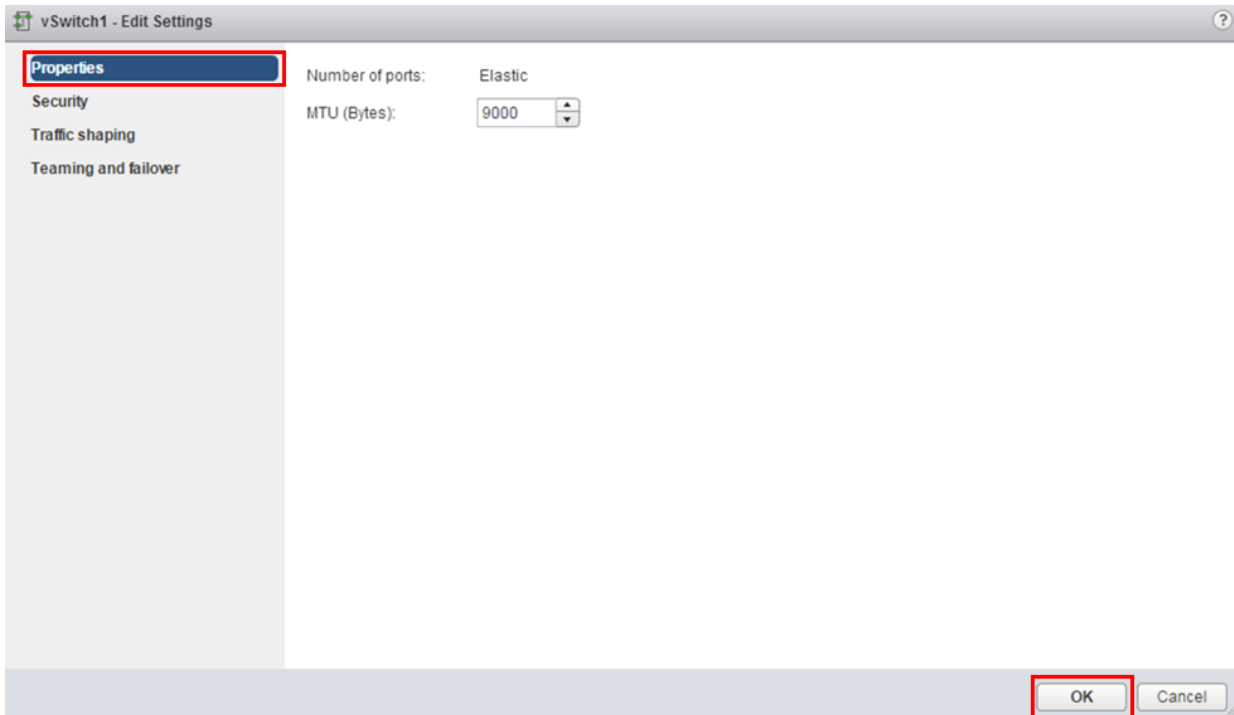
To change the MTU on a virtual switch from the VMware vSphere web client, follow these steps:

1. Select the ESXi host and go to Configure > Networking > Virtual Switches.
2. Select the virtual switch and click the Edit Settings (pencil) icon.



3. In the vSwitch1 – Edit Settings window, select Properties and change the MTU to 9000. Click OK.

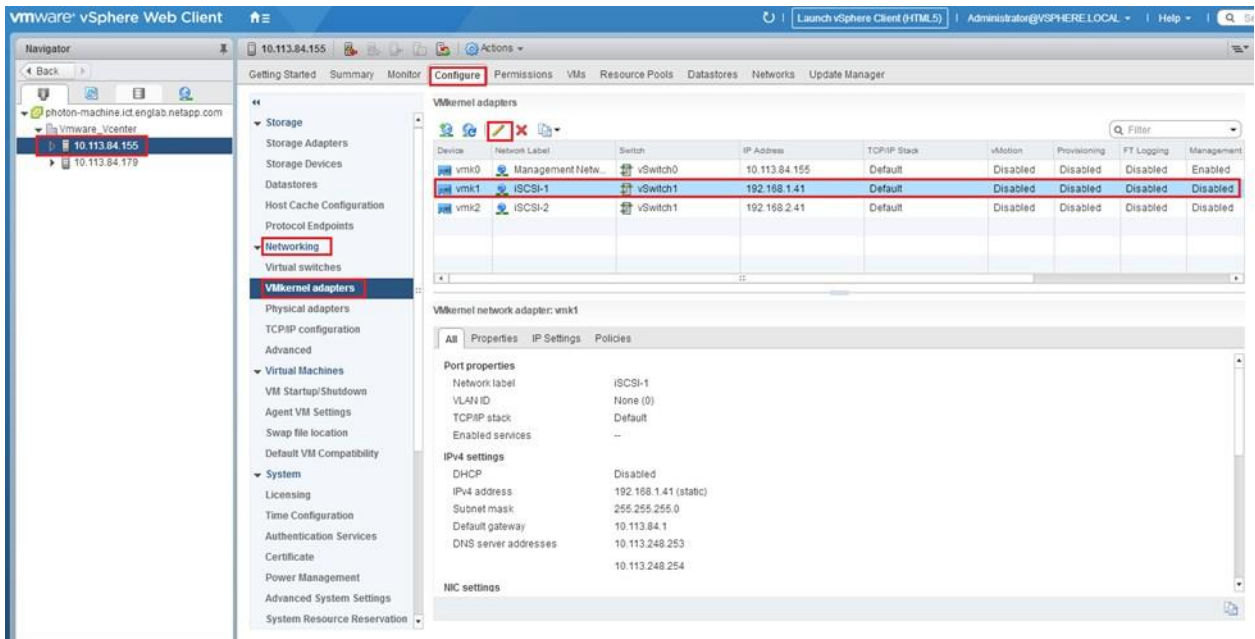




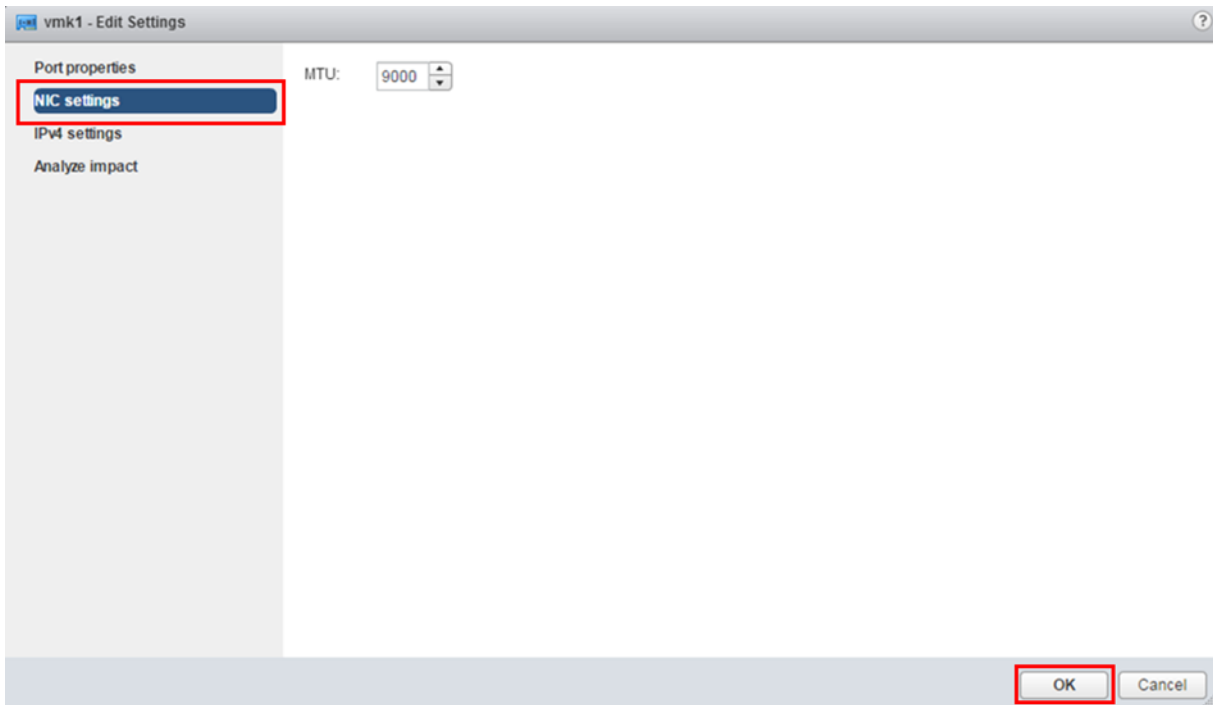
## Change the MTU on VMkernel Adapters

To change the MTU on VMkernel Adapters, complete the following steps:

1. Select the ESXi host and go to Configure > Networking > VMkernel Adapters.
2. Select the iSCSI adapter and click the Edit Settings (pencil) icon.



3. In the vmk1 – Edit Settings window, select NIC Settings and change the MTU to 9000. Click OK.




**Note:** Be sure to change the MTU size on all iSCSI VMkernel adapters.

## Appendix B: Configuring iSCSI CHAP Authentication

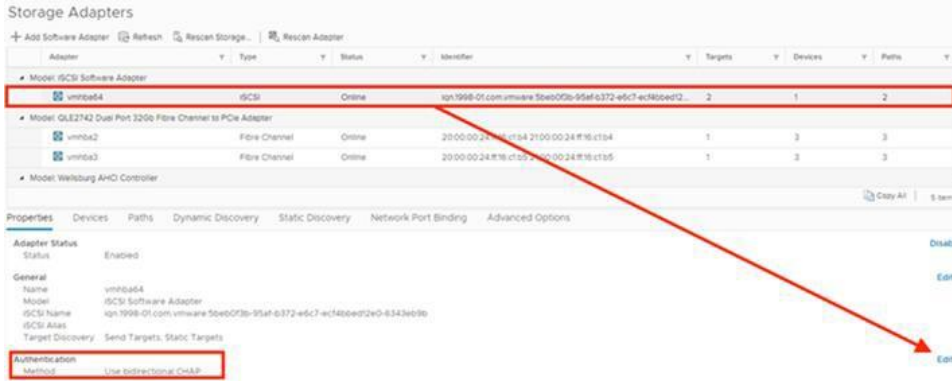
### VMware vSphere Web Client View

In VMware, the iSCSI Challenge-Handshake Authentication Protocol (CHAP) can be configured using vSphere on either the parent iSCSI software adapter or on individual iSCSI targets listed under the Dynamic Discovery and Static Discovery tabs. The following steps show how to configure iSCSI CHAP using the vSphere Web Client.


1. Select the desired ESXi host from the vSphere Cluster inventory and then select Configure > Storage Adapters.



2. **Option A.** To configure iSCSI CHAP on the parent iSCSI software adapter, select the adapter and then click Edit. This option applies to all iSCSI target connections.

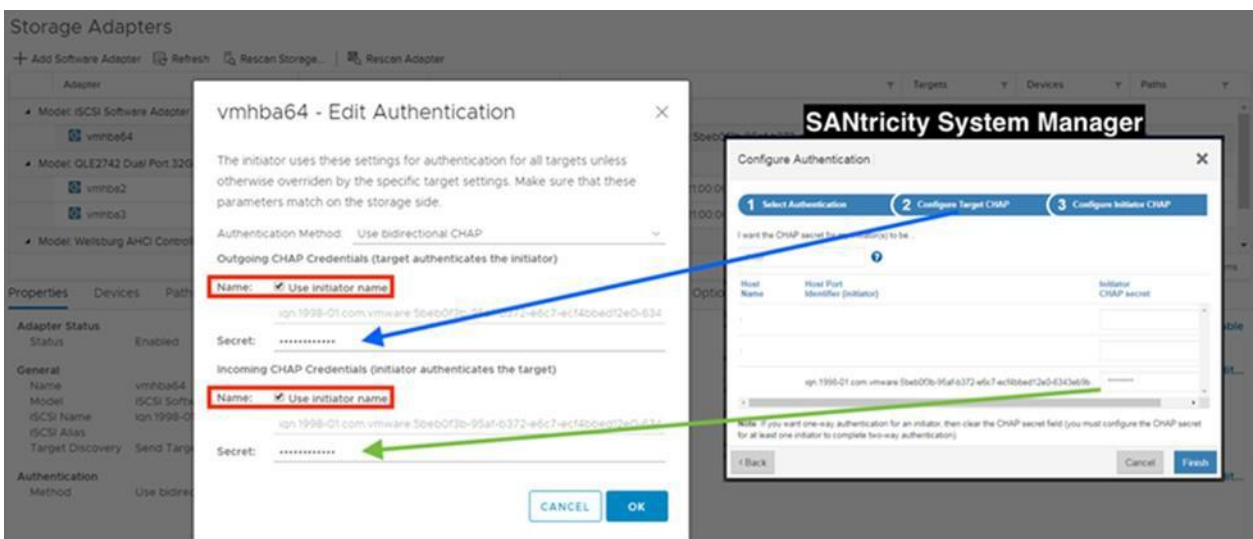


**Option B.** To configure iSCSI CHAP on individual iSCSI servers (targets) listed under the Dynamic Discovery or Static Discovery tabs, select the server and then click Authentication.



In the subsequent Edit Authentication dialogue, fill out the configuration defined on the E-Series array (see step 3 for details). If you are editing the CHAP configuration on individual iSCSI servers, you might need to uncheck Inherit Settings from Parent.

3. To configure iSCSI authentication on the E-Series array, complete the following steps:
  - a. In SANtricity System Manager, select Settings in the left-hand pane.
  - b. Click the System icon.
  - c. From the iSCSI settings section click Configure Authentication.
  - d. Complete the resulting wizard depicted in the following screenshot to correlate SANtricity System Manager and ESXi configuration parameters.



**Note:** This screenshot shows the last configuration step, but the target CHAP secret is defined in step 2.



## Related Resources

- [Configuring iSCSI Authentication \(E-Series\)](#)
- [Configuring CHAP Parameters for iSCSI Adapters \(VMware\)](#)

## Where to Find Additional Information

To learn more about the information that is described in this document, review the following documents and/or websites:

- NetApp E-Series and SANtricity 11 Documentation Center (VMware Express Configuration)  
[https://docs.netapp.com/ess-11/index.jsp?topic=%2Fcom.netapp.doc.ssm-exp-ic-vm%2Fhome.html&cp=3\\_1](https://docs.netapp.com/ess-11/index.jsp?topic=%2Fcom.netapp.doc.ssm-exp-ic-vm%2Fhome.html&cp=3_1)
- NetApp Hardware Universe  
<https://hwu.netapp.com/Controller/Index?platformTypeId=2357027>
- VMware KB article: Testing VMkernel Network Connectivity with the vmkping Command (1003728)  
<https://kb.vmware.com/s/article/1003728>
- NetApp KB article: Performance Degradation with Data Assurance enabled Volumes and iSCSI  
[https://kb.netapp.com/app/answers/answer\\_view/a\\_id/1074155/~/performance-degradation-with-data-assurance-enabled-volumes-and-iscsi-](https://kb.netapp.com/app/answers/answer_view/a_id/1074155/~/performance-degradation-with-data-assurance-enabled-volumes-and-iscsi-)
- VMware KB article: Considerations for Using Software iSCSI Port Binding in ESX/ESXi (2038869)  
<https://kb.vmware.com/s/article/2038869>
- VMware KB article: Adjusting Round Robin IOPS Limit from Default 1,000 to 1 (2069356)  
<https://kb.vmware.com/s/article/2069356>
- VMware KB article: ESX/ESXi Hosts Might Experience Read or Write Performance Issues with Certain Storage Arrays (1002598)  
<https://kb.vmware.com/s/article/1002598>
- NetApp technical report: TR-4725: Introduction to NetApp E-Series E2800 Arrays  
<https://www.netapp.com/us/media/tr-4725.pdf>
- NetApp technical report: TR-4724: Introduction to NetApp E-Series E5700 Arrays  
<https://www.netapp.com/us/media/tr-4724.pdf>
- VMware Storage and Availability Technical Documents  
<https://storagehub.vmware.com>
- NetApp TechComm TV video: VMware Configuration Guide for E-Series Integration with ESXi  
<https://www.youtube.com/watch?v=qgZr5LkK144>

## Version History

Version	Date	Document Version history
Version 1.1	November 2020	Added Host Block Size Requirements section.
Version 1.0	July 2019	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.

### **Copyright Information**

Copyright © 2020 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

### **Trademark Information**

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

TR-4789-1120