



# CHAPTER 13

## Configuring iSCSI Multipath

This chapter describes how to configure iSCSI multipath for multiple routes between a server and its storage devices.

### iSCSI Multipath Description

The iSCSI multipath feature sets up multiple routes between a server and its storage devices for maintaining a constant connection and balancing the traffic load. The multipathing software handles all input and output requests and passes them through on the best possible path. Traffic from host servers is transported to shared storage using the iSCSI protocol that packages SCSI commands into iSCSI packets and transmits them on the Ethernet network.

iSCSI multipath provides failover. In the event a path, or any of its components, fails, the server selects another available path. In addition to path failover, multipathing load balances by distributing storage loads across multiple physical paths to reduce or remove potential bottlenecks.

The following are descriptions of the iSCSI multipath processes:

iSCSI Multipath Process	Provided by Cisco Nexus 1000V	Provided by VMware code in VMkernel	Description
<b>Uplink Pinning</b>	X		Each VMKernel port created for iSCSI access are pinned to one physical NIC. This overrides any NIC teaming policy or port bundling policy. All traffic from the VMKernel port uses only the pinned uplink to reach the upstream switch.
<b>Storage Binding</b>		X	The ESX/ESXi host creates iSCSI adapters, also called VMware iSCSI host bus adapter (VMHBA) adapters, for the physical NICs. For Software iSCSI, only one VMHBA is created for all the physical NICs. For Hardware iSCSI, one VMHBA is created for each physical NIC that supports iSCSI offload in hardware. Storage binding is the step in which each VMKernel port is pinned to the VMHBA associated with the physical NIC to which the VMKernel port is pinned.
<b>Path selection</b>		X	

iSCSI Multipath Process	Provided by Cisco Nexus 1000V	Provided by VMware code in VMkernel	Description
Path failover		X	In the event of failure, in the path or any of its components, the server selects another available path.
Load Balancing			Reduces or removes bottlenecks by distributing storage loads across multiple physical paths.

The Cisco Nexus 1000V DVS performs iSCSI multipathing regardless of the iSCSI target. The iSCSI daemon on an ESX server communicates with a target in multiple sessions by pinning two or more host VMkernel NICs on the same Cisco Nexus 1000V. Therefore, the Cisco Nexus 1000V should have a minimum of two physical NICs.

Once you enable iSCSI multipath on the port profile, using the command **capability iscsi-multipath**, the VEM automatically pins the VMkernel NICs to VM NICs.

Standard NICs connect the host to a remote iSCSI target on the IP network. The software iSCSI adapter that is built into ESX/ESXi communicates with the physical NICs through the network stack.

For detailed information about how to use VMware® ESX™ and VMware ESXi systems with an iSCSI storage area network (SAN), see the [iSCSI SAN Configuration Guide](#).

## iSCSI Dependent Hardware Adapter

A third-party dependent hardware iSCSI adapter, such as a Broadcom Ethernet NIC with iSCSI offload functions, offloads iSCSI and network processing from your host leveraging VMware iSCSI management and configuration interfaces.

After this NIC is installed, the dependent hardware iSCSI adapter is loaded onto the host and appears in the list of storage adapters. Although the adapter is enabled by default, to make it functional, VMkernel networking must be set up for the iSCSI traffic and the adapter must be bound to a VMkernel iSCSI port.

Only one VMHBA adapter is available for each NIC that is capable of iSCSI offload. The VMkernel port is bound to the VMHBA adapter of the hardware NIC (vnic) to which it is pinned. In [Figure 13-1](#), VMkernel NIC 1 is bound to vmhba33.

## iSCSI Adapter on the VMware vSwitch

VMkernel iSCSI ports, for the traffic between the iSCSI adapter and the physical NIC, must be created before enabling or configuring the software or hardware iSCSI for multipathing. Only one adapter (VMHBA) is created for software iSCSI and all VMkernel ports are bound to it.

## iSCSI Multipath Setup on the VMware Switch

Before enabling or configuring multipathing, networking must be configured for the software or hardware iSCSI adapter. This involves opening a VMkernel iSCSI port for the traffic between the iSCSI adapter and the physical NIC.

For software iSCSI, only one adapter is required for the entire implementation. All VMkernel ports are bound to this adapter. For example, in Figure 13-1, the following adapters and NICs are used.

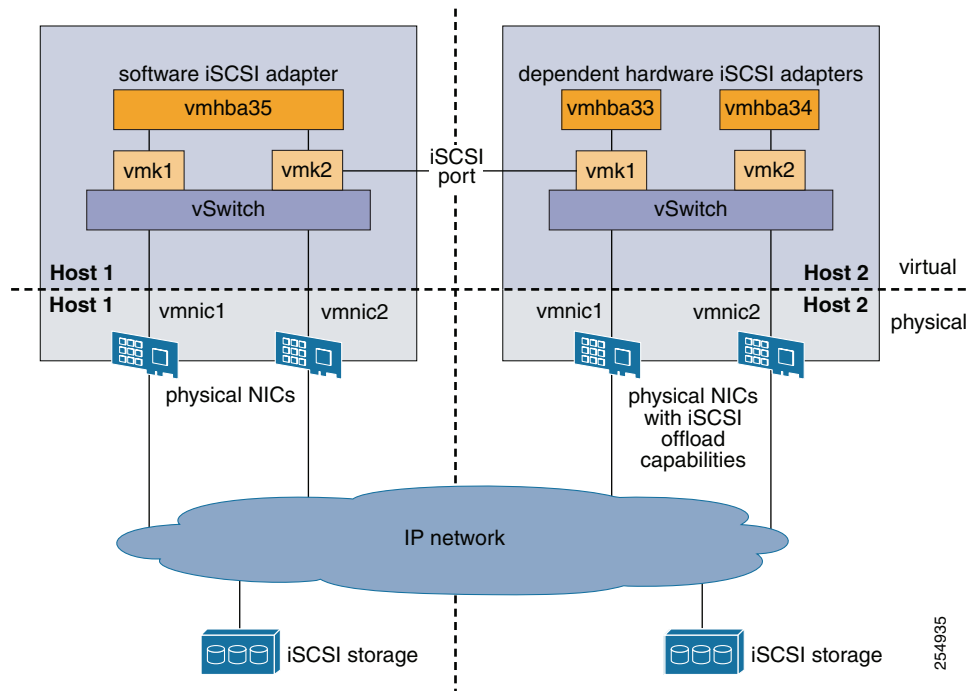
Hardware iSCSI Adapter	VMkernel NIC	VM NIC
VMHBA	1	1
	2	2

For hardware iSCSI, a separate adapter is required for each NIC. Each VMkernel port is bound to the adapter of the physical VM NIC to which it is pinned. For example, in Figure 13-1, the following adapters and NICs are used.

Hardware iSCSI Adapter	VMkernel NIC	VM NIC
vmhba33	1	1
vmhba34	2	2

Figure 13-1 shows the setup of iSCSI multipath on a VMware virtual switch.

Figure 13-1 iSCSI Multipath on VMware Virtual Switch



254935

## Guidelines and Limitations

The following are guidelines and limitations for the iSCSI multipath feature.

- Only port profiles of type vEthernet can be configured with **capability iscsi-multipath**.
- The port profile used for iSCSI multipath must be an access port profile. It cannot be a trunk port profile.

- The following are not allowed on a port profile configured with **capability iscsi-multipath**
  - The port profile cannot also be configured with **capability l3 control**.
  - A system VLAN change when the port profile is inherited by VMkernel NIC.
  - An access VLAN change.
  - A port mode change to trunk mode.
- Only VMkernel NIC ports can inherit a port profile configured with **capability iscsi-multipath**.
- A VMware Kernel NIC can only be pinned or assigned to one physical NIC.
- A physical NIC can have multiple VMware Kernel NICs pinned or assigned to it.
- The iSCSI initiators and storage must already be operational.
- ESX 4.0 Update1 or later supports software iSCSI multipathing only.
- ESX 4.1 or later supports both software and hardware iSCSI multipathing.
- The iSCSI adapter must be enabled and bound to a VMkernel iSCSI port.
- VMkernel networking must be functioning for the iSCSI traffic.

## Prerequisites

The iSCSI Multipath feature has the following prerequisites.

- You must understand VMware iSCSI SAN storage virtualization.
- You must know how to set up the iSCSI Initiator software on your VMware ESX/ESXi host.
- The host is already functioning with the VMware ESX 4.0.1 Update 01 software release.
- You must understand iSCSI multipathing and path failover.
- A system VLAN is created on the Cisco Nexus 1000V.
  - One of the uplink ports must already have this VLAN in its system VLAN range.
- The host is configured with one port channel that includes two or more physical NICs.
- VMware kernel NICs configured to access the SAN external storage are required.

## Default Settings

Table 13-1 lists the default settings in the iSCSI Multipath configuration.

**Table 13-1** *iSCSI Multipath Defaults*

Parameter	Default
Type (port-profile)	vEthernet
Description (port-profile)	None
VMware port group name (port-profile)	The name of the port profile
Switchport mode (port-profile)	Access
State (port-profile)	Disabled

# Configuring a Port Profile for iSCSI Multipath

Use this section to configure communication multipathing between hosts and targets over iSCSI protocol by assigning the vEthernet interface to an iSCSI multipath port profile configured with a system VLAN.

## Creating a Port Profile for VMkernel Ports

You can use this procedure to create a port profile for VMkernel ports.

### BEFORE YOU BEGIN

Before starting this procedure, you must know or do the following.

- You have already configured the host with one port channel that includes two or more physical NICs.
- You have already created VMware kernel NICs to access the SAN external storage.
- A VMware Kernel NIC can only be pinned or assigned to one physical NIC.
- A physical NIC can have multiple VMware Kernel NICs pinned or assigned to it.
- Multipathing must be configured on the interface by using this procedure to create an iSCSI multipath port profile and then assigning the interface to it.
- You are logged in to the CLI in EXEC mode.
- You know the VLAN ID for the VLAN you are adding to this iSCSI multipath port profile.
  - The VLAN must already be created on the Cisco Nexus 1000V.
  - The VLAN that you assign to this iSCSI multipath port profile must be a system VLAN.
  - One of the uplink ports must already have this VLAN in its system VLAN range.
- The port profile must be an access port profile. It cannot be a trunk port profile. This procedure includes steps to configure the port profile as an access port profile.

### SUMMARY STEPS

1. **config t**
2. **port-profile type vethernet***name*
3. **vmware port-group** [*name*]
4. **switchport mode access**
5. **switchport access vlan** *vlanID*
6. **no shutdown**
7. (Optional) **system vlan** *vlanID*
8. **capability iscsi-multipath**
9. **state enabled**
10. (Optional) **show port-profile** *name*
11. (Optional) **copy running-config startup-config**

## DETAILED STEPS

	Command	Purpose
Step 1	<b>config t</b>  <b>Example:</b> <pre>n1000v# config t n1000v(config)#</pre>	Places you in the CLI Global Configuration mode.
Step 2	<b>port-profile type vethernet name</b>  <b>Example:</b> <pre>n1000v(config)# port-profile type vethernet VMK-port-profile n1000v(config-port-prof)#</pre>	Places you into the CLI Port Profile Configuration mode for the specified port profile. <ul style="list-style-type: none"> <li><b>type:</b> Defines the port-profile as Ethernet or vEthernet type. Once configured, this setting cannot be changed. The default is vEthernet type.</li> </ul> <b>Note</b> If a port profile is configured as an Ethernet type, then it cannot be used to configure VMware virtual ports. <ul style="list-style-type: none"> <li><i>name:</i> The port profile name can be up to 80 characters and must be unique for each port profile on the Cisco Nexus 1000V.</li> </ul>
Step 3	<b>description profiledescription</b>  <b>Example:</b> <pre>n1000v(config-port-prof)# description "Port Profile for iSCSI multipath" n1000v(config-port-prof)#</pre>	Adds a description to the port profile. This description is automatically pushed to the vCenter Server. <i>profile description:</i> up to 80 ASCII characters <b>Note</b> If the description includes spaces, it must be surrounded by quotations.
Step 4	<b>vmware port-group [name]</b>  <b>Example:</b> <pre>n1000v(config-port-prof)# vmware port-group VMK-port-profile n1000v(config-port-prof)#</pre>	Designates the port-profile as a VMware port group.  The port profile is mapped to a VMware port group of the same name. When a vCenter Server connection is established, the port group created in Cisco Nexus 1000V is then distributed to the virtual switch on the vCenter Server.  <i>name:</i> The VMware port group name.If you want to map the port profile to a different port group name, use the alternate name.
Step 5	<b>switchport mode access ]</b>  <b>Example:</b> <pre>n1000v(config-port-prof)# switchport mode access n1000v(config-port-prof)#</pre>	Designates that the interfaces are switch access ports (the default).
Step 6	<b>switchport access vlan vlanID</b>  <b>Example:</b> <pre>n1000v(config-port-prof)# switchport access vlan 254 n1000v(config-port-prof)#</pre>	Assigns the system VLAN ID to the access port for this port profile. <b>Note</b> The VLAN assigned to this iSCSI port profile must be a system VLAN.
Step 7	<b>no shutdown</b>  <b>Example:</b> <pre>n1000v(config-port-prof)# no shutdown n1000v(config-port-prof)#</pre>	Administratively enables all ports in the profile.

	Command	Purpose
Step 8	<b>system vlan <i>vlanID</i></b>  <b>Example:</b> n1000v(config-port-prof)# system vlan 254 n1000v(config-port-prof)#	Adds the system VLAN to this port profile.  This ensures that, when the host is added for the first time or rebooted later, the VEM will be able to reach the VSM. One of the uplink ports must have this VLAN in its system VLAN range.
Step 9	<b>capability iscsi-multipath</b>  <b>Example:</b> n1000v(config-port-prof)# capability iscsi-multipath n1000v(config-port-prof)#	Allows the port to be used for iSCSI multipathing.  In vCenter Server, the iSCSI Multipath port profile must be selected and assigned to the VM kernel NIC port.
Step 10	<b>state enabled</b>  <b>Example:</b> n1000v(config-port-prof)# state enabled n1000v(config-port-prof)#	Enables the port profile.  The configuration for this port profile is applied to the assigned ports, and the port group is created in the VMware vSwitch on the vCenter Server.
Step 11	<b>show port-profile name <i>name</i></b>  <b>Example:</b> n1000v(config-port-prof)# show port-profile name multipath-profile n1000v(config-port-prof)#	(Optional) Displays the current configuration for the port profile.
Step 12	<b>copy running-config startup-config</b>  <b>Example:</b> n1000v(config-port-prof)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

## Creating VMKernel Ports and Attaching the Port Profile

You can use this procedure to create VMkernel ports and attach a port profile to them.

### BEFORE YOU BEGIN

Before starting this procedure, you must know or do the following:

- You have already created a port profile by following the [“Creating a Port Profile for VMkernel Ports” procedure on page 13-5](#).
- The VMKernel ports are created directly on the vSphere client.
- You cannot have two different port profiles for iSCSI traffic. All VMware VMkernel NICs for iSCSI must be attached to the same port profile.
- Create one VMkernel NIC for each physical NIC that carries the iSCSI VLAN. The number of paths to the storage device is the same as the number of VMkernel NICs created. The number of VMKs must match the number of physical uplink ports that are attached to the Cisco Nexus 1000V. Then, the Cisco Nexus 1000V can assign a VMK and pin it to a physical uplink.
- The VMkernel NICs you create in this procedure may also carry other VLANs.

---

**Step 1** Create one VMkernel NIC for each physical NIC that carries the iSCSI VLAN.

For example, if you want to configure two paths, create two physical NICs on the Cisco Nexus 1000V DVS to carry the iSCSI VLAN. Create two vmknics for 2 paths.

**Step 2** Attach the port profile configured with **capability iscsi-multipath** to the VMkernel ports.

## Related Documents

Related Topic	Document Title
VMware SAN Configuration	<a href="#">VMware SAN Configuration Guide</a>
Port Profile Configuration	<i>Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.0(4)SV1(3)</i>
Interface Configuration	<i>Cisco Nexus 1000V Interface Configuration Guide, Release 4.0(4)SV1(3)</i>
Complete command syntax, command modes, command history, defaults, usage guidelines, and examples for all Cisco Nexus 1000V commands.	<i>Cisco Nexus 1000V Command Reference, Release 4.0(4)SV1(3)</i>

## Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

## Feature History for iSCSI Multipath

[Table 13-2](#) lists the release history for the iSCSI Multipath feature.

**Table 13-2** Feature History for iSCSI Multipath

Feature Name	Releases	Feature Information
iSCSI Multipath	4.0(4)SV1(2)	The iSCSI Multipath feature was added.