

QLogic Adapters and Cisco Nexus 5000 Series Switches: Fibre Channel over Ethernet Design Guide

Contents

Introduction	3
FCoE Protocol	3
Implementing Baseline FCoE	4
Solution Topology 1	5
Configuring the Cisco Nexus 5000 Series in NPIV Mode	5
Dynamic Load Balancing	8
NPV Traffic Management with the NPV Traffic Map	8
List of External Interfaces per Server Interface	8
Solution Topology 2	10
Configuring the Cisco Nexus 5000 Series in Switch Mode	10
Configuring FCoE	18
Configuring VLAN-to-VSAN Mapping	18
Configuring 10 Gigabit Ethernet Interfaces for FCoE	18
Configuring Virtual Fibre Channel Interfaces	19
Assigning vFC Interfaces to the VSAN	20
Configuring Management and Access Control	21
Configuring Role-Based Access	21
Configuring TACACS+ Authentication	22
Configuring TACACS+ Accounting	25
Configuring the QLogic 10 Gigabit Enhanced Ethernet CNA	26
Configuring the Adapter Side	27
Bringing Up the FCoE Link	28
Debugging	32
Using the QLogic SANsurfer CLI	34
Managing the Ethernet (Networking) Interface of the CNA	36
Windows Platform	36
Using the SANsurfer CNA Networking CLI	36
Creating Teams or Bonds	38
Creating VLANs	39
Using Microsoft Windows Property Pages	40
Linux Platform	42
SuSE Platform	43
Creating Teams or Bonds	43
Creating VLANs	45
RHEL Platform	45
Creating Teams or Bonds	45
Creating VLANs	46
VMware Platform	47
Creating Teams or Bonds	48
Creating VLANs	51

Introduction

This document provides design guidelines for implementing a unified data center fabric using Cisco Nexus[®] 5000 Series Switches and QLogic second-generation converged network adapters (CNAs). Unified fabric implies that multiple types of network traffic can be present on the same physical link: for example, the fabric may transport both LAN and SAN traffic simultaneously over 10 Gigabit Ethernet. To achieve this unified approach, Fibre Channel over Ethernet (FCoE) is used to transport storage traffic over an Ethernet network.

The Cisco Nexus 5000 Series offers the first unified fabric switches in the networking industry. They are field proven, with more than 1000 customers and more than a year in the market. These powerful 10 Gigabit Ethernet LAN switches can operate as fully functional SAN switches. Along with other products in the Cisco Nexus Family, the Cisco Nexus 5000 Series can provide a data center access layer solution for connectivity to Gigabit Ethernet, 10 Gigabit Ethernet, and Fibre Channel ports.

CNAs are multifunction adapters with the characteristics of traditional network interface cards (NICs) and Fibre Channel host bus adapters (HBAs). The QLogic QLE8100 family of CNAs is leading the second generation of CNAs with significant enhancements in performance, power utilization, and server-platform applicability. Together, QLE8100 adapters and the Cisco Nexus 5000 Series create a consolidated and optimized solution for server connectivity in the data center.

This design guide provides a brief protocol discussion as well as configuration guidelines for implementing a unified fabric with FCoE. For a more complete description of the FCoE protocol defined in the FC-BB-5 standard, please visit <http://www.t11.org/index.html>.

FCoE Protocol

FCoE provides a mechanism for transporting Fibre Channel frames on top of an Ethernet infrastructure. At a high level, unaltered Fibre Channel frames are encapsulated in an Ethernet header and sent along a lossless Ethernet fabric and decapsulated when they reach their target. Because no protocol conversion or state tables are required, FCoE is considered to be a gatewayless technology. The FCoE architecture is completely based on the Fibre Channel protocol. It provides the same host-to-switch and switch-to-switch connectivity as Fibre Channel fabrics. FCoE also provides the same level of management and security found in Fibre Channel today with the use of zoning, port world wide name (pWWN)-based port security, etc.

Because the transmission of Fibre Channel frames requires a lossless transport, an Ethernet network used to transport FCoE frames must also provide the same lossless characteristics found in today's Fibre Channel SANs. These characteristics are provided by using Priority Flow Control (PFC), a revised implementation of the IEEE 802.3X Ethernet standard known as Pause and currently passing through the standards body as IEEE 802.1Qbb.

PFC allows you to pause specific classes of service (CoSs) based on an IEEE 802.1p CoS value. FCoE traffic is marked with a specific CoS value that corresponds to a no-drop class. When congestion occurs in the network, a pause frame for the specified CoS value is sent to the server adapter, pausing the FCoE traffic. This approach is similar to the buffer-credit flow control used in Fibre Channel fabrics.

In native Fibre Channel, initiators and targets log into the domain and name servers in Fibre Channel networks to receive their Fibre Channel ID (FCID) so they can begin communicating on the fabric. In the same way, FCoE hosts log into an FCoE fabric through a switching element found in the Cisco Nexus 5000 Series Switches. This element is called the Fibre Channel forwarder (FCF). The FCF acts as the logical Fibre Channel switch for the end-host login requests. After hosts are logged into an FCF, all traffic must pass through that FCF to reach its target. The Cisco Nexus 5000 Series Switches can run in both switch mode and N-port virtualization (NPV) mode, which will be discussed in the solution sections of this guide.

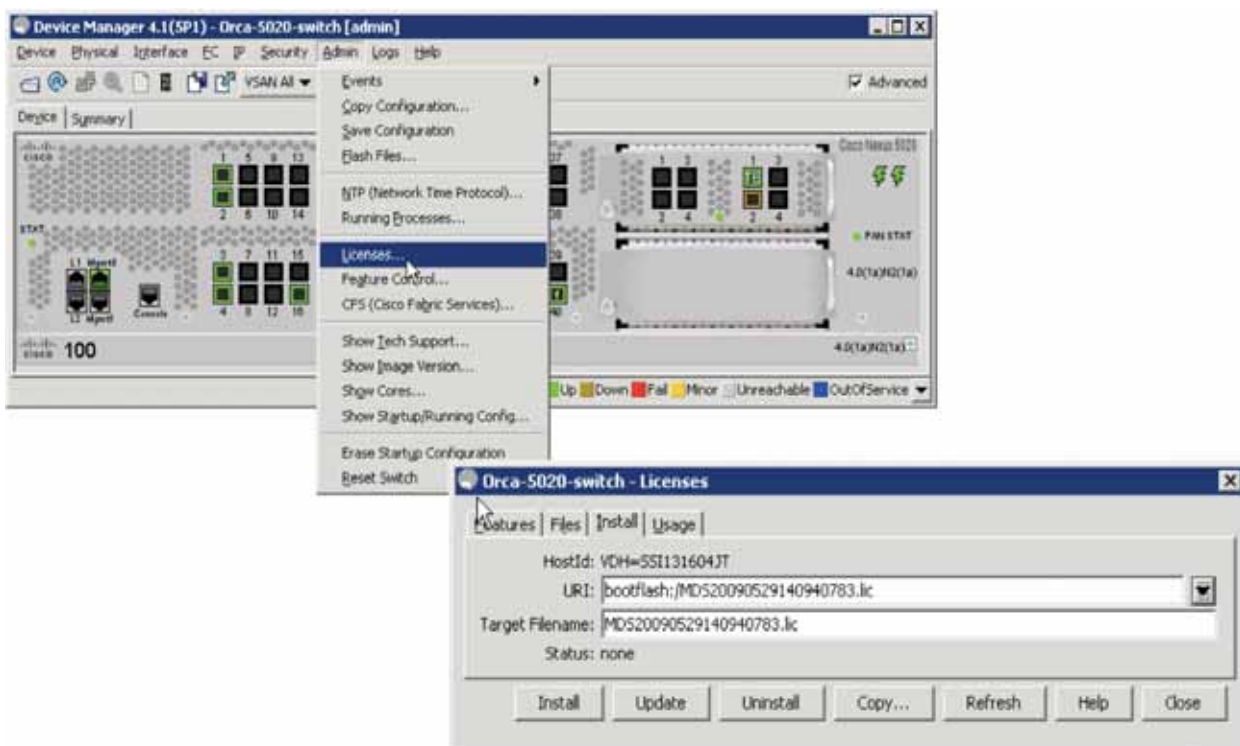
Implementing Baseline FCoE

To enable FCoE on the Cisco Nexus 5000 Series, an FCoE license is required. (An FCoE configuration can be applied without installing the license first by taking advantage of the 120-day license grace period.) After the license file is received, use FTP or TFTP to send the file to the Cisco Nexus 5000 Series Switch and place it in the bootflash directory. When the file transfer is complete, apply the following configuration:

```
switch# config terminal
switch(config) # install license bootflash:license_file.lic
```

To install the FCoE license through the GUI, follow these steps:

1. In the Device Manager, choose Admin > Licenses.
2. In the License dialog box, click the Install tab and enter the URL and filename in the text boxes.
3. Click Install.



After the license is activated, enable the FCoE feature.

4. In the Device Manager, choose Admin > Feature Control.

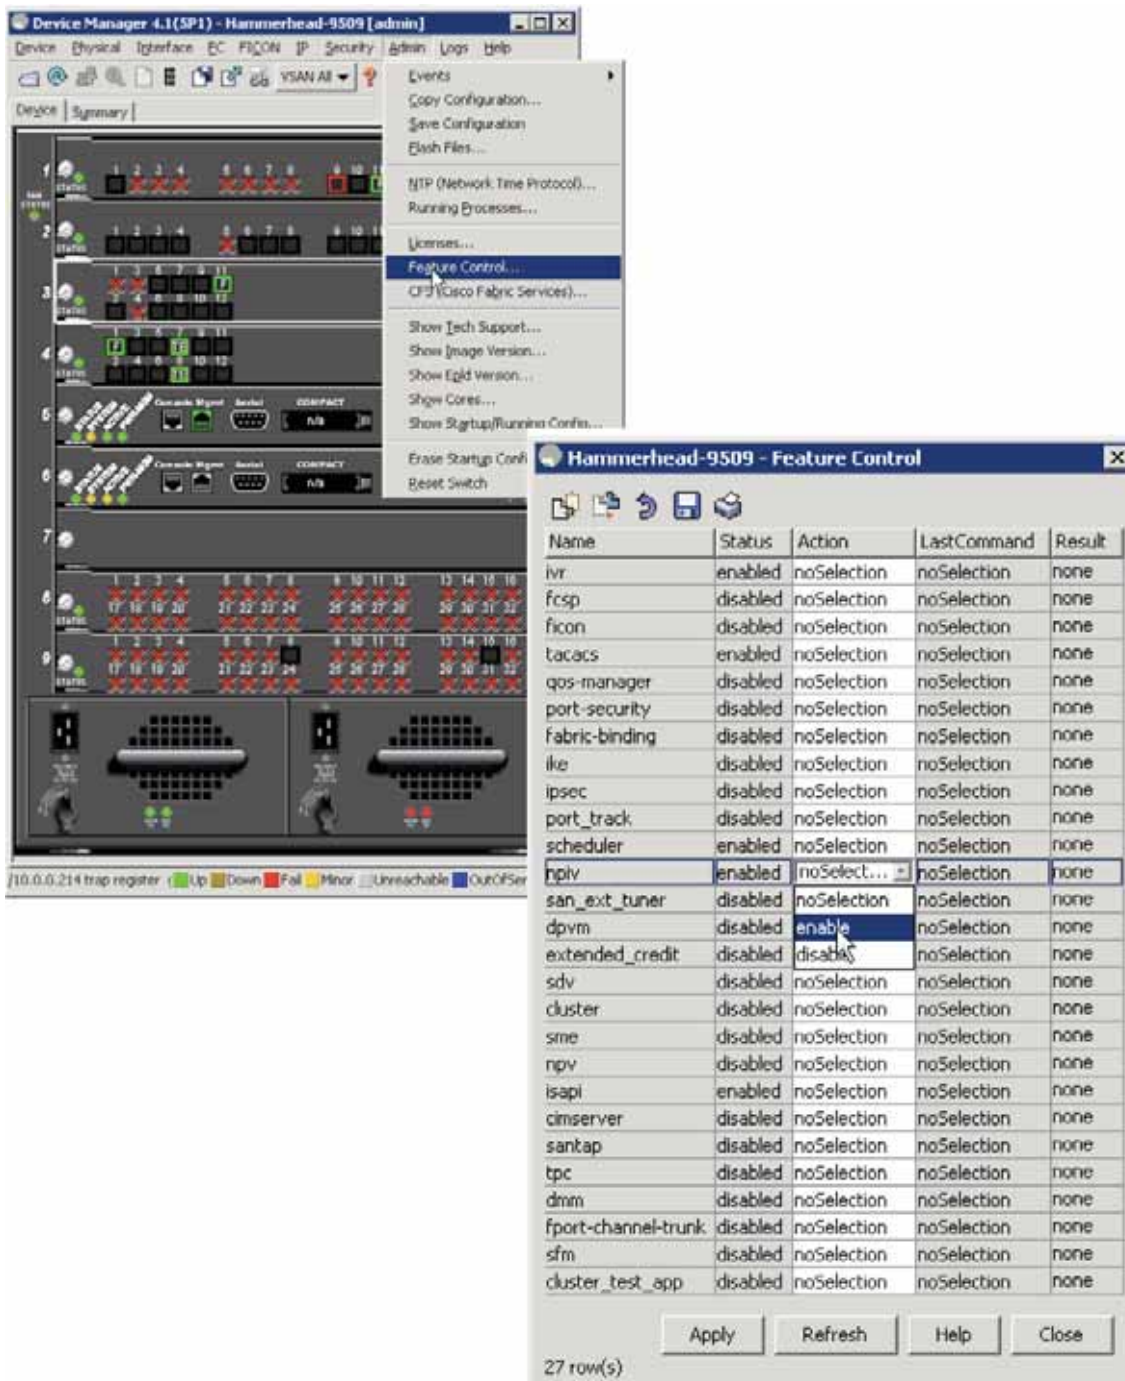
In the Feature Control dialog box, select FCoE and choose Action > enable

After the FCoE feature is enabled, save the configuration and reboot the Cisco Nexus 5000 Series Switch to activate the FCoE and native Fibre Channel ports. (This reboot is no longer necessary if the Cisco Nexus 5000 Series Switch is running Cisco® NX-O Software Release 4.1(3) N1 code or later). By default, the Fibre Channel switching element of the Cisco Nexus 5000 Series Switch will come up in switch mode:

```
switch #config terminal
switch(config) # feature fcoe
```

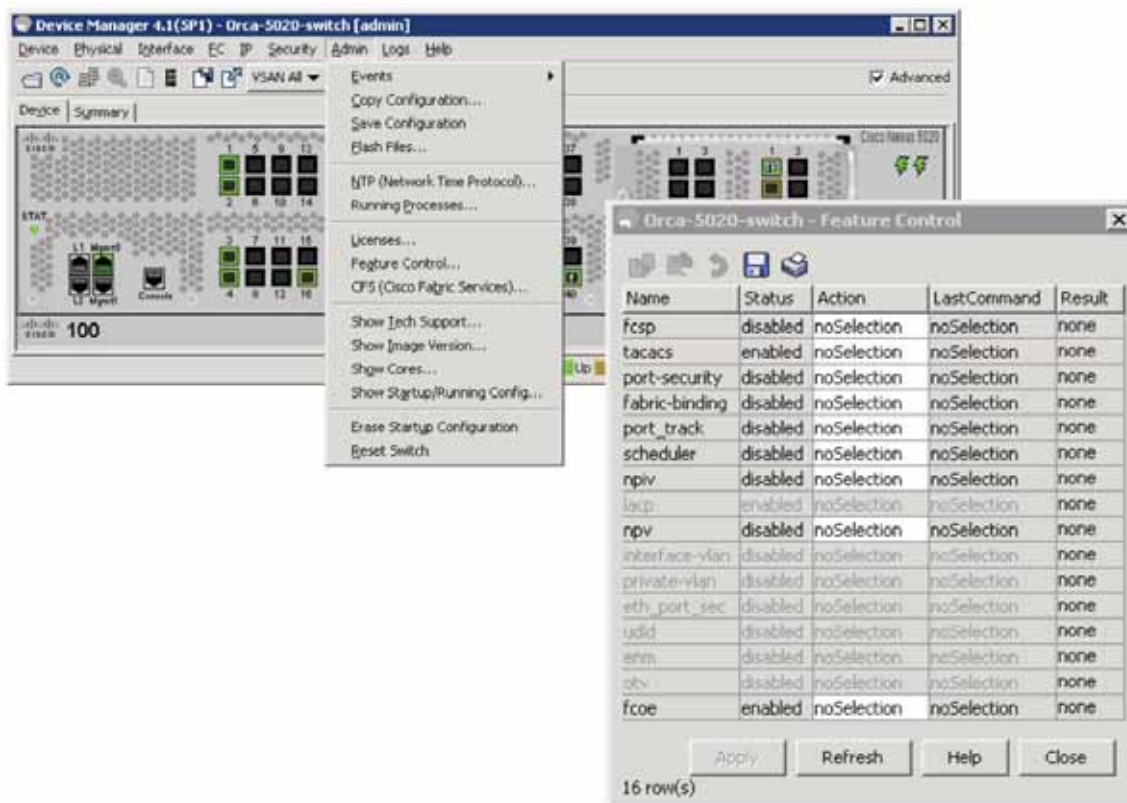

To use the GUI to enable NPIV on a core director, follow these steps:

1. In the Device Manager, choose Admin > Feature Control
2. In the Feature Control dialog box, select npiv and choose Action > enable.



When you set the Cisco Nexus 5000 Series Switch in NPV mode, the switch erases the configuration and needs to be rebooted. To enable NPV on the Cisco Nexus 5000 Series, enter:

```
switch# config terminal
switch(config)# npv enable
```

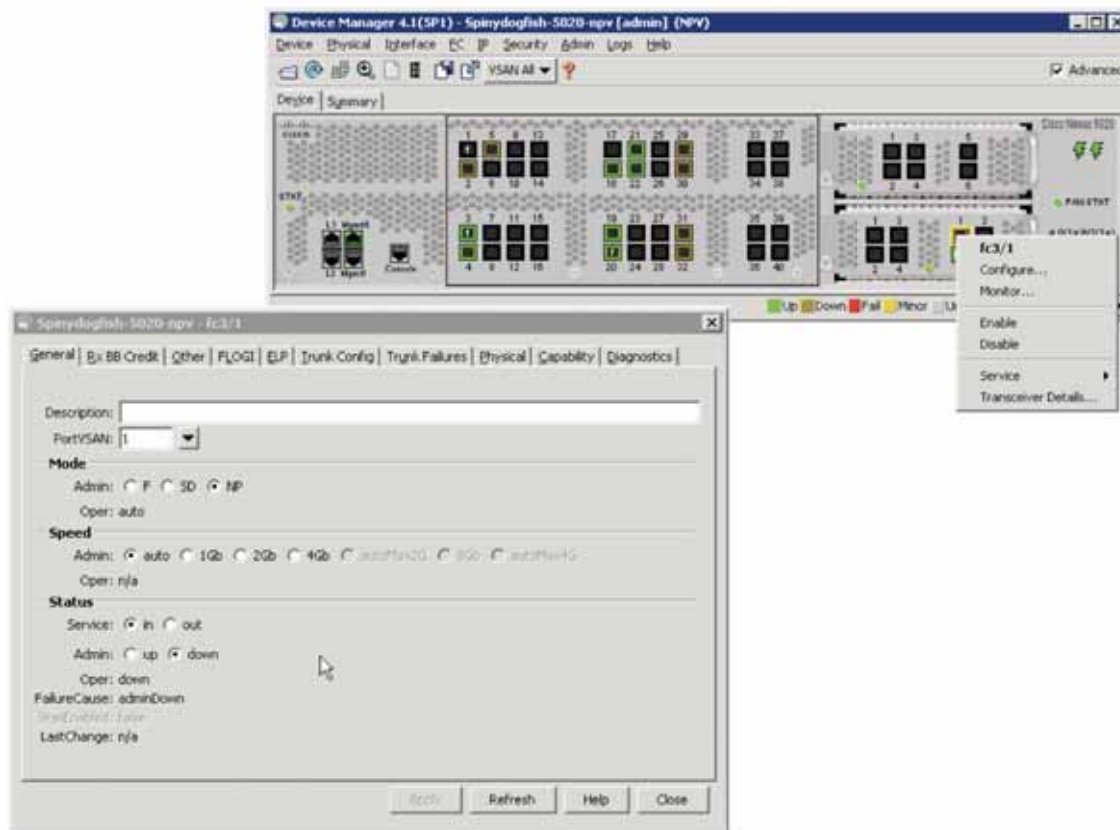


When the switch recovers from the reboot, configure the uplink ports of the core director in NP mode. To configure the uplink ports in NP mode, enter:

```
switch# config terminal
switch(config) # interface fc <slot>/<port>
switch(config-if)# switchport mode NP
```

To use the GUI to configure the Cisco Nexus 5000 Series Fibre Channel port as an NP uplink port, follow these steps:

1. In the Device Manager, on the Device tab, right-click the FC port, and choose Configure.
2. On the General tab, in the Mode section, select Admin: NP; in the Status section, select Admin: up.



If multiple NP uplink ports are connected to and configured for the core director, you have three options for configuring link selection for host logins and failover mechanisms. These are discussed in the following sections.

Dynamic Load Balancing

By default, the Cisco Nexus 5000 Series Switch will send fabric login messages down the available NP uplinks in a round-robin fashion. Should an NP uplink fail, hosts that were logged into the fabric through the failed NP uplink will need to log in again to the fabric and be load balanced across the other available NP uplinks

NPV Traffic Management with the NPV Traffic Map

The NPV traffic map associates one or more server-facing interfaces with an NP uplink interface. This feature allows the administrator to perform load balancing on a per-host basis. Here is the configuration:

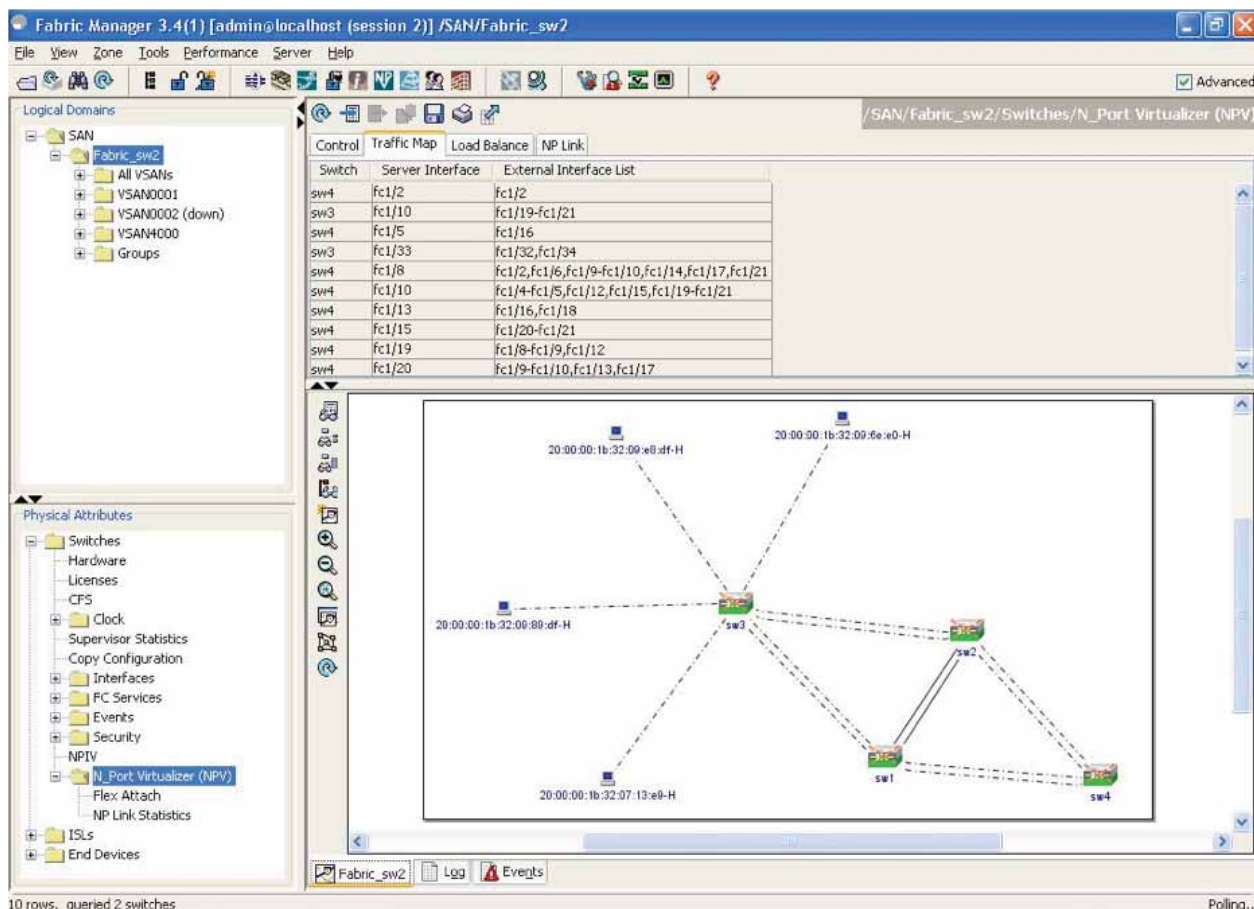
```
switch# config t
switch(config)# npv traffic-map server-interface [ fc (slot/port) or vfc
(slot/port)] external-interface fc slot/port
```

List of External Interfaces per Server Interface

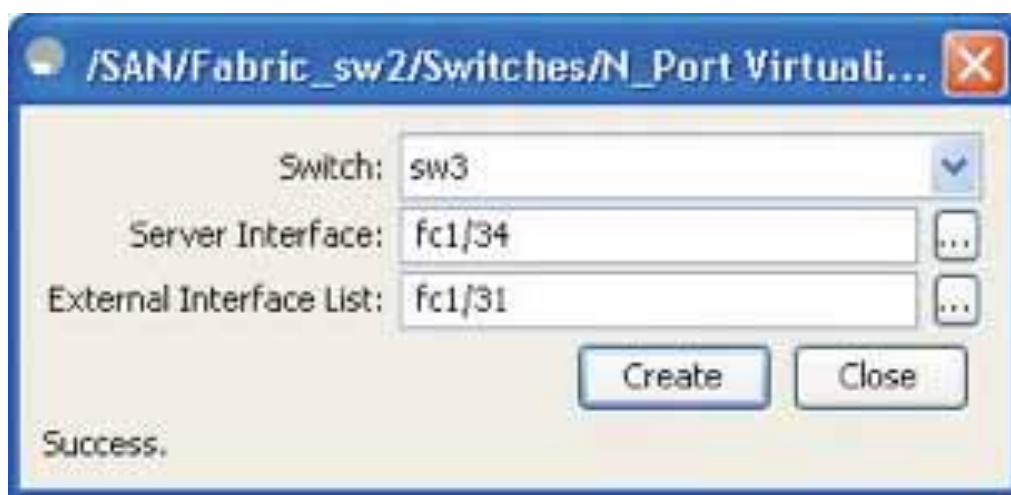
When the server interface is down, or when the specified external interface list includes the external interface already in use, a list of external interfaces linked to the server interfaces is used.

To configure the list of external interfaces per server interface using the GUI, follow these steps:

1. In the Fabric Manager, choose Physical Attributes > Switches > N_Port Virtualizer (NPV).



2. Click the Traffic Map tab.
3. Click the icon in the toolbar or right-click and then choose Create Row.
4. In the Map Entry dialog box, click the drop-down Switch list and choose the switch.



5. Type the port numbers or click the [...] button (not available on blade server switches) to select the server interface and external interfaces.



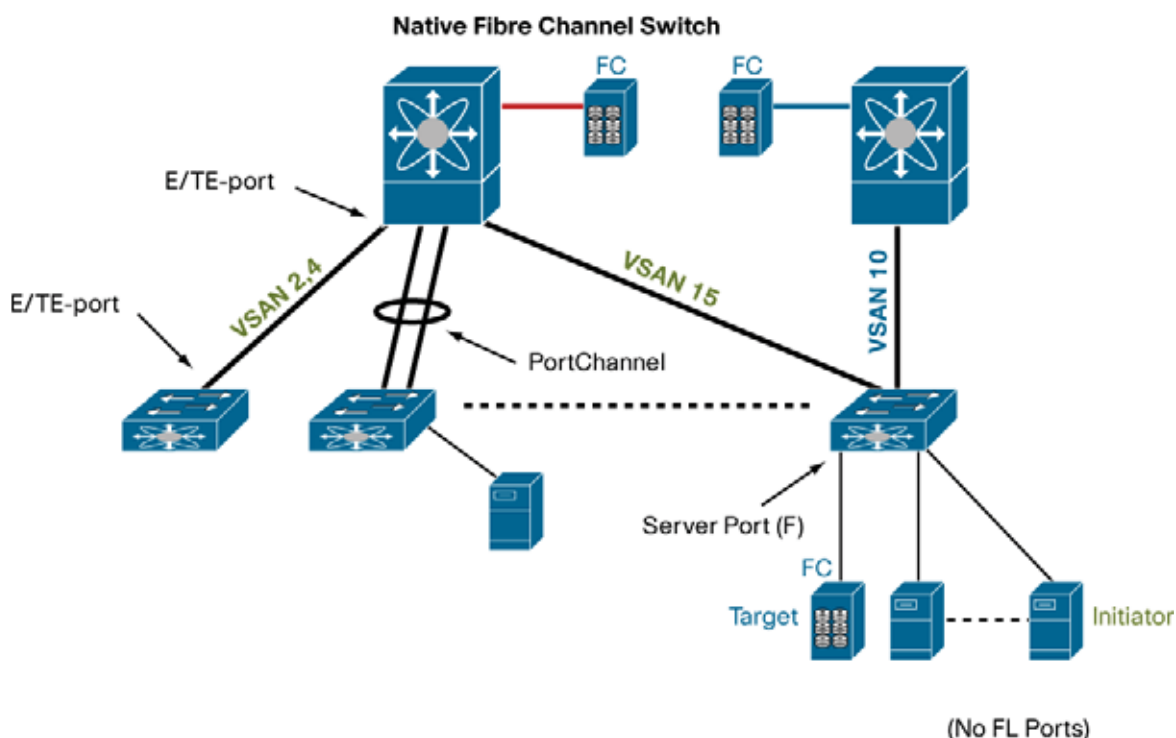
Note: Only one server interface can be selected at a time. Multiple external interfaces can be mapped to a single server interface. Previously selected ports are disabled and cannot be selected.

To delete the map entry, select the row from the Traffic Map tab and click the delete icon on the toolbar or right-click and choose Delete Row.

Solution Topology 2

Figure 2 shows the second solution topology.

Figure 2. Topology 2



Configuring the Cisco Nexus 5000 Series in Switch Mode

The Cisco Nexus 5000 Series Switch has two modes of operation: NPV mode and switch mode. When the device is in switch mode, the Cisco Nexus 5000 Series Switch provides standard Fibre Channel services and switching capability. Each Cisco Nexus 5000 Series Switch has a unique domain identifier on the Fibre Channel fabric, and as more switches are deployed in the fabric, the domain count grows. As the Cisco Nexus 5000 Series Switch attaches to the core Fibre Channel network, the ports on both the Cisco Nexus 5000 Series Switch and the core switch come up as expansion ports, or E-ports.

As mentioned previously, after the FCoE feature is enabled, the Cisco Nexus 5000 Series configuration will need to be saved, and the switch must be rebooted to activate the FCoE capabilities and native Fibre Channel ports. When the switch recovers from the reboot, it comes up in switch mode by default. At this point, you can connect the native Fibre Channel ports to a Fibre Channel switch. The ports will follow standard Fibre Channel switching and come up as E/TE ports.

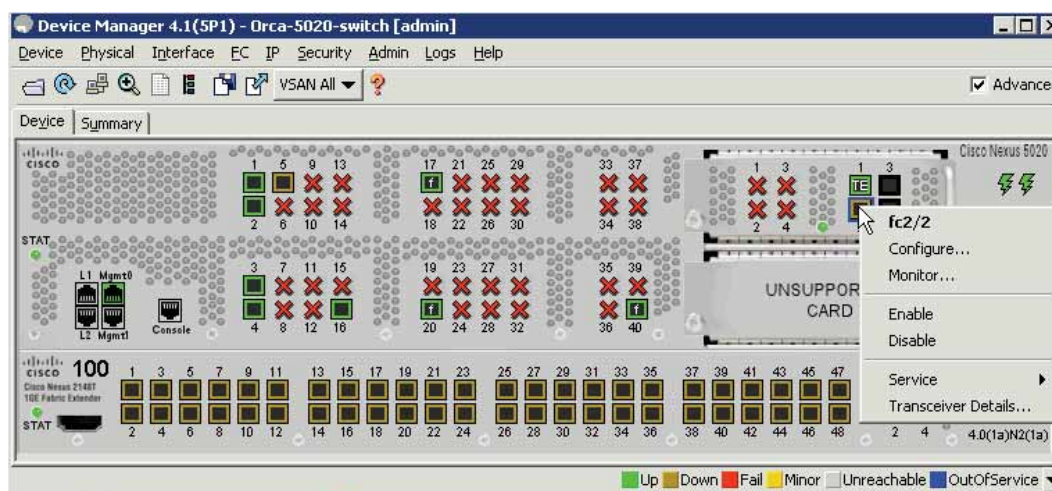
The ports on the Fibre Channel expansion module are autosensing for port types and should negotiate E-port connectivity with the core switch.

To manually configure the Fibre Channel ports on the expansion module, enter:


```
switch-5K# config terminal
switch-5K(config) interface fc <slot>/<port>
switch-5K (config-if)# switchport mode E
```

To configure the ports through the GUI, follow these steps:

1. Start the Device Manager and right-click the port to be configured.
2. In the menu that appears, choose Configure.



3. It is a best practice to hard code the Mode, Speed, and Status options for the ports when configuring ports in the fabric. Select Mode: E (for E-port), Speed: 4Gb (for 4 Gbps), and Status: Admin: up (for administratively up). Then click Apply to commit the configuration for the port.

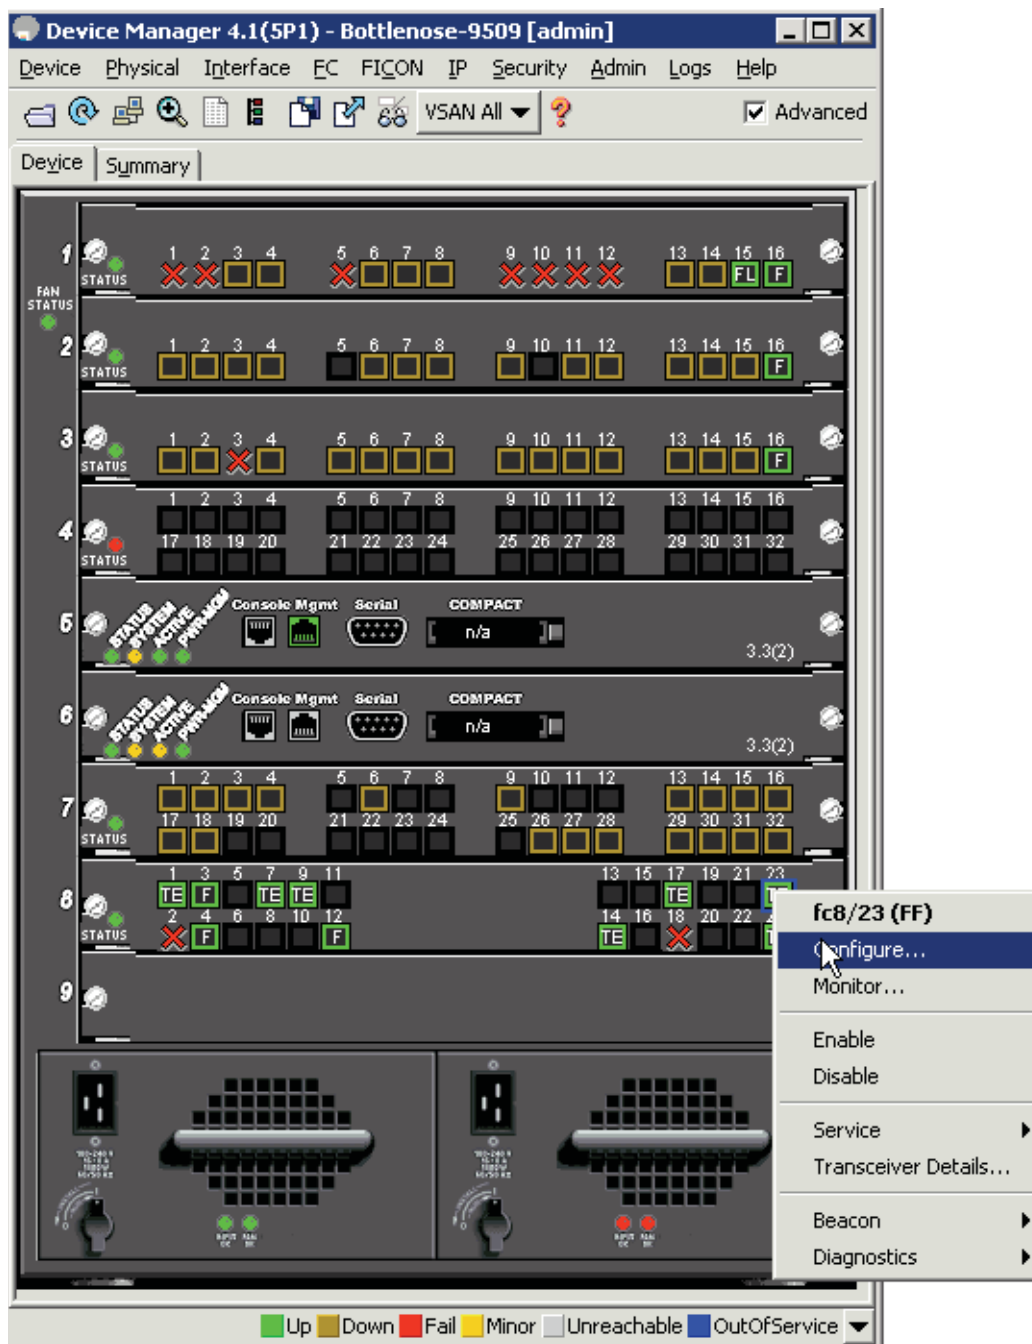


If the interfaces are in auto mode on the core director, the interface will come up in expansion or E-port mode:

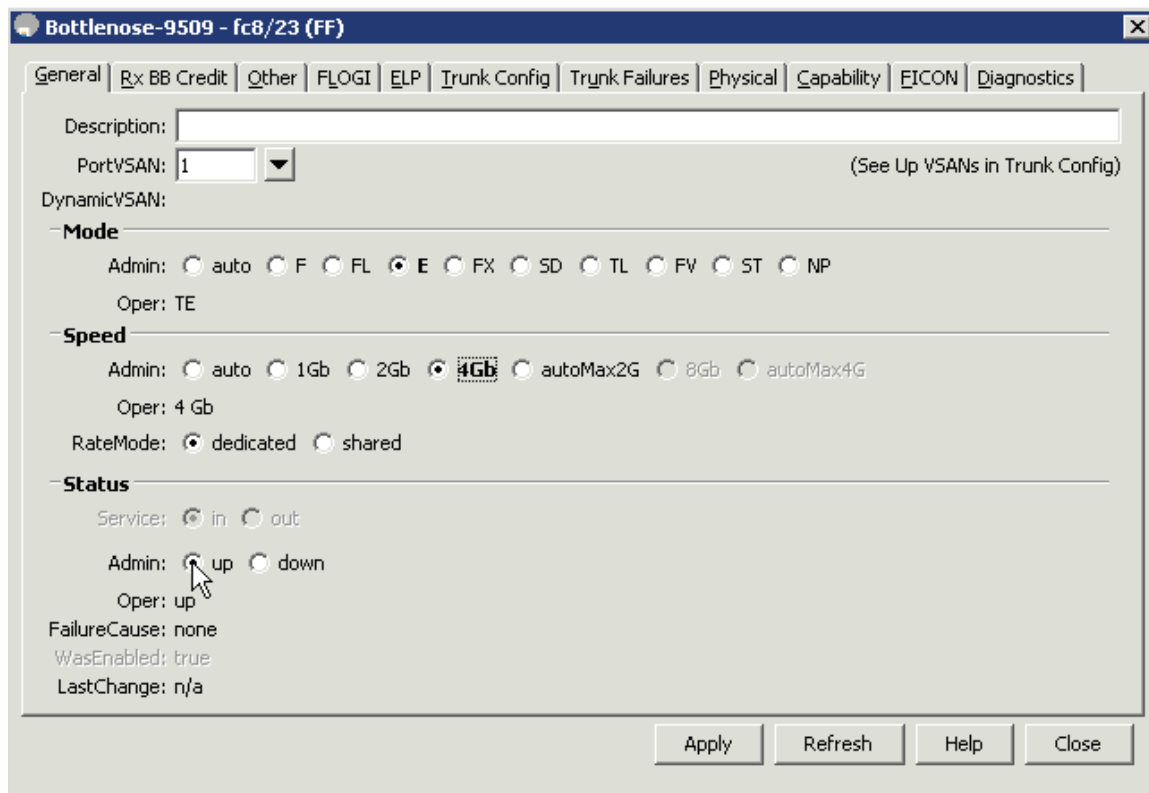
```
switch-core# config terminal
switch-core(config)# interface fc <slot>/<port>
switch-core(config-if)# switchport mode E
```

To configure the ports in the GUI, follow these steps:

1. In the Device Manager, right-click the port to be configured and choose Configure.



2. In the interface configuration dialog box select Mode: E (for E-port), Speed: 4Gb (for 4 Gbps), and Status: Admin: up (for administratively up). Then click Apply to commit the configuration for the port.



Bottlenose-9509 - fc8/23 (FF)

General | Rx BB Credit | Other | FLOGI | ELP | Trunk Config | Trunk Failures | Physical | Capability | FICON | Diagnostics

Description:

PortVSAN: 1 (See Up VSANs in Trunk Config)

DynamicVSAN:

Mode

Admin: auto F FL E FX SD TL FV ST NP

Oper: TE

Speed

Admin: auto 1Gb 2Gb 4Gb autoMax2G 8Gb autoMax4G

Oper: 4 Gb

RateMode: dedicated shared

Status

Service: in out

Admin: up down

Oper: up

FailureCause: none

WasEnabled: true

LastChange: n/a

Apply Refresh Help Close

When the Cisco Nexus 5000 Series Switch is operating in switch mode with FCoE and all uplink ports are connected to the same director, it is a best practice to bundle the links in a PortChannel. This approach enables high availability and better bandwidth utilization due to the exchange-based load balancing employed over a PortChannel.

```
switch-5k# config terminal
switch-5k(config)# interface san-port-channel 1
switch-5k(config-if)# exit
switch-5k(config)# interface fc 2/1
switch-5k(config-if)# switchport mode E
switch-5k(config-if)# channel-group 1
fc2/1 added to san-port-channel 1 and disabled
please do the same operation on the switch at the other end of the channel,
then do "no shutdown" at both ends to bring them up
switch-5k(config-if)# no shutdown
switch-5k(config-if)# interface fc 2/2
switch-5k(config-if)# switchport mode E
switch-5k(config-if)# channel-group 1
fc2/1 added to san-port-channel 1 and disabled
please do the same operation on the switch at the other end of the channel,
then do "no shutdown" at both ends to bring them up
switch-5k(config-if)# no shutdown
```

To configure the core switch uplink ports into a PortChannel on the Cisco MDS 9000 Family core director, enter:

```
switch-core# config terminal
switch-core(config)# interface port-channel 1
switch-core(config-if)# exit
```

```
switch-core(config)# interface fc 8/1
switch-core(config-if)# switchport rate-mode dedicated
switch-core(config-if)# switchport mode E
switch-core(config-if)# channel-group 1
fc3/1 added to port-channel 1 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both end to bring them up
switchport-core(config-if)#no shutdown

switch-core(config-if)# interface fc 4/1
switch-core(config-if)# switchport rate-mode dedicated
switch-core(config-if)# switchport mode E
switch-core(config-if)# channel-group 1
fc3/1 added to port-channel 1 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both end to bring them up
switchport-core(config-if)#no shutdown
switchport-core(config-if)# end
```

After the PortChannels are created, if more than one VSAN is being used on the Cisco Nexus 5000 Series Switch, trunking must be configured to allow the multiple VSANs to transverse the PortChannel. When configuring trunking, it is a best practice to allow only the required VSANs on the PortChannel rather than all VSANs. To add VSANs to a trunked interface, enter:

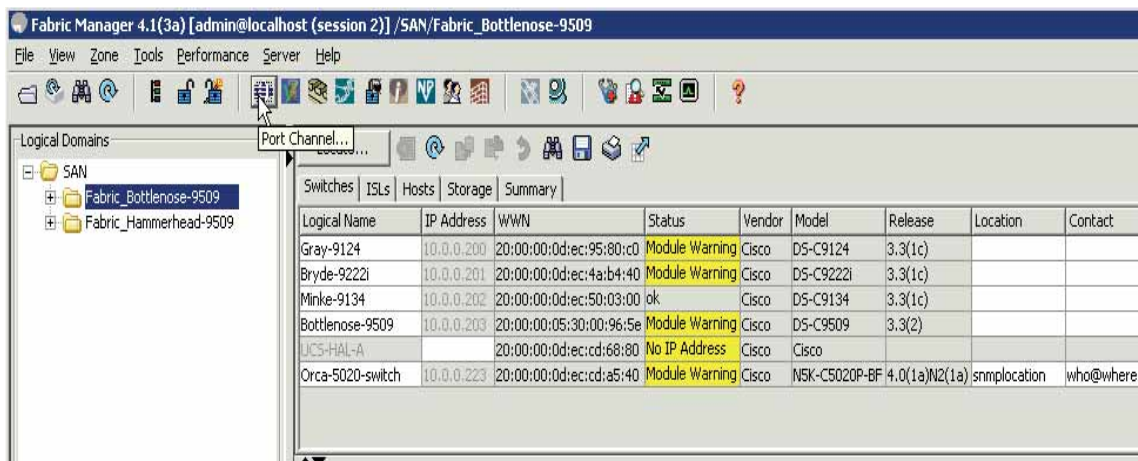
```
switch-5k# config terminal
switch-5k(config)# interface san-port-channel 1
switch-5k(config-if)# switchport trunk allowed vsan add 2
switch-5k(config-if)# switchport trunk allowed vsan add 4
switch-5k(config-if)# end
```

Configure the VSAN trunking parameters on the core switch so that the trunk is fully active:

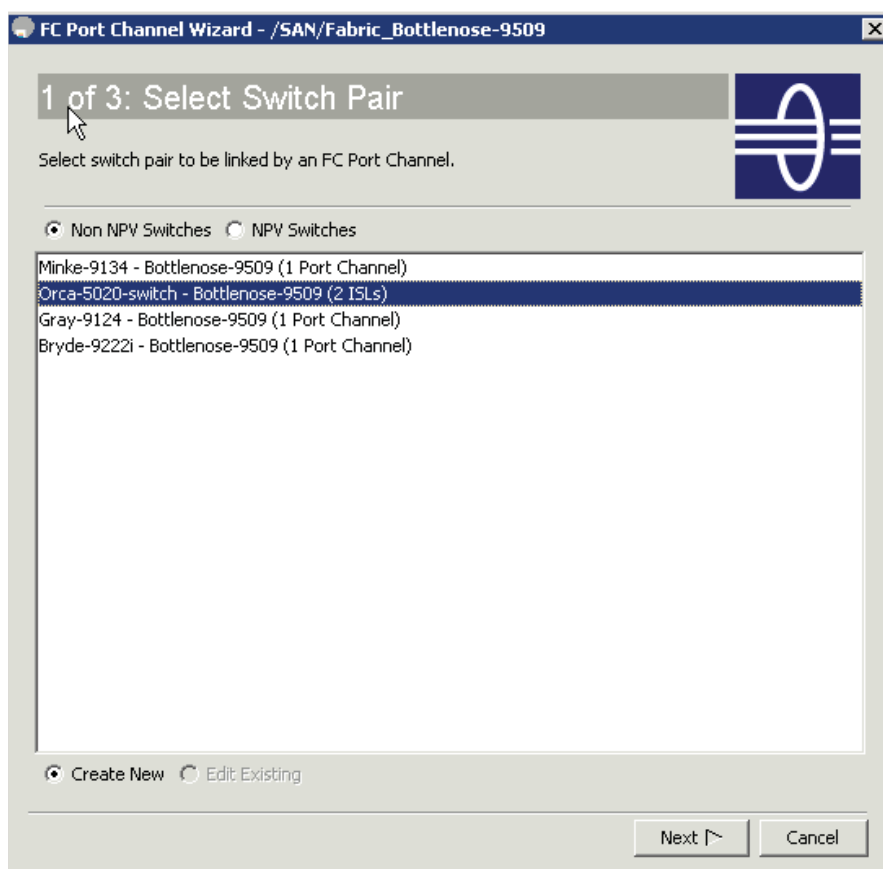
```
switch-core# config terminal
switch-core(config)# interface port-channel 1
switch-core(config-if)# switchport trunk allowed vsan add 2
switch-core(config-if)# switchport trunk allowed vsan add 4
switch-core(config-if)# end
```

To configure PortChannel interfaces, you can use the Device Manager on the two devices or use the Fabric Manager and the PortChannel wizard. To configure PortChannel interfaces, follow these steps:

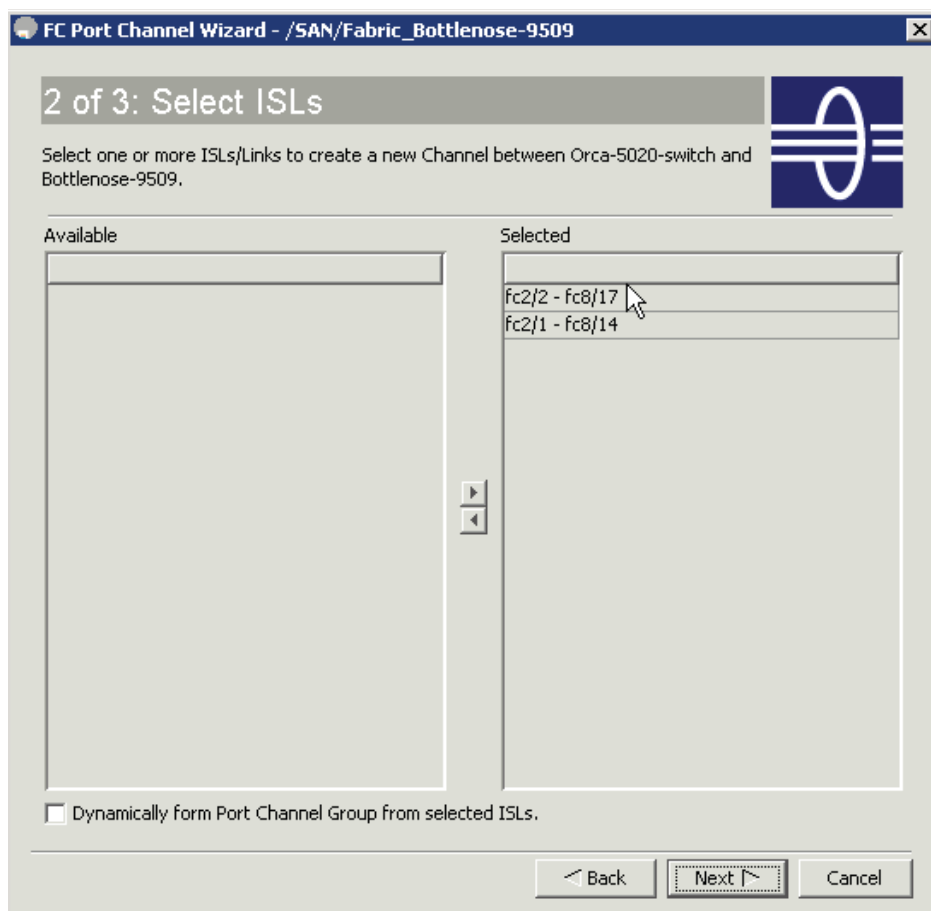
1. In the Fabric Manager, click the PortChannel wizard icon.



2. The first of three configuration dialog boxes will open. Select the two switches on which the PortChannel will be created. Then click Next.



3. Select the links between the switches that should be configured on the PortChannel interface. Then click Next.



4. Specify the following PortChannel parameters: Channel Id, Description, and Port VSAN. Select Trunk Mode: trunk and Speed: auto. Select the check box for "Force Admin, Trunk, Speed, and VSAN attributes to be identical" to force compatibility configuration on the links.

Configuring FCoE

Configuring VLAN-to-VSAN Mapping

When configuring an FCoE fabric, the first step is to create VLAN-to-VSAN mapping, which allows the logical Fibre Channel fabric to transverse the Ethernet network. It is a best practice to have isolated FCoE VLANs that are dedicated to FCoE traffic and separate from the rest of the native Ethernet VLANs. You also should not assign VLAN 1 or VSAN 1 to the FCoE network. Typically, that VLAN and VSAN are used for management traffic or for devices that have no other VLAN or VSAN assignment.

As a result of the practice of isolating FCoE traffic on its own VLAN, any host running FCoE needs to be configured as a trunk port rather than an access port. The host port on the Cisco Nexus 5000 Series Switch must also be configured as a trunk port (see “Configuring the Adapter Side” later in this document).

To configure the VLAN-to-VSAN mapping, enter:

```
switch# config terminal
switch(config)# vlan 2
switch(config-vlan)# fcoe vsan 2
switch(config-vlan)# end
```

Note: The FCoE VSAN must be configured and in the VSAN database of the Cisco Nexus 5000 Series Switch before it is mapped to a VLAN.

Configuring 10 Gigabit Ethernet Interfaces for FCoE

As stated earlier, the host-facing ports of the Cisco Nexus 5000 Series Switch need to be configured as trunk ports when carrying FCoE traffic. To configure trunking on these ports, enter:

```

switch# config terminal
switch(config)# interface Ethernet 1/1
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 1,2
switch(config-if)# end
  
```

Configuring Virtual Fibre Channel Interfaces

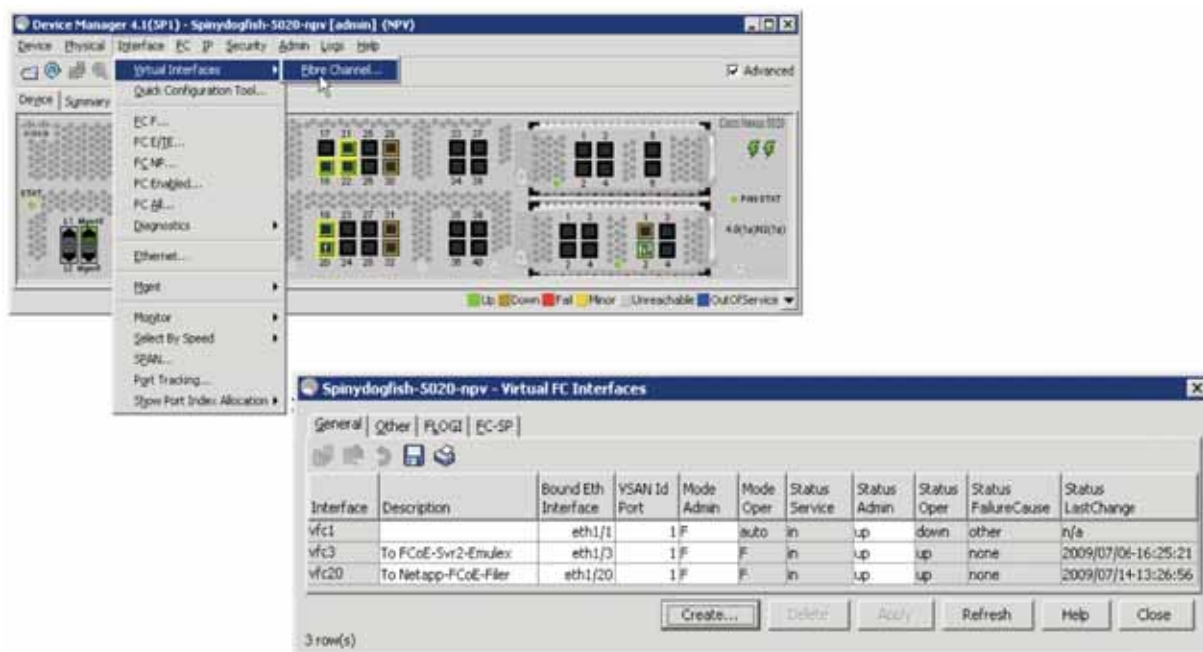
The Fibre Channel portion of FCoE is enabled on the host-facing ports by creating virtual Fibre Channel (vFC) interfaces and binding them to physical Ethernet ports on the Cisco Nexus 5000 Series Switch:

```


switch# config terminal
switch(config)# interface vfc 1
switch(config-if)# switchport mode F
switch(config-if)# bind interface Ethernet 1/1
switch(config-if)# exit
  
```

To configure the vFC interface through the GUI, follow these steps:

1. In the Device Manager, choose Virtual Interfaces > Fibre Channel.



2. In the Create Virtual FC Interfaces dialog box that opens, click Create to open the interface parameters.
3. Enter the vFC interface you are creating. Click the Bind Interface button to open the list of 10 Gigabit Ethernet ports to bind the vFC interface.
4. Click the Create button to create the vFC interface.



Assigning vFC Interfaces to the VSAN

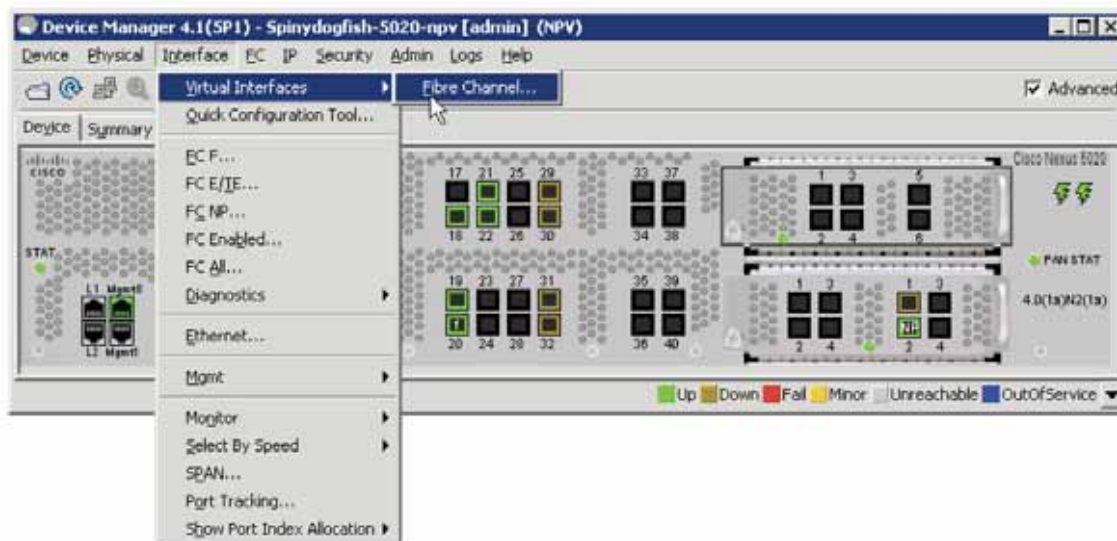
After the vFC interfaces are configured, they need to be assigned and configured as members of a VSAN. The VLAN-to-VSAN mapping must already be configured so that FCoE traffic can pass from the initiator up to the Cisco Nexus 5000 Series Switch.

To configure the mapping from the command line, enter:

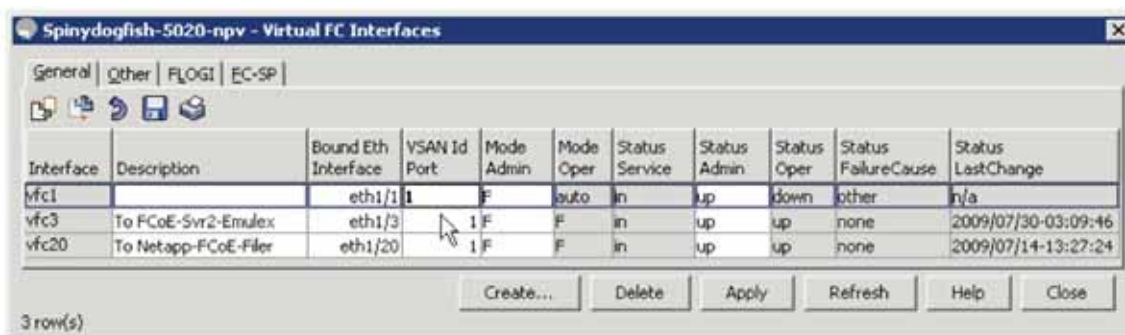
```
switch# config t
switch(config)# vsan database
switch(config-vsan)# vsan 2 interface vfc 1
switch(config-vsan)# end
```

To add the vFC interfaces to the correct VSANs using the GUI, follow these steps:

1. In the Device Manager, choose Interface > Virtual Interfaces > Fibre Channel.



2. In the Virtual FC Interfaces dialog box that appears, select the VSAN Id Port column and enter the VSAN to which the vFC interface should be configured.



Configuring Management and Access Control

Configuring Role-Based Access

Role-based access control (RBAC) provides the necessary framework to create and deploy roles for different users. User roles are defined by rules that specify the access permissions each person assigned to that role is allowed. Each user role can contain multiple rules, and each user can belong to more than one role. For example, if role1 allows access only to configuration operations, and role2 allows access only to debug operations, then users who belong to both role1 and role2 can access both configuration and debug operations.

User roles can be defined to limit the switch resources that the user can access as well as to limit more granular access to entities such as interfaces, VLANs, and VSANs.

Rules are the basic element of a role. A rule specifies the operations a user is permitted to perform. Rules can be assigned on a command, feature, or feature-group basis. To learn more about the configuration parameters, please refer the system management section of the Cisco Nexus 5000 Series Switches configuration guide:

http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli/sec_rbac.html

Configuring TACACS+ Authentication

The Cisco Nexus 5000 Series supports both the RADIUS and TACACS+ protocols. It also supports authentication, authorization, and accounting (AAA) services to verify the identity of, grant access to, and track the actions of engineers who manage the Cisco Nexus 5000 Series fabric.

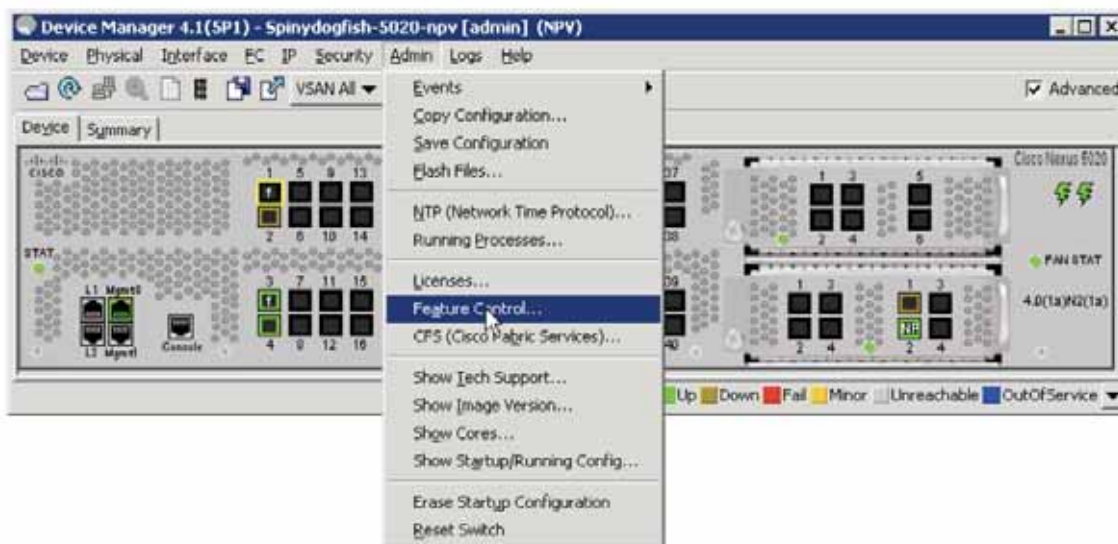
The TACACS+ or RADIUS server authenticates the user and returns to the switch the correct role in which the user should be placed. The user will then be limited to the functions and features contained in that role.

The TACACS+ feature is disabled by default. To enable TACACS+ from the CLI, enter:

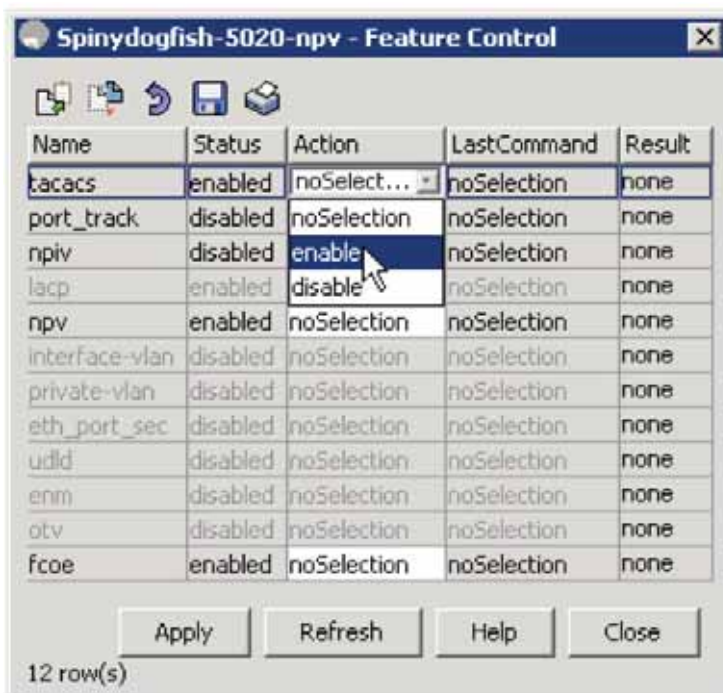
```
switch# config terminal
switch(config)# feature tacacs+
switch(config)# end
```

To enable the feature through the GUI, follow these steps:

1. In the Device Manager, choose Admin > Feature Control.



2. In the Feature Control dialog box that appears, select TACACS and choose Action > enable to enable the TACACS+ feature.



The next step in enabling TACACS+ for authentication is to configure the servers that are running the TACACS+ protocol along with the preshared key. To enable TACACS+ for authentication, enter:

```
switch# config t
switch(config)# tacacs-server host X.X.X.X key 7 cisco123
switch(config)# aaa group server tacacs+ TacServer
switch(config-tacacs)# server X.X.X.X
switch(config-tacacs)# use-vrf management
```

After configuring the server and the server groups, configure the applications that you want to have TACACS+ authenticate. To authenticate all applications, enter:

```
switch# config t
switch(config)# aaa authentication login default group <group name>
```

To configure TACACS+ through the GUI, follow these steps:

1. In the Device Manager, choose Security > AAA.
2. In the TACACS+ server configuration dialog box that opens, click the Create button to create a new TACACS+ server.



- Enter the IP address or the DNS name of the server and configure the type of security preshared key to be used. You can also enter a test user and password for TACACS+ server health checks. When you are done, click Create.

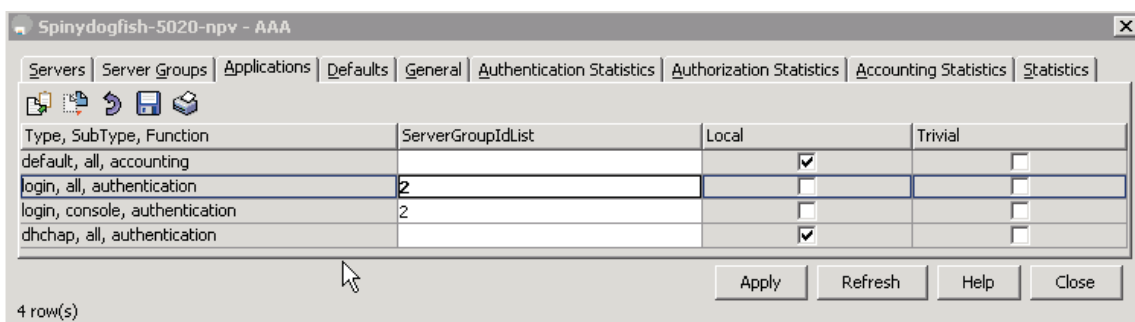
- In the AAA dialog box that appears after the server is configured, select the Server Groups tab to create the server group. Then click Create.

Id	Name	Protocol	ServerIdList	DeadTime
1	radius	radius		0
2	NY-DC-ACS	tacacs+	10.0.0.27	0
3	tacacs	tacacs+		0
4	Test	tacacs+		0

- In the Create AAA Group dialog box that opens, enter in the name of the group, the type of server, and the server list. Then click Create to create the server group.



6. Configure the applications that need authentication. You should authenticate everything through TACACS+. In the main AAA configuration dialog box, select the Applications tab.
7. Enter the server group ID that was created under the group configuration. Then click Apply and then Close.



Configuring TACACS+ Accounting


In a unified fabric environment in which you have two different technology groups potentially managing the Cisco Nexus 5000 Series Switch, TACACS+ accounting should be configured for user command tracking purposes. To enable TACACS+ accounting from the command line, enter:

```
switch# config t
switch(config)# aaa accounting default group <group-name>
```

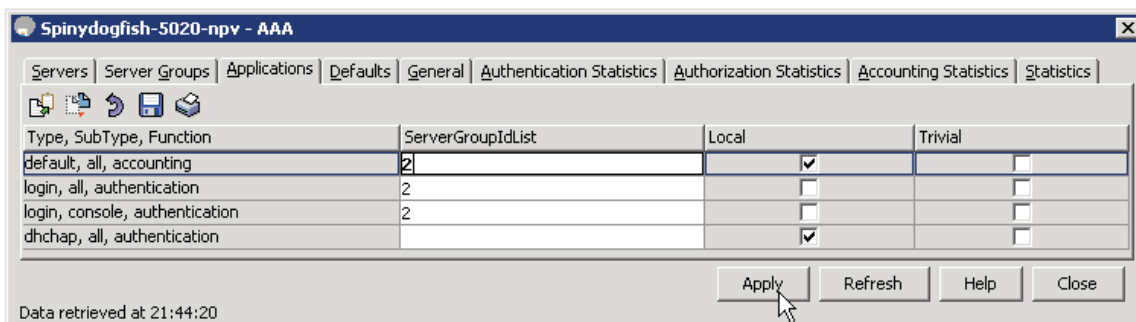
Where the group name matches the TACACS+ server group configured under the AAA group configuration.

To configure AAA accounting through the GUI, follow these steps:

1. In the Device Manager, choose Security > AAA.



- In the AAA configuration dialog box that appears, click the Applications tab and enter the server group ID that matches the TACACS+ server group configured under the AAA group configuration. Click Apply and then click Close.



Configuring the QLogic 10 Gigabit Enhanced Ethernet CNA

The QLogic 8100 Series CNAs are low-profile, 10 Gigabit Enhanced Ethernet PCI Express products that are available in single- and dual-port versions. The adapters use a single integrated QLogic ASIC that offers line-rate performance for both storage and data networking traffic and reduces CPU utilization through hardware offload for FCoE protocol processing (Figure 3).

Figure 3. QLogic QLE8152 Dual-Port CNA



The QLogic 8100 Series is available for channel and original equipment manufacturer (OEM) customers in six SKUs as described in Table 1, identified by transceiver type (copper cable, short-reach [SR] optical, and long-reach [LR] optical) and number of ports.

Table 1. 8100 Series Technical Specifications

SKU	QLE8150-CU QLE8152-CU	QLE8140-SR QLE8142-SR	QLE8140-LR QLE8142-LR
Transceiver type	Ships with empty Small Form-Factor Pluggable Plus (SFP+) enclosures Intended for use with active and passive copper cables, which are sold separately by FCoE switch vendors	SR optical	LR optical
Maximum cable reach	5m A list of supported cables is posted at QLogic's public website in the products section	300m	10 km
Host Bus Connectivity			
PCI Express lane support	PCIe Gen2x4 or Gen1x8		
Hardware platform	IA32 (x86), Intel64 and AMD64 (x64), IA64, SPARC, and PowerPC		
Ethernet Specifications			
Ethernet throughput	10 Gbps		
Autonegotiate 10000 MB, 100 MB, and 10Mb	No		
Ethernet frame size	1500 or 9000 bytes (jumbo frames)		
Stateless offload	IP, TCP, and User Datagram Protocol (UDP) checksum offloads; large send offload (LSO); giant send offload(GSO); receive-side scaling (RSS); and header-data split (I/OAT) support		
Enhanced Ethernet	PFC (IEEE 802.1Qbb rev 0), Enhanced Transmission Selection (ETS; IEEE 802.1Qaz rev 0), and Data Center Bridging Exchange (DCBX) Protocol (IEEE 802.1Qaz rev 0)		
Fibre Channel Specifications			
Fibre Channel throughput	10 Gbps		
I/O operations per second (IOPS)	250,000 per port		
Class of service	Class 3		
Protocols	SCSI-FCP and FC-TAPE		
Tools and Utilities			
QLogic SANsurfer Fibre Channel and FCoE GUI and CLI, and SANsurfer CNA Networking CLI and native OS networking tools			
LAN and SAN boot (PXE, UEFI, and FCode)			

Configuring the Adapter Side

To configure the adapter side, do the following:

- Download and install drivers for the appropriate platform from the QLogic website. Install both FCoE and Ethernet (Networking)¹ drivers.
- Download and install the QLogic SANsurfer Management Suite (SMS) bundle from the website for the given platform. Be sure to install the QLogic SMS FCoE GUI and command-line interface (CLI) along with the QLogic SMS CNA Networking CLI. The QLogic SMS FCoE GUI and CLI can be used to configure the FCoE interface; to configure the networking interface, use either native platform tools or the QLogic SMS CNA Networking CLI. The QLogic SMS CNA Networking CLI is currently supported on Microsoft Windows platforms only.
- Make sure that the cable (copper or fiber) is connected from the CNA port to the Cisco Nexus Family switch port.

¹ For the Ethernet (Networking) interface utilities to work, the FCoE driver must be installed.

- When connecting using optical fiber, make sure that the SFPs associated with the adapter ports have been qualified for QLogic. Installing a nonqualified SFP in a CNA port will cause the link to fail.
- Be sure that the Cisco Nexus Family switch has been updated with the latest firmware and that the switch port has been configured as follows:
 - The FCoE feature is enabled on the port or across the switch.
 - Switch Fibre Channel ports are configured.
 - The VSAN and VLAN for FCoE and Ethernet have been configured.
 - Ethernet Interfaces and vFC interfaces have been configured.

For switch-side configuration, please refer to the section “Configuring FCoE” earlier in this document.

After the switch side has been configured, the next step is to bring up the FCoE link. For QLogic 8100 Series CNAs, all DCBX-related operational parameters, including ETS, PFC, VLAN, and CoS, are exchanged with the switch. No configuration is required for DCBX parameters at the adapter end as long as the CNA has the correct version of the firmware², drivers are loaded, and a cable connection exists between the adapter and switch.

The QLogic SANsurfer Fibre Channel and FCoE management suite is supported on Microsoft Windows as well as Linux platforms. The instructions that follow for bringing up the FCoE link on Windows platforms also apply to Linux platforms.

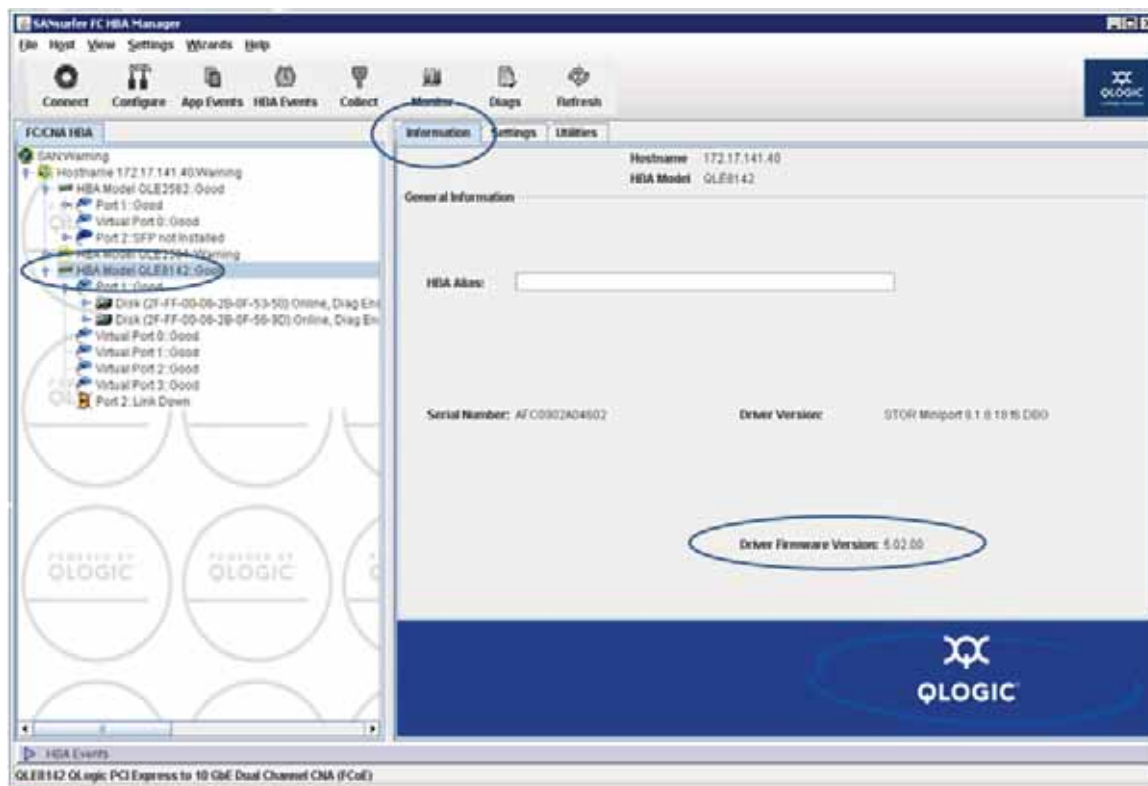
Bringing Up the FCoE Link

To bring up the FCoE link, follow these steps:

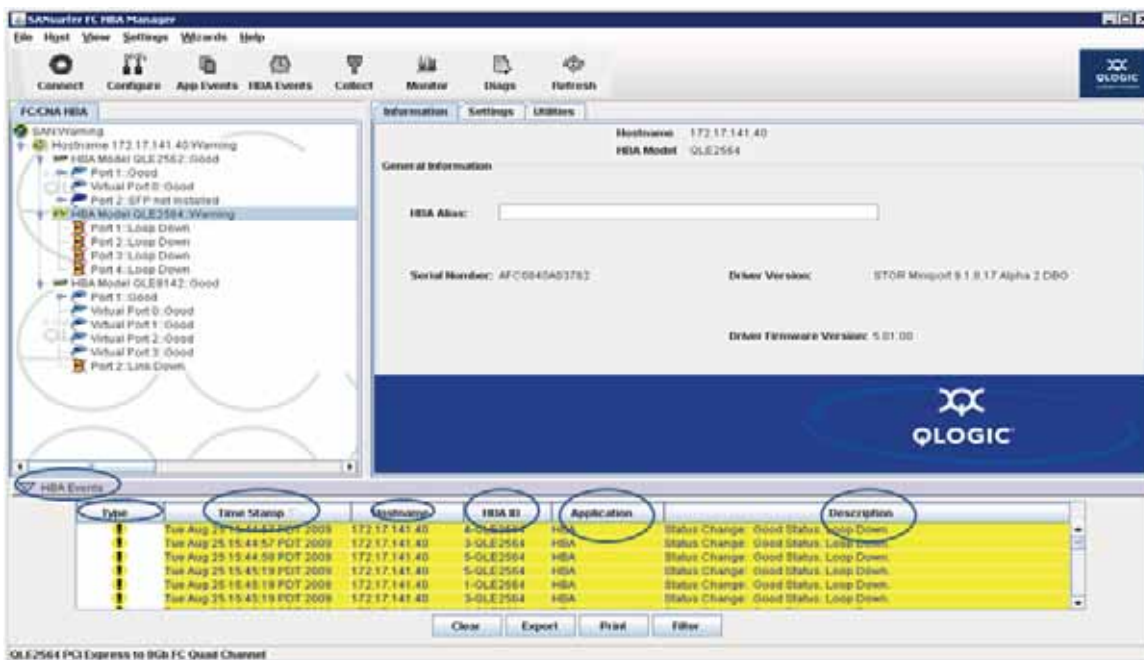
1. Invoke the QLogic SANsurfer Fibre Channel HBA Manager³; the manager is installed in the QLogic Management Suite directory.

² The minimum version of firmware on the CNA should be Version 5.01.03.

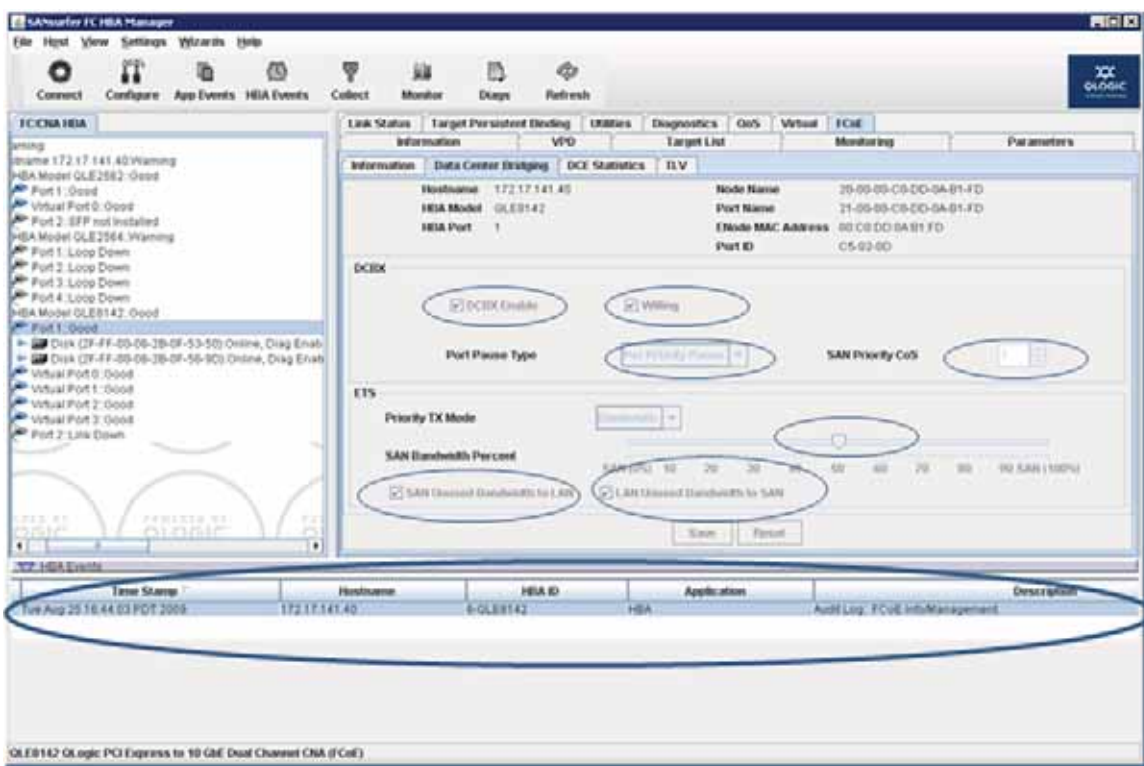
³ Make sure that the following minimum versions are used: QLogic SANsurfer Fibre Channel HBA Manager 5.0.1 build 57, SANsurfer CLI version 1.7.3 build 14, and SANsurfer CNA Networking CLI version 1.0.00.31 build 2.



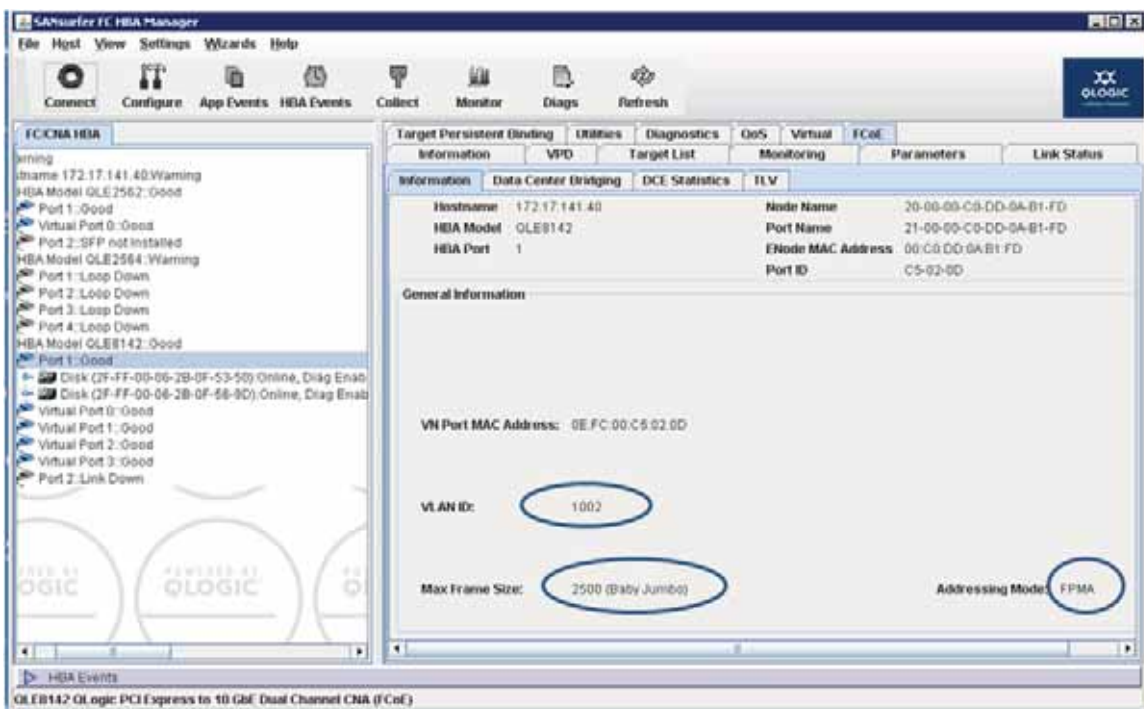
2. SANsurfer by default discovers all the QLogic Fibre Channel adapters and CNAs locally as shown in the screen image. To discover remote adapters automatically, select the Enable Auto Connect option when you launch QLogic SANsurfer. For remote discovery, click Connect and in the screen that appears, enter the IP address of the remote host. To discover and display all the adapters at the remote end, you need to install the QLogic QLremote agent on the remote hosts; otherwise, the QLogic SANsurfer Fibre Channel HBA Manager will not discover the given adapter on the host.
3. Open the HBA Events menu to get all CNA- and port-related events. The events are information regarding any issue with the CNA or port. For instance, if a Fibre Channel port is down, the status of the port is reported in the HBA Events pane. All severe events are shown with a red background, and all warnings are shown with a yellow background. If the switch is configured correctly, the QLogic SANsurfer Fibre Channel HBA Manager should not display any severe events for the given port.




- Click the FCoE tab and review the Data Center Bridging pane. It should display the values shown in the following screen image. DCBX-related parameters such as ETS, PFC, and SAN Priority CoS, are exchanged from the switch, and the exchanged values are displayed in the Data Center Bridging pane. For the QLogic 8100 Series CNAs, these values cannot be configured with management tools. Instead, the switch will push all the configuration data with DCBX Protocol.



- Open the information pane of the FCoE tab. It should display the values shown in the following screen image; be especially sure that the Addressing Mode field is specified correctly. The default values are VLAN ID: 1002, Max Frame Size: 2500 (Baby Jumbo), and Addressing Mode: FPMA.



- Verify that the FCoE link is now up. In QLogic SANsurfer, the status of the corresponding port should be displayed as Good, and the **show interface** command for the corresponding vFC port should display an FCID associated with the port.



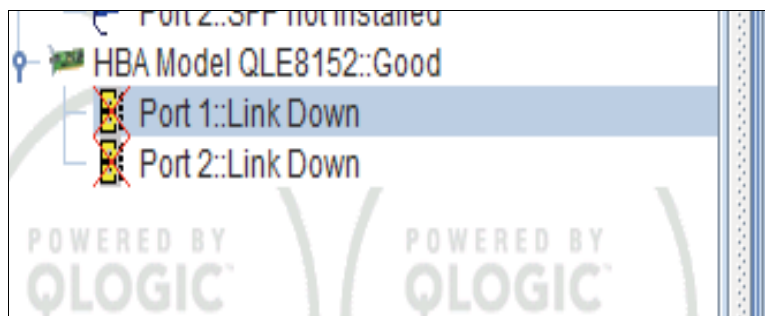
```

PAE(config)# show interface vfcl
vfcl is up
  Bound interface is Ethernet1/1
  Hardware is GigabitEthernet
  Port WWN is 20:00:00:0d:ec:a4:90:3f
  Admin port mode is F
  snmp link state traps are enabled
  Port mode is F, FCID is 0x0f0005
  Port vsan is 2
  Beacon is turned unknown
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  7688183503 frames input, 10079822721920 bytes
    0 discards, 0 errors
  6517613903 frames output, 5035781001484 bytes
    0 discards, 0 errors

```

Debugging

If the FCoE link is not up, QLogic SANsurfer will display the corresponding status as Link Down, along with warning events. Additional diagnostics may be needed to identify root cause.



Type	Time Stamp	Hostname	HBA ID	Application	Description
!	Mon Sep 14 14:27:00 PDT 2...	win-mh76cg8g...	2-QLE8152	HBA	Status Change: Good Status. Loop Down.
!	Mon Sep 14 14:27:18 PDT 2...	win-mh76cg8g...	2-QLE8152	HBA	Status Change: Good Status. Loop Down.
!	Mon Sep 14 14:27:22 PDT 2...	win-mh76cg8g...	2-QLE8152	HBA	Status Change: Good Status. Loop Down.
!	Mon Sep 14 14:27:38 PDT 2...	win-mh76cg8g...	2-QLE8152	HBA	Status Change: Good Status. Loop Down.
!	Mon Sep 14 14:27:42 PDT 2...	win-mh76cg8g...	2-QLE8152	HBA	Status Change: Good Status. Loop Down.
!	Mon Sep 14 14:27:59 PDT 2...	win-mh76cg8g...	2-QLE8152	HBA	Status Change: Good Status. Loop Down.

The adapter and switch exchange various states before declaring the FCoE link on each side. By tracing the state transition between the adapter and the switch, it is possible to identify the state that prevented the link from coming up.

Since all the states are managed at the Cisco Nexus Family switch level, run the commands that follow at the switch level. A current state of **FCOE_MGR_ETH_ST_UP** means that DCBX was exchanged successfully.

If DCBX exchange was not successful, the problem may be a firewall mismatch between the adapter and the switch. In this case, the current state is set to **FCOE_MGR_ETH_ST_INIT**, and no progress can be made with respect to state transition. For this kind of scenario, please contact Cisco or QLogic support personnel.

Run the following command on the switch for a given port interface to obtain the current state of the port along with the history associated with state transition.

```
show platform software fcoe_mgr info int...

Orca-5020-switch# show platform software fcoe_mgr info interface eth 1/40

Eth1/40(0x81f3434), if_index: 0x1a027000, RID Eth1/40
FSM current state: FCOE_MGR_ETH_ST_UP
PSS Runtime Config:-
PSS Runtime Data:-
  Oper State: up
  VFC IF Index: vfc3
```



```
FCOE on ? : TRUE
LLS supported ? : TRUE

>>>>FSM: <Ethernet1/40> has 14 logged transitions<<<<<

1) FSM:<Ethernet1/40> Transition at 805974 usecs after Wed Sep 2 02:37:00 2009
  Previous state: [FCOE_MGR_ETH_ST_INIT]
  Triggered event: [FCOE_MGR_ETH_EV_IF_CREATED]
  Next state: [FCOE_MGR_ETH_ST_DOWN]

2) FSM:<Ethernet1/40> Transition at 296126 usecs after Wed Sep 2 02:37:10 2009
  Previous state: [FCOE_MGR_ETH_ST_DOWN]
  Triggered event: [FCOE_MGR_ETH_EV_FCOE_CHANGE]
  Next state: [FSM_ST_NO_CHANGE]

3) FSM:<Ethernet1/40> Transition at 345753 usecs after Wed Sep 2 02:37:10 2009
  Previous state: [FCOE_MGR_ETH_ST_DOWN]
  Triggered event: [FCOE_MGR_ETH_EV_IF_UP]
  Next state: [FCOE_MGR_ETH_ST_UP]

4) FSM:<Ethernet1/40> Transition at 372977 usecs after Wed Sep 2 02:37:10 2009
  Previous state: [FCOE_MGR_ETH_ST_UP]
  Triggered event: [FCOE_MGR_ETH_EV_IF_L2_STATE_CHANGE]
  Next state: [FSM_ST_NO_CHANGE]

5) FSM:<Ethernet1/40> Transition at 505700 usecs after Wed Sep 2 02:37:11 2009
  Previous state: [FCOE_MGR_ETH_ST_UP]
  Triggered event: [FCOE_MGR_ETH_EV_LLS_CHANGE]
  Next state: [FSM_ST_NO_CHANGE]

6) FSM:<Ethernet1/40> Transition at 871231 usecs after Wed Sep 2 02:37:13 2009
  Previous state: [FCOE_MGR_ETH_ST_UP]
  Triggered event: [FCOE_MGR_ETH_EV_LLS_CHANGE]
  Next state: [FSM_ST_NO_CHANGE]

7) FSM:<Ethernet1/40> Transition at 924476 usecs after Wed Sep 2 05:04:47 2009
  Previous state: [FCOE_MGR_ETH_ST_UP]
  Triggered event: [FCOE_MGR_ETH_EV_IF_L2_STATE_CHANGE]
  Next state: [FSM_ST_NO_CHANGE]

8) FSM:<Ethernet1/40> Transition at 965102 usecs after Wed Sep 2 05:04:47 2009
  Previous state: [FCOE_MGR_ETH_ST_UP]
  Triggered event: [FCOE_MGR_ETH_EV_IF_LOGICAL_DOWN]
  Next state: [FSM_ST_NO_CHANGE]
```

```
9) FSM:<Ethernet1/40> Transition at 14002 usecs after Wed Sep 2 05:04:49 2009
  Previous state: [FCOE_MGR_ETH_ST_UP]
  Triggered event: [FCOE_MGR_ETH_EV_IF_DOWN]
  Next state: [FCOE_MGR_ETH_ST_DOWN]

10) FSM:<Ethernet1/40> Transition at 293007 usecs after Wed Sep 2 05:04:54 2009
  Previous state: [FCOE_MGR_ETH_ST_DOWN]
  Triggered event: [FCOE_MGR_ETH_EV_IF_UP]
  Next state: [FCOE_MGR_ETH_ST_UP]

11) FSM:<Ethernet1/40> Transition at 340253 usecs after Wed Sep 2 05:04:54 2009
  Previous state: [FCOE_MGR_ETH_ST_UP]
  Triggered event: [FCOE_MGR_ETH_EV_IF_L2_STATE_CHANGE]
  Next state: [FSM_ST_NO_CHANGE]

12) FSM:<Ethernet1/40> Transition at 214035 usecs after Wed Sep 2 05:04:55 2009
  Previous state: [FCOE_MGR_ETH_ST_UP]
  Triggered event: [FCOE_MGR_ETH_EV_LLS_CHANGE]
  Next state: [FSM_ST_NO_CHANGE]

13) FSM:<Ethernet1/40> Transition at 222109 usecs after Wed Sep 2 05:04:55 2009
  Previous state: [FCOE_MGR_ETH_ST_UP]
  Triggered event: [FCOE_MGR_ETH_EV_FCOE_CHANGE]
  Next state: [FSM_ST_NO_CHANGE]

14) FSM:<Ethernet1/40> Transition at 213533 usecs after Wed Sep 2 05:04:56 2009
  Previous state: [FCOE_MGR_ETH_ST_UP]
  Triggered event: [FCOE_MGR_ETH_EV_LLS_CHANGE]
  Next state: [FSM_ST_NO_CHANGE]

  Curr state: [FCOE_MGR_ETH_ST_UP]
```

Using the QLogic SANsurfer CLI

All the above tasks performed using the QLogic SANsurfer Fibre Channel HBA Manager can also be performed using the QLogic SANsurfer CLI. However, the QLogic SANsurfer CLI cannot discover and manage remote HBAs and CNAs, whereas the QLogic SANsurfer Fibre Channel HBA Manager can.

1. Open the QLogic SANsurfer CLI; it is installed in the QLogic Management Suite directory.
2. To display the DCB-related parameters, select FCoE.

```

Enter Selection:
Error: Invalid selection!

SANSurfer FC/CNA HBA CLI
v1.7.3 Build 16

Main Menu

1: General Information
2: HBA Information
3: HBA Parameters
4: Target/LUN List
5: iDMA Settings
6: Boot Device
7: Utilities
8: Beacon
9: Diagnostics
10: FCoE
11: Help
12: Exit

Enter Selection:
Error: Invalid selection!
    
```

- For the port, select option 2, Data Center Bridging. The CLI will display DCB-related parameters. As with the GUI, the parameters displayed in CLI include SAN Priority CoS, ETS, and Port Pause Type priority.

```

ENODE MAC Addr : 00:00:00:12:0E:B1
WPN           : 21-00-00-C0-DD-12-0E-B1
Desc         : QLE8142 IBM PCI Express to 10 GbE Dual Channel CNA (FCoE)

1: Information
2: Data Center Bridging
3: DCB Statistics
4: TLV
5: Return to Previous Menu

Note: 0 to return to Main Menu
Enter Selection: 2
-----
Data Center Bridging Port 0 Configuration Parameters
-----
Host Name           : eng_rhel5
HBA Instance       : 0
HBA Model          : QLE8142
HBA Description    : QLE8142 IBM PCI Express to 10 GbE Dual Chanr
-----
DCBX
-----
DCBX Enable        : Enabled
Willing            : Enabled
Port Pause Type   : Per Priority Pause
SAN Priority COS   : 3
-----
ETS
-----
Priority Tx Mode   : Bandwidth
SAN Bandwidth Percent : 50
SAN Unused Bw To LAN : Enabled
LAN Unused Bw To SAN : Enabled
    
```

- To obtain information details, select option 1. Similar to the GUI Information tab, the Information option in the CLI displays the operational values of parameters, such as VLAN ID, Max Frame Size, and Addressing Mode. These values cannot be configured with the management tool.

```

VI.7.3 Build 16
FCoE Utilities Menu
HBA Instance 0 (QLE8142 Port 1) : Online
ENode MAC Addr: 00:C0:DD:12:0E:B1
WWPN      : 21-00-00-C0-DD-12-0E-B1
Desc      : QLE8142 IBM PCI Express to 10 GbE Dual Channel CNA (FCoE)
1: Information
2: Data Center Bridging
3: DCE Statistics
4: TLV
5: Return to Previous Menu

Note: 0 to return to Main Menu
Enter Selection: 1
-----
HBA Instance 0: QLE8142 Port 1 WWPN 21-00-00-C0-DD-12-0E-B1 PortID C5-02-1F
-----
General Info
-----
VN Port MAC Address      : 0E:FC:00:C5:02:1F
VLAN ID                  : 1002
Max Frame Size           : 2500 (Baby Jumbo)
Addressing Mode          : FPMA
-----
Hit <RETURN> to continue
    
```

Managing the Ethernet (Networking) Interface of the CNA

Windows Platform

The Microsoft Windows platform supports the QLogic SANsurfer CNA Networking CLI. The look and feel of the networking CLI is similar to that of the QLogic SANsurfer Fibre Channel HBA CLI. However, the tool manages the networking interface of the CNA. In addition to managing the networking interface, the networking CLI, similar to the Fibre Channel GUI and CLI, displays DCB port-level information.

In addition to the QLogic SANsurfer CNA Networking CLI, the Windows Property page is supported to manage the networking portion of the CNA. However, the Windows Property page does not support advanced features such as VLANs and teaming.

Using the SANsurfer CNA Networking CLI

To use the QLogic SANsurfer CNA Networking CLI, follow these steps:

1. Invoke the QLogic SANsurfer CNA Networking CLI. It is installed in the QLogic Management Suite directory on Microsoft Windows platforms. When the CLI is invoked, it detects CNAs in the local server and provides a list of options, including support for advanced features such as VLANs and teaming.

```

main Interactive Menu
-----
1. CNA: 1 Port: 1           I/F Scope ID : 11
Program Version : 1.0.00.31   Driver Version: 01.00.00
MPI FW Version  : 1.34.00     Model       : QLE8142
Physical MAC    : 00:c0:dd:0a:b1:a8  LAA MAC    : 00:c0:dd:0a:b1:a8
-----

1. Display Program Version Information
2. Host Level Info & Operations
3. Configure Teams Menu
4. Configure VLANs Menu
5. CNA Level Info & Operations
6. Port Level Info & Operations
7. List All QLogic CNA Ports detected
8. Help
9. Select CNA Port
10. Refresh
11. Exit
enter selection: _
  
```

2. To create a team or insert a VLAN, be sure to load the teaming or VLAN driver before performing any VLAN or teaming operations. To install or update the driver, from the main menu select option 2, Host Level Info & Operations. Then from the next screen, select option 2, Install/Update VLAN/Teaming driver, All adapters. Next, the tool will ask "Do you want to use external source for VLAN/Teaming driver?". To install the VLAN/Teaming driver for the first time, select "no". To update the existing VLAN/Teaming driver, select "yes" followed by the path name to the driver files. In either case, the tool will next ask "Proceed with installation/update of VLAN/Teaming driver". Select "yes" and ensure the drivers are installed properly by checking the final message "Successfully installed vt driver: <version number>".

```

1. Display General system information
2. Install/Update VLAN/Teaming driver, All adapters
3. Select CNA Port
4. Refresh
5. Exit
enter selection: 2
Do you want to use external source for VLAN/Teaming driver? (yes, no) [no]:
Installed vt driver version      : 01.00.00.11
To be installed vt driver version: 01.00.00.11

Proceed with installation/update of VLAN/Teaming driver? (yes, no) [yes]:
Versions are equal.
About to install vt driver: 01.00.00.11

About to save Teaming/Vlan current state.
Refreshing interfaces ... Please wait ...
*** No Teams available ***
VLAN List:
No VLANs to display.
Done with saving teaming/Vlan current state.
About to remove VLAN(s).
VLAN List:
Refreshing interfaces ... Please wait ...

Nothing to display.
Successfully removed VLAN(s) or no VLANS(s) were present.
About to remove teams(s).
Refreshing interfaces ... Please wait ...
*** No teams available ***
Successfully removed team(s) or no teams(s) were present.
About to install the driver
Installing VLAN and teaming drivers ... Done.
Successfully installed vt driver: 01.00.00.11

Attempting to restore TEAMS and VLANS if configured prior to vtdriver update.
*** Teams:
*** Ports:
Successfully restored TEAMS/VLANS
Press the Enter key to continue.
    
```

After the teaming or VLAN driver has been installed or updated successfully, you can create the team or insert the new VLAN.

Creating Teams or Bonds

To create teams or bonds, follow these steps:

1. From the QLogic SANsurfer CNA Networking CLI main menu, select option 3, Configure New Team.
2. On the next screen, select option 3, Configure New Team. Specify the teaming parameters such as the team type and ports associated with the team. The tool will create a team based on these entries, and the result of the operation will be reported as successful or failed.

```

3. Configure New Team
4. Delete Team
5. Refresh
6. Exit
enter selection: 3
Select Team Type (1=Fail Over; 2=Load Balanced) [1]: 1
1. CNA: 1 CNA Port: 1 CNA Model: QLE8142 PortID: 1
   Mac Phys. Address: 00:c0:dd:0a:b1:a8 Loc. Mac: 00:c0:dd:0a:b1:a8
   CNA Serial Number: AFC0902A04526 MPI FW Version: 1.34.00 Adapter Alias: None
   Port Alias: Port_Alias_MY
   IPv4 Address: 10.65.68.200
   IPv6 Addresses: fe80::2c0:ddff:fe0a:b1a8
2. CNA: 1 CNA Port: 2 CNA Model: QLE8142 PortID: 1
   Mac Phys. Address: 00:c0:dd:0a:b1:a9 Loc. Mac: 00:c0:dd:0a:b1:a9
   CNA Serial Number: AFC0902A04526 MPI FW Version: 1.34.00 Adapter Alias: None
   Port Alias: None
   IPv4 Address: 10.65.68.100
   IPv6 Addresses: fe80::2c0:ddff:fe0a:b1a9
Select two or more CNA Port Indices [1,2 or ALL]: ALL

Attempting to create new team:

Team Description:
Team Type: Fail Over
Selected ports : 1, 2

CNA: 1 CNA Port: 1 MAC: 00:c0:dd:0a:b1:a8 Description: QLogic 10Gb PCI Ethernet Adapt
3 - Network Load Balancing Filter Device
   QLogic 10Gb PCI Ethernet Adapter #3 - Network Load Balancing Filter Device
CNA: 1 CNA Port: 2 MAC: 00:c0:dd:0a:b1:a9 Description: QLogic 10Gb PCI Ethernet Adapt
4 - Network Load Balancing Filter Device
   QLogic 10Gb PCI Ethernet Adapter #4 - Network Load Balancing Filter Device

About to create the team. Please wait ..
Successfully created new team.
Press the Enter key to continue.

```

Creating VLANs

To create VLANs, follow these steps:

1. From the QLogic SANsurfer CNA Networking CLI main menu, select option 4, Configure VLAN Menu.
2. On the next screen, select option 3, Add VLAN to Port or Team. Select the port or team's ListIndex, followed by the VLAN ID to be inserted for the given CNA port. After you specify the VLAN ID, the driver will display a message indicating whether the VLAN ID insertion was successful.

```

SANsurfer CNA Networking CLI
Configure VLANs Menu
-----
1. CNA: 1 Port: 1 I/F Scope ID : 11
Program Version : 1.0.00.31 Driver Version: 01.00.00
MPI FW Version : 1.34.00 Model : QLE8142
Physical MAC : 00:c0:dd:0a:b1:a8 LAA MAC : 00:c0:dd:0a:b1:a8
-----

1. Display VLANs List
2. Display VLANs Information
3. Add VLAN to Port or Team
4. Remove VLAN from Port or Team
5. Refresh
6. Exit
Enter selection: 3
VLAN List:
ListIndex: 1 CNA: 1 CNA Port: 1 VLAN ID: None MAC: 00:c0:dd:0a:b1:a8 Description: QLo
mic 10Gb PCI Ethernet Adapter #3 - Network Load Balancing Filter Device
ListIndex: 2 CNA: 1 CNA Port: 2 VLAN ID: None MAC: 00:c0:dd:0a:b1:a9 Description: QLo
mic 10Gb PCI Ethernet Adapter #4 - Network Load Balancing Filter Device

Select one or more ListIndices from the list [1, 2]: 1
Enter VLAN ID (1..4095): 1023
About to set VLAN ID: 1023 for ListIndex: 1
Successfully set VLAN ID.
Press the Enter key to continue.
  
```

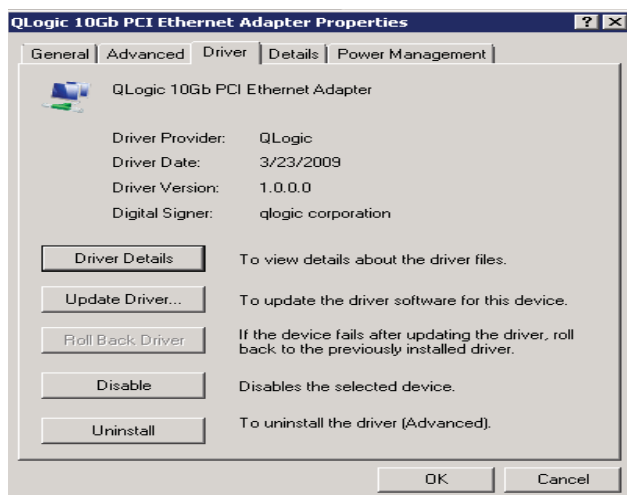
Using Microsoft Windows Property Pages

The Microsoft Windows Property pages comprise a native-platform tool for managing the networking interface of CNAs. This tool is mainly a configuration tool, with options for asset and inventory management. The basic features include driver updates, device asset and inventory information, and support for advanced features such as enabling and disabling of checksum offloads for protocols such as TCP, IP, and UDP.

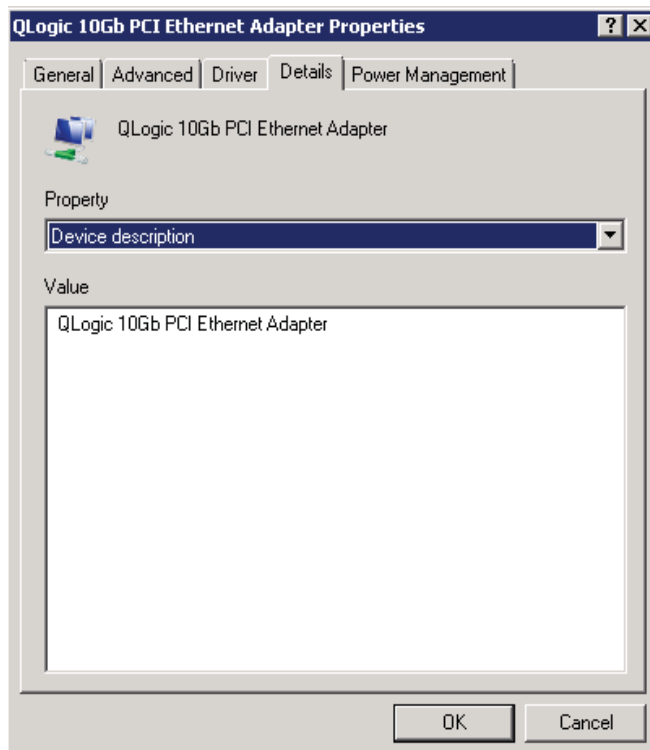
The current version of the Windows Property page does not support teaming or VLAN features.

Windows Property page options for managing the networking interface include:

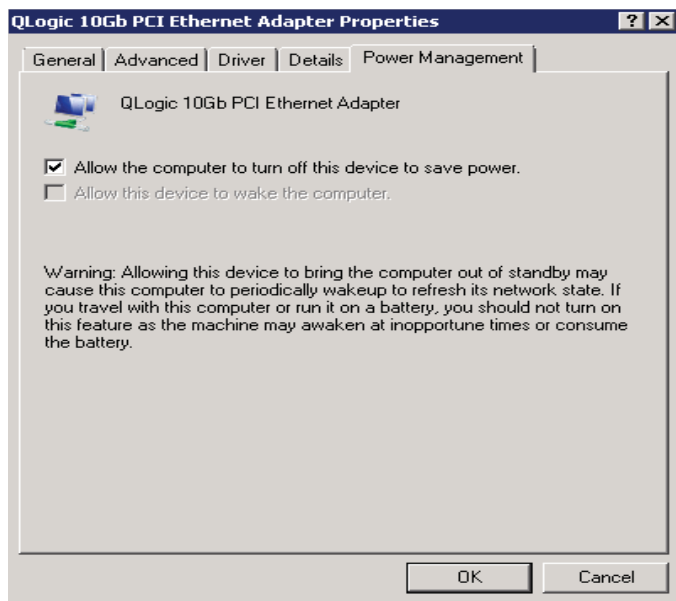
- Driver properties



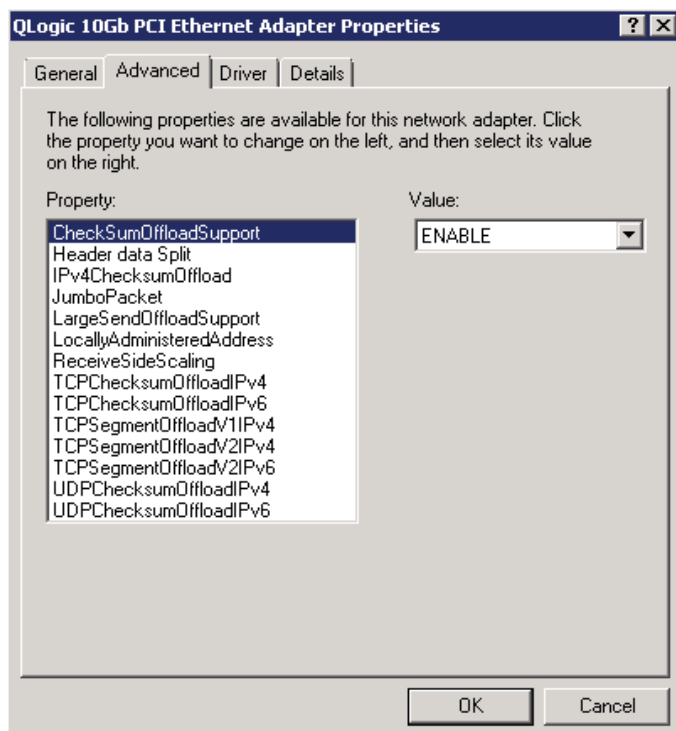
- Device details



- Power management properties



- Advanced properties



Linux Platform

The current release of the QLogic SANsurfer CNA Networking CLI is supported on Microsoft Windows platforms only. To configure the networking interface of a CNA on a Linux platform, use Linux native tools such as ethtool and ifconfig. The following sections provide instructions for configuring teaming or bonding and VLANs on QLogic 8100 Series CNAs for SuSE and RHEL Linux platforms.

SuSE Platform

To configure network settings using the SuSE Linux platform, you invoke the YaST2 command line with LAN options at the Linux prompt. The command will display a YaST2 window. YaST2 will discover all existing interfaces on the server. The YaST2 setup tool is the System Management utility of the SuSE Linux distribution. (It is the second version of the former YaST.)

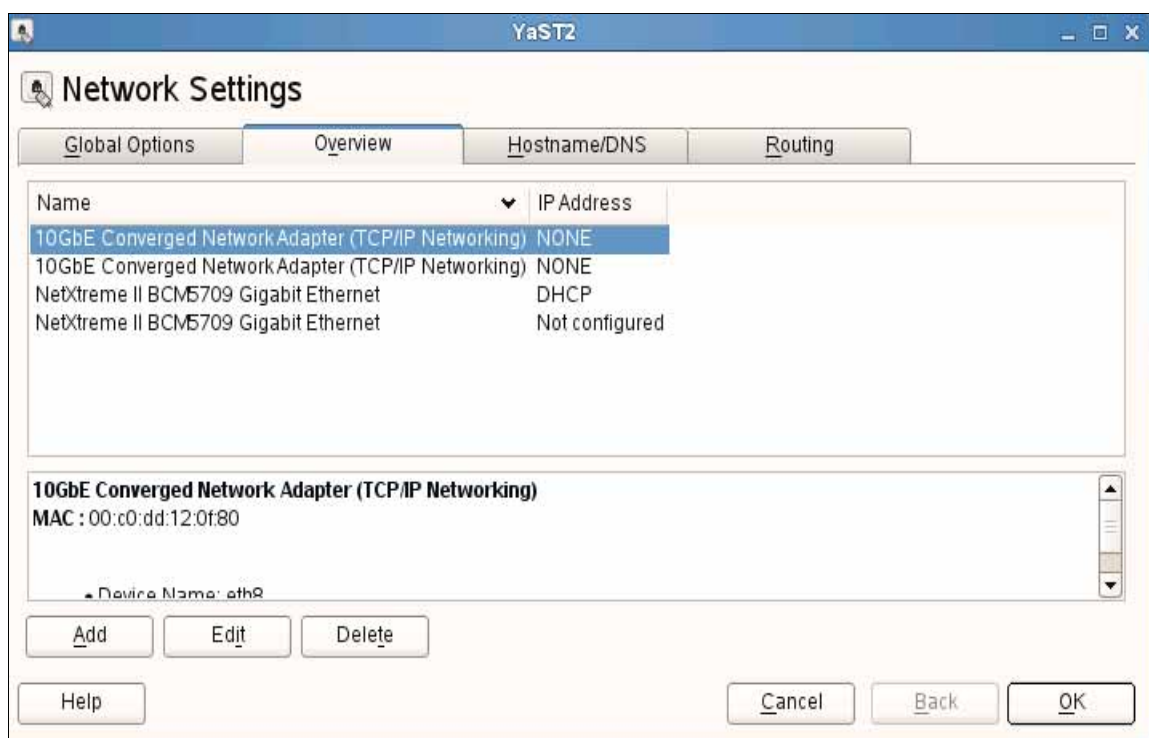
For teaming or bonding to work, the Linux bonding driver must be loaded. Run the following command to make sure that the bonding driver is loaded:

```
modprobe bonding
```

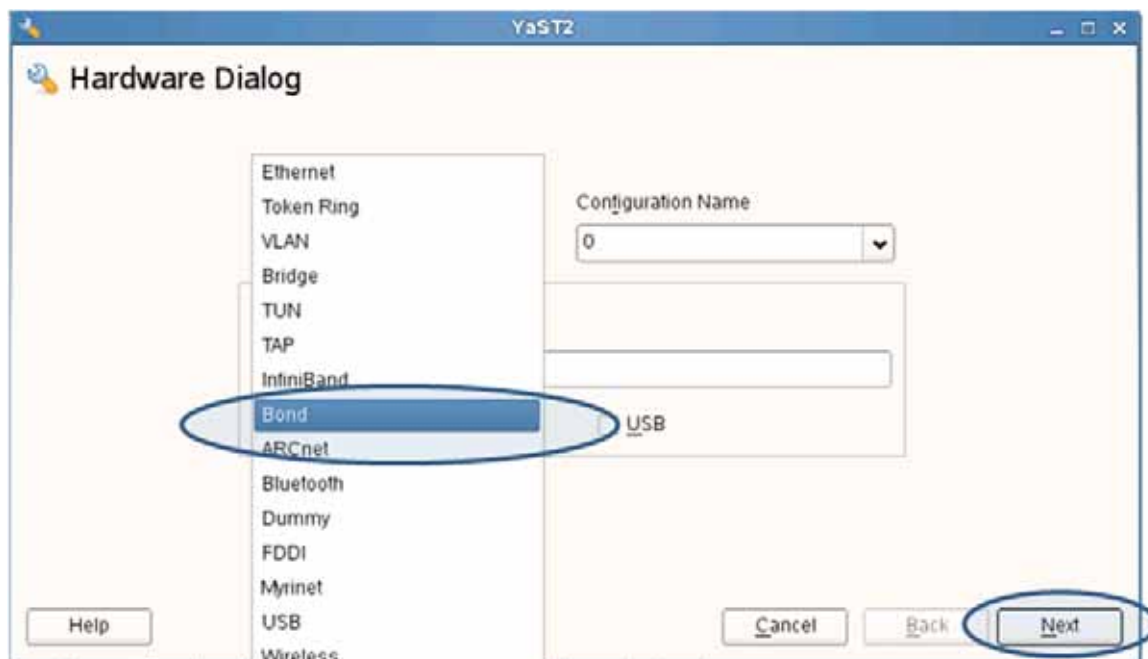
Creating Teams or Bonds

To create teams or bonds using SuSE Linux, follow these steps:

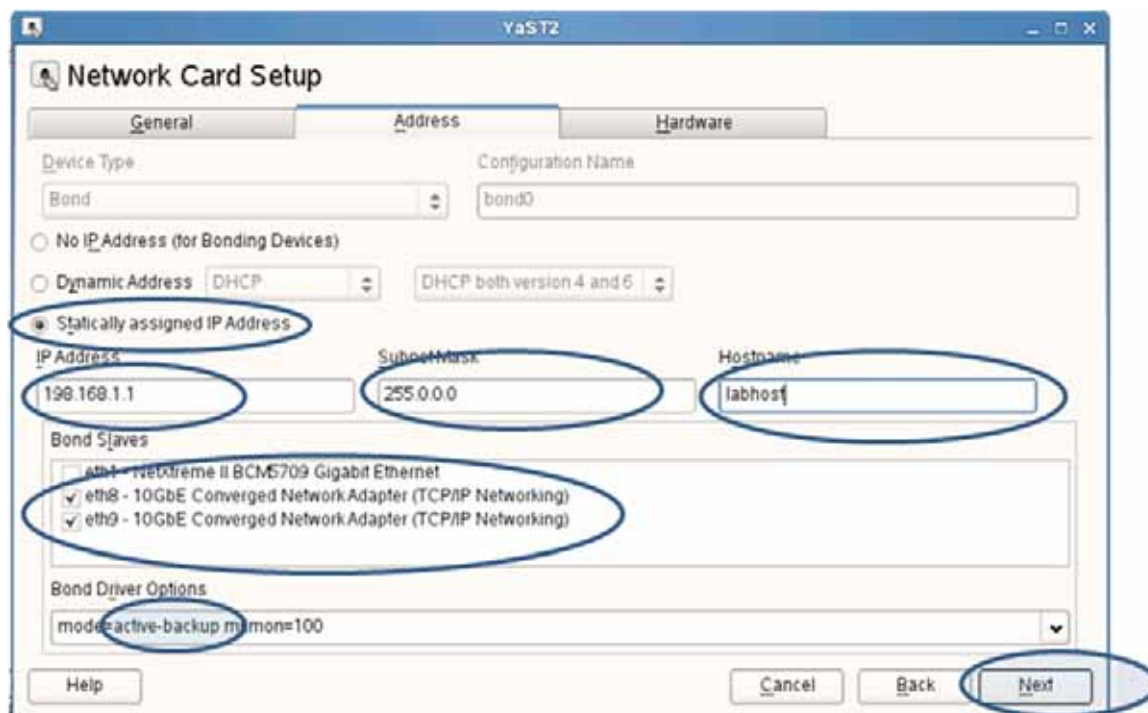
1. From the Linux prompt, invoke the YaST2 command line with LAN options. YaST2 will discover all existing interfaces on the server.



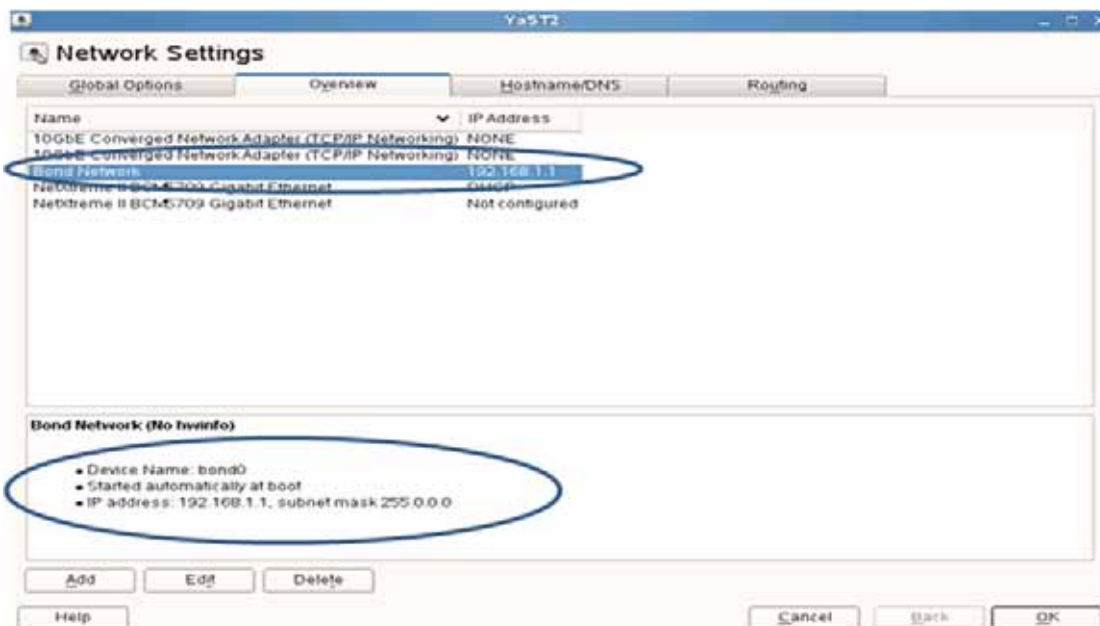
2. Select one of the ports and click the Add button.
3. The tool will display various options, including Bond and VLAN. For bonding, select the Bond option. Then click Next.



- YaST2 displays various options for bonding. Select the type of IP address to be assigned, the subnet mask, and the bonding ports. The mode can be either active-backup for failure or balance-RR for load balancing.



- Click Next. YaST2 will create a teaming interface. The newly created teaming interface will be visible among other interfaces of YaST2 tool.



Creating VLANs

To add or delete a VLAN on a given port, choose VLAN instead of Bond on the YaST2 Hardware Dialog screen.

RHEL Platform

On RHEL platforms, use the following steps to create teams and VLANs.

Creating Teams or Bonds

To create teams on RHEL platforms, follow these steps:

1. At the CLI, enter


```
cd /etc/sysconfig/network-scripts
```
2. Enter


```
cp ifcfg-ethX ethX.bak
```

 where X = the Ethernet interface number.
3. Create the bond interface:


```
cp ifcfg-ethX ifcfg-bond0
vi ifcfg-bond0
DEVICE=bond0
IPADDR=192.168.x.x
NETMASK=255.255.248.0
ONBOOT=yes
BOOTPROTO=none
USERCTL=no
MII_NOT_SUPPORTED=yes
```
4. Edit the Ethernet ports that you want to add to the bond:


```
vi ifcfg-ethX
DEVICE=ethX
USERCTL=no
ONBOOT=yes
```

```
MASTER=bond0
SLAVE=yes
BOOTPROTO=none
MII_NOT_SUPPORTED=yes
```

5. Edit the /etc/modprobe.conf file:

```
alias bond0 bonding
options bond0 max_bonds=2 miimon=100 mode=1 (0=round robin; 1=active/backup)
```

6. Load the bond driver module:

```
modprobe bonding
```

7. Restart the network:

```
service network restart
```

8. To view everything in the bond, enter:

```
less /proc/net/bonding/bond0
if mode=0 is selected it will read "bonding mode: load balancing (round-robin)
if mode=1 is selected it will read "bonding mode: fault-tolerance (round-robin)
```

Creating VLANs

To create VLANs on RHEL platforms, follow these steps:

1. At the CLI, remove the IP address information on the ethX interface:

```
ifconfig ethX 0.0.0.0
ifconfig ethX up
```

2. Configure a VLAN on the ethX interface using vconfig as shown here (1023 is the VLAN ID). If the IEEE 8021q.o module is not loaded, the vconfig command (when invoked for the first time) will automatically load the module.

```
vconfig add ethX 1023
```

3. Configure IP on the VLAN interfaces:

```
ifconfig ethX.1023 192.128.x.x netmask 255.255.248.0 up
```

4. Preserve the L=VLAN configuration across reboots by adding it to the configuration files:

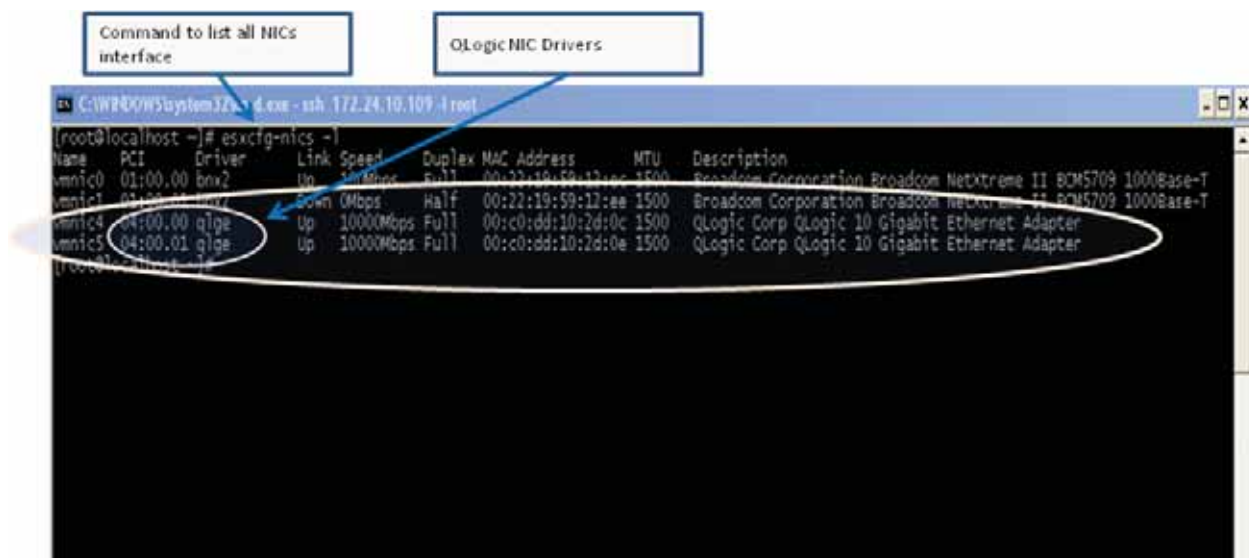
```
cd /etc/sysconfig/network-scripts/
vi ifcfg-ethX
DEVICE=ethX
ONBOOT=no
TYPE=Ethernet
vi ifcfg-ethX.1023
DEVICE=ethX.1023
IPADDR=192.128.x.x
NETMASK=255.255.248.0
ONBOOT=yes
BOOTPROTO=none
```

VMware Platform

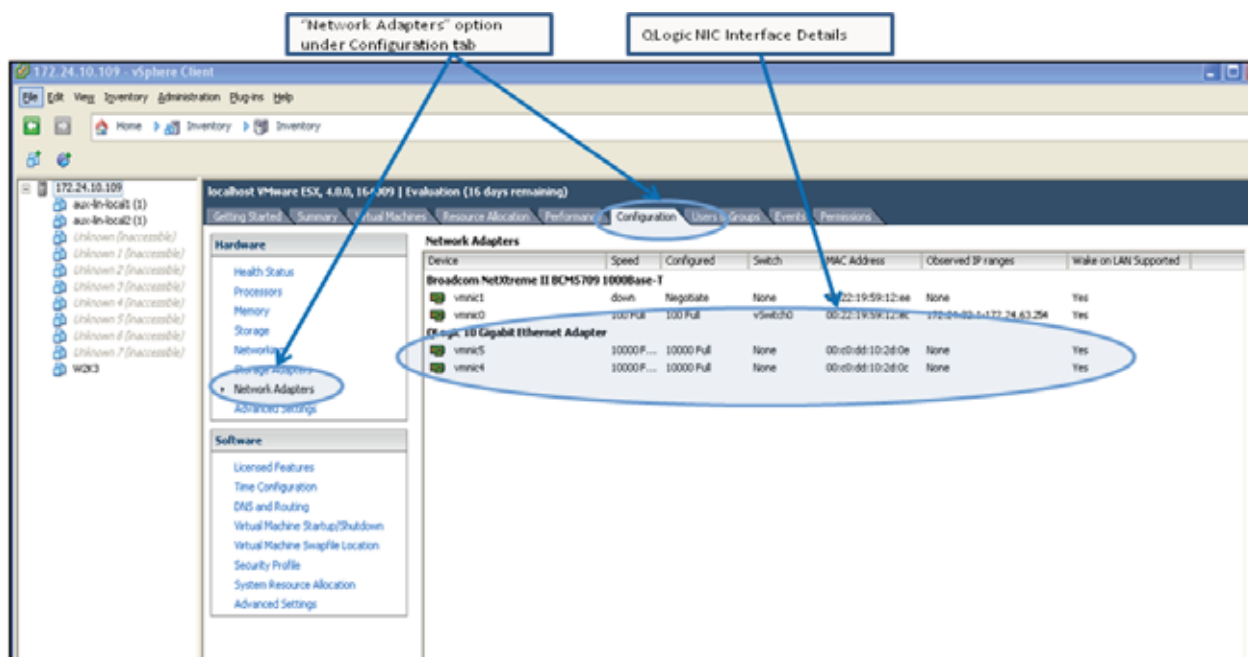
Teaming & VLAN configuration is accomplished using native VMware ESX tools. This includes the GUI-based vSphere Client, as well as native CLI tools such as `esxcfg-nics` and `esxcfg-vswitch`.

Port Groups are typically created to specify virtual NIC attributes in terms of Bandwidth, Security, VLAN and Teaming. The Virtual Machines (VMs) are then mapped to one of the Port Groups based on VM networking requirements. The remainder of this section assumes that the user is comfortable with basic working knowledge of ESX networking in terms of vSwitches, Port Groups, vSphere client, Service console etc.

Before executing any NIC related commands, confirm that the QLogic 81XX networking drivers are installed. This can be verified by executing “`esxcfg-nics -l`” from the service console. If drivers have been properly installed, the command will list all the QLE81xx networking interfaces as shown in the figure below.



The installation of networking drivers can also be verified via the vSphere client by clicking on the “Configuration” tab followed by the “Network Adapters” option. If the drivers have been installed successfully, the “Network Adapters” section will list all Networking Interfaces of the QLE81xx adapters as show below in the figure.



Creating Teams or Bonds

VMware ESX supports two modes of teaming:

- Load Balancing
- Network Failure Detection.

ESX provides the following four modes of load balancing (a) Virtual Switch Port based (b) MAC based (c) IP Hash based and (d) Explicit Failover Order based. In any ESX environment, load balancing only applies to transmitted traffic. Enabling load balancing on received traffic requires proper configuration of the external physical switch.

Virtual Switch Port based load balancing:

In a VMware topology, Virtual Machines (VMs) contain Virtual NICs (vNICs) which are attached to Virtual Switches. Each vNIC connects to a Virtual Switch through a Virtual Port, which is analogous to a physical port on a physical switch. Virtual Switch Port based load balancing routes vNIC networking traffic to/from a physical NIC (pNIC) based on Virtual Port assignment. This means that all networking traffic for a particular vNIC passes through the same pNIC. Virtual Port-to-pNIC assignment is done in a round-robin fashion to ensure even distribution of traffic amongst physical NICs. This load balancing mode requires no external switch configuration.

MAC based load balancing:

MAC based load balancing routes vNIC networking traffic to/from a pNIC based on the vNIC's MAC address. This is similar to Virtual Switch Port based load balancing in that all networking traffic for a particular vNIC always passes through the same pNIC and no external switch configuration is required.

IP Hash based load balancing:

IP Hash based load balancing routes vNIC networking traffic to/from a pNIC based on the source and destination IP addresses of each transmitted packet. All networking traffic for a particular source and destination IP address pair will pass through the same pNIC. Assuming a vNIC is passing traffic to multiple IP destinations, the vNIC's traffic will be sent and received through multiple pNICs. IP hash based load balancing enables link aggregation, whereby multiple pNICs can be grouped to provide a fatter networking pipe with aggregated bandwidth. This requires that the external switch is configured for 802.3ad teaming in static mode.

Explicit Failover Order based load balancing:

Explicit failover order routes all vNIC traffic to the highest priority pNIC. In case of a failover event on the highest priority pNIC, the user can designate a prioritized list of alternate pNICs to take its place.

Network Failure Detection provides two different methods for determining whether a link to a physical NIC has failed via (a) Link Status Only or (b) Beacon Probing.

Link Status Only:

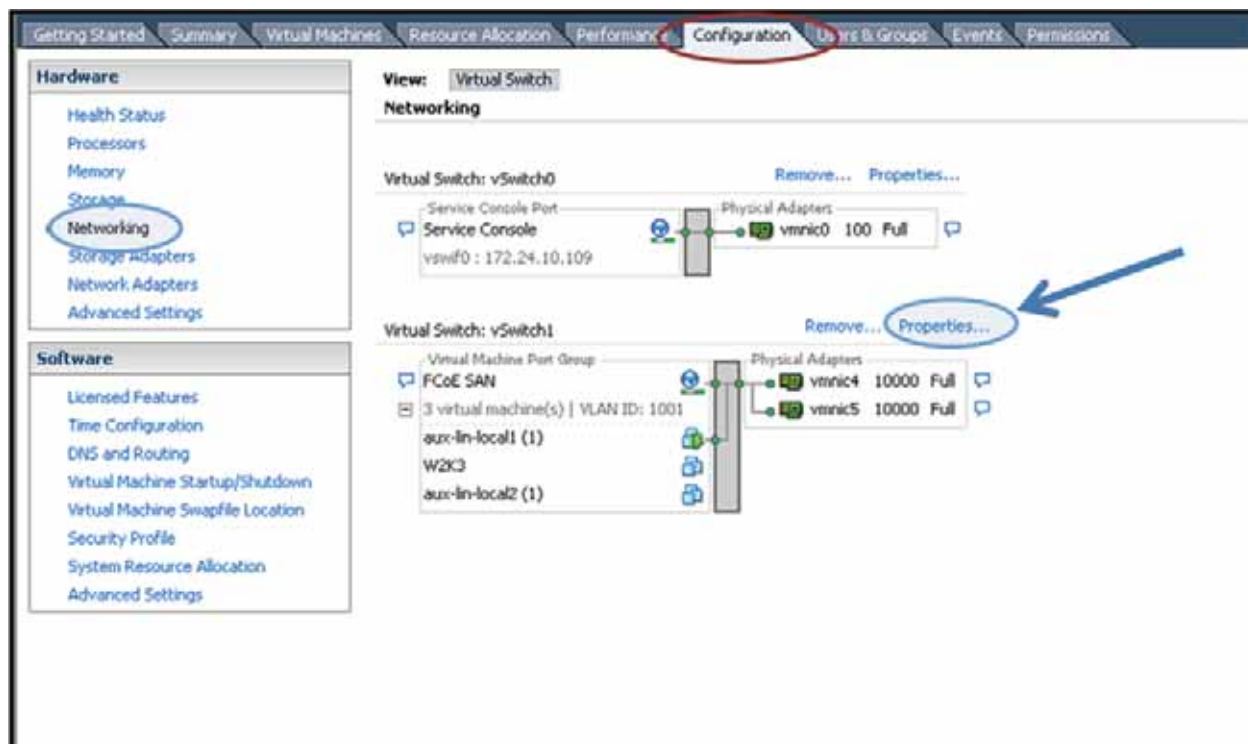
In Link Status Only failure detection, ESX relies solely on the pNIC to signal that the link has failed. This will detect some failures, such as cable pulls and physical switch power failures, but cannot detect many configuration errors.

Beacon Probing:

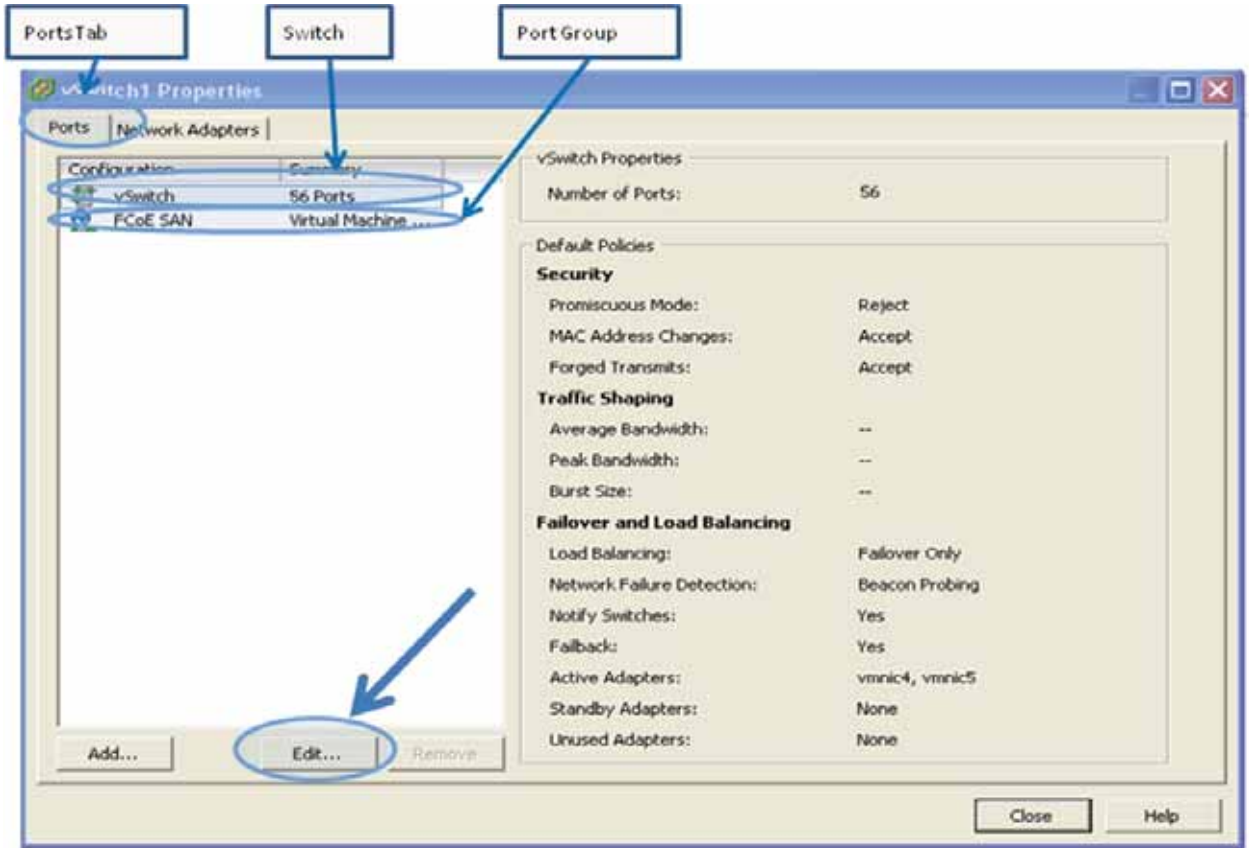
Beacon Probing failure detection transmits and listens for Ethernet broadcast frames known as beacon probes that are sent by other Ethernet adapters. ESX relies on beacon probing results, as well as link status, to identify when a link has failed. This will capture both link failures as well as configuration errors.

The following example shows how to create an IP hash based load balancing team with beacon probing failure detection using vSphere client. The same steps can be followed to create other modes of teaming on a virtual switch or port groups. The QLE81xx VMware drivers support all the above load balancing and network failure detection modes. In the example below, it is assumed that the port group “FCoE SAN” template is inherited from existing vSwitch properties.

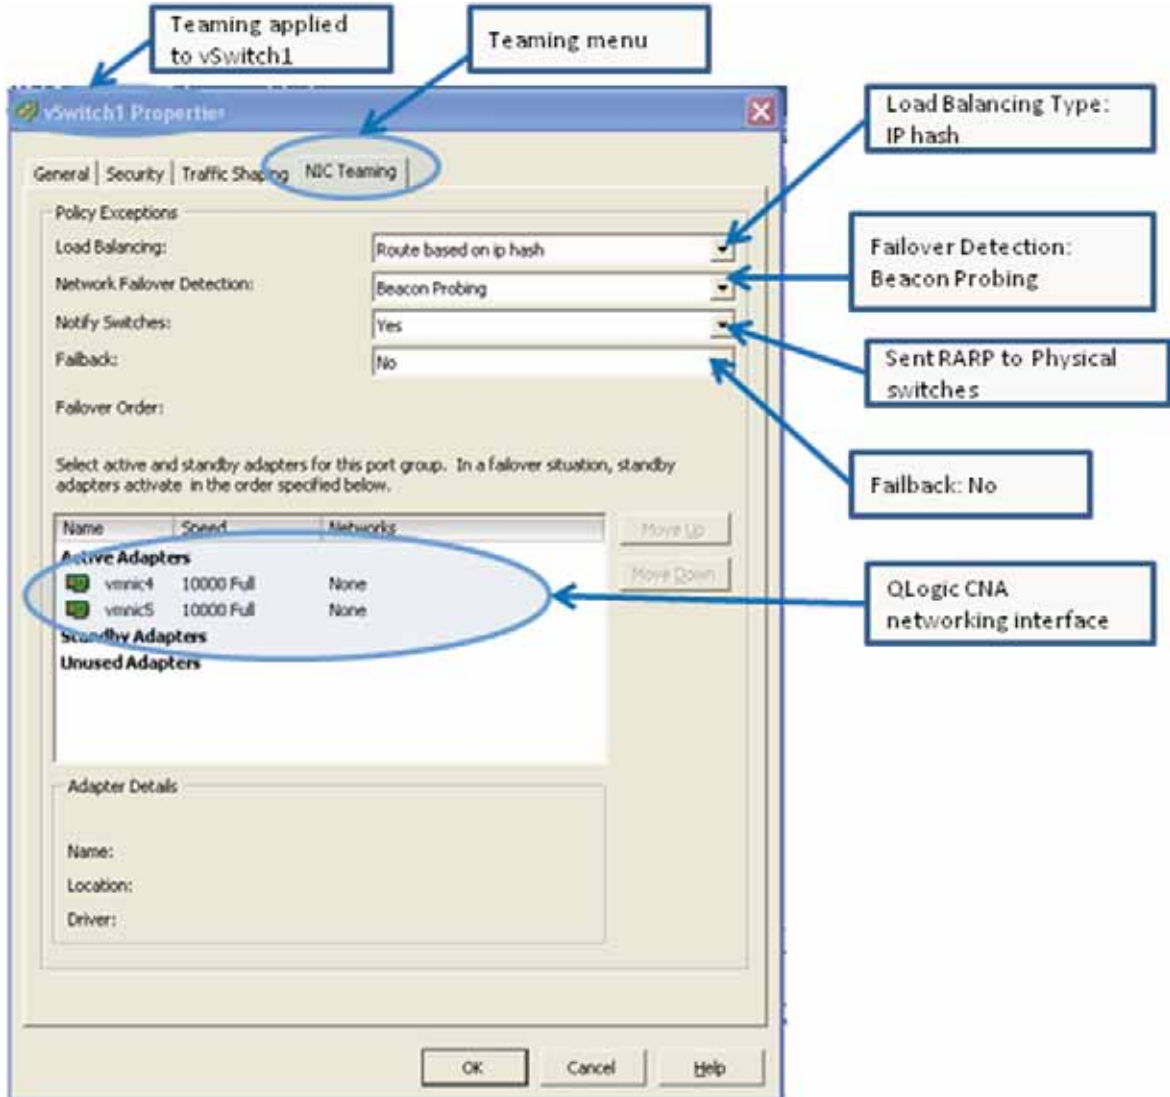
1. Click on the properties menu to bring up the vSwitch1 properties dialog box. The “Virtual Switch” pane is found under the “Networking” option of the “Configuration” tab.



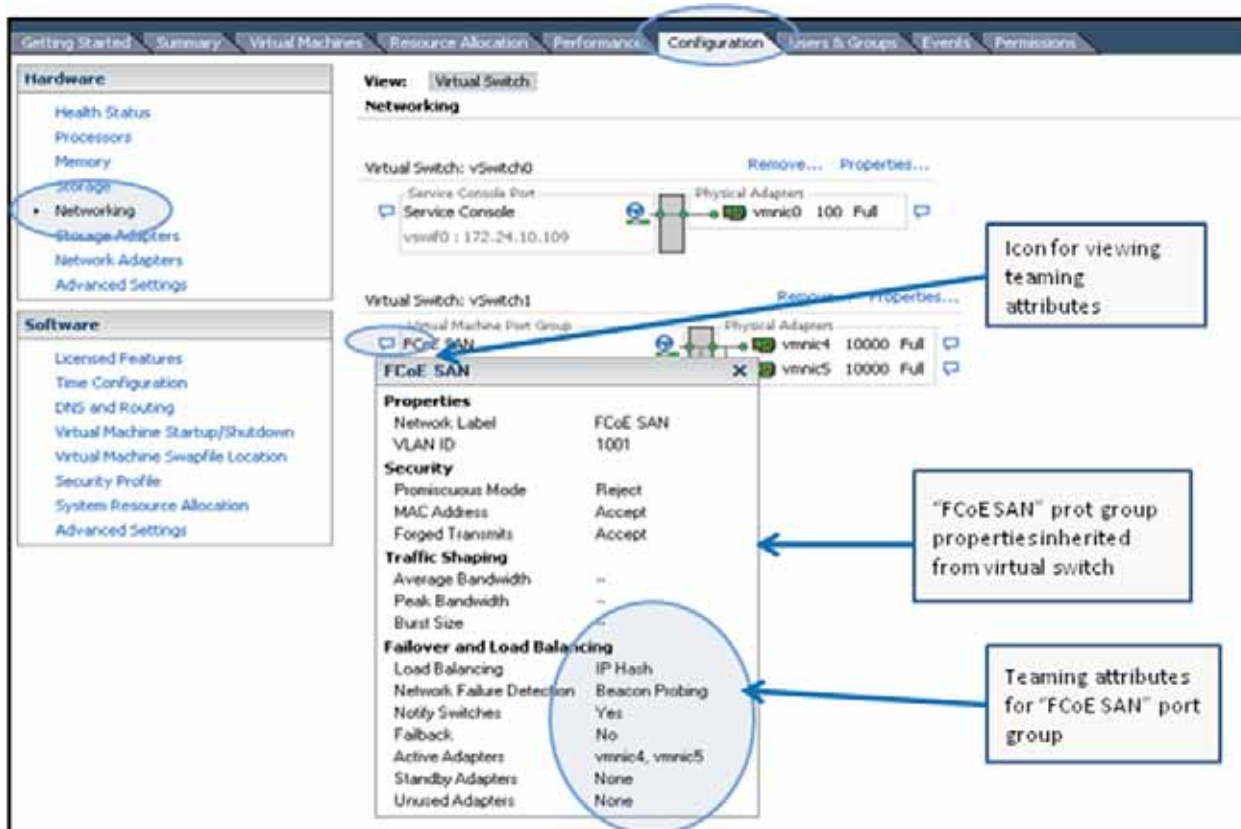
2. On the "Ports" tab, click on "vSwitch" followed by the "Edit" button. This will pop up a new window with various tabs to configure vSwitch attributes. Click on the "NIC Teaming" tab.



- Under the “NIC Teaming” tab, select “Route based on IP hash” for the Load Balancing option. For Network Failover Detection, select “Beacon Probing”, allowing NIC redundancy in the event of external failure. The “Notify Switches” option has been set to “Yes”, which will result in transmission of a RARP packet to update the physical switch look up table. The last option “Failback” has been set to “No”, meaning that if a failed pNIC comes back online, it will not be allowed back into active duty for the team.



- In order to verify the creation of the team, click on the icon next to the “FCoE SAN” port group in the “Networking” option of the Configuration tab. This will pop up a new window showing attributes of the “FCoE SAN” port group. At the bottom of the pop-up menu, under the “Failure and Load Balancing” heading, all the teaming attributes are displayed.



Creating VLANs

VMware ESX supports three modes of 802.1Q VLAN tagging.

- VM Guest Tagging (VGT)
- External Switch Tagging (EST)
- Virtual Switch Tagging (VST)

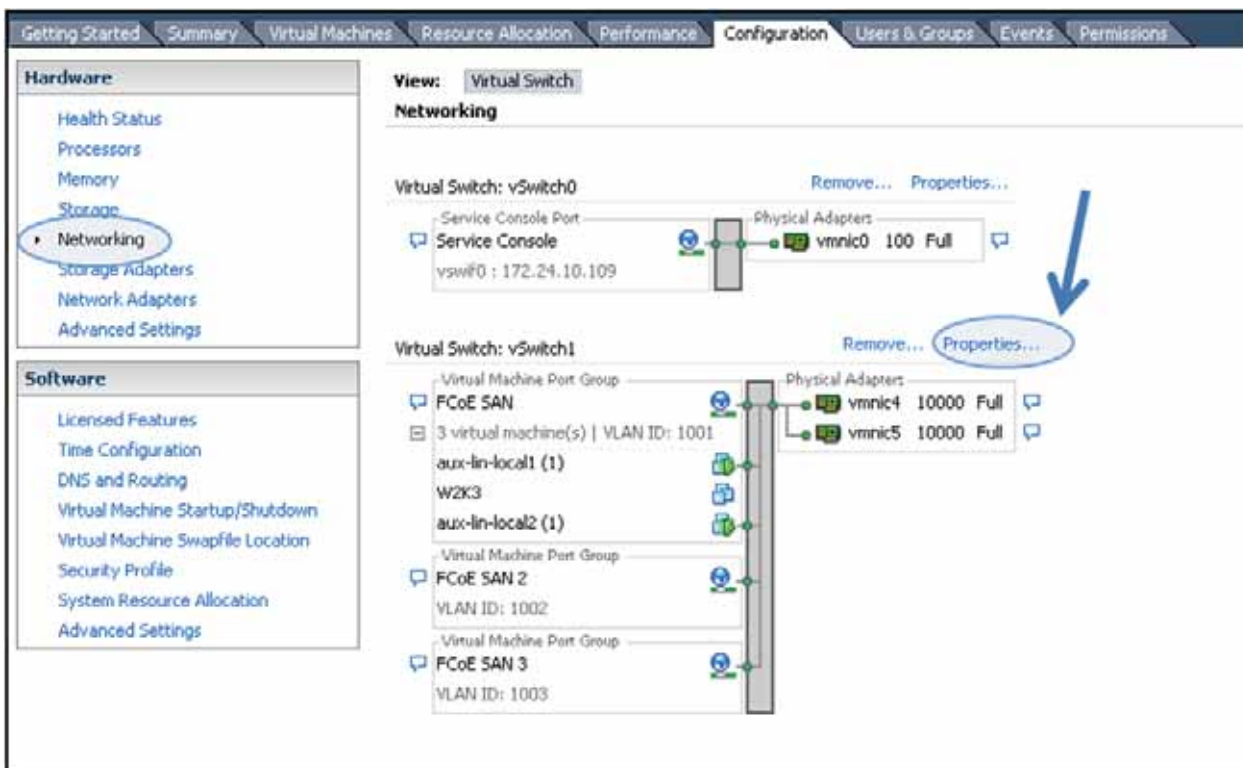
In VGT mode, the VLAN tagging and untagging is performed by the Virtual Machine using a VLAN-enabled driver. VLAN tags are preserved all the way from the VM to the external switch. The physical link between pNICs and the external switch must be configured in “Trunk Mode”. One advantage to this mode is it allows a single VM to participate in more than four VLANs.

In EST mode, VLAN tagging and untagging is performed by the physical switch. One drawback to this approach is that the number of VLANs supported per physical ESX host is limited to the number of installed pNICs.

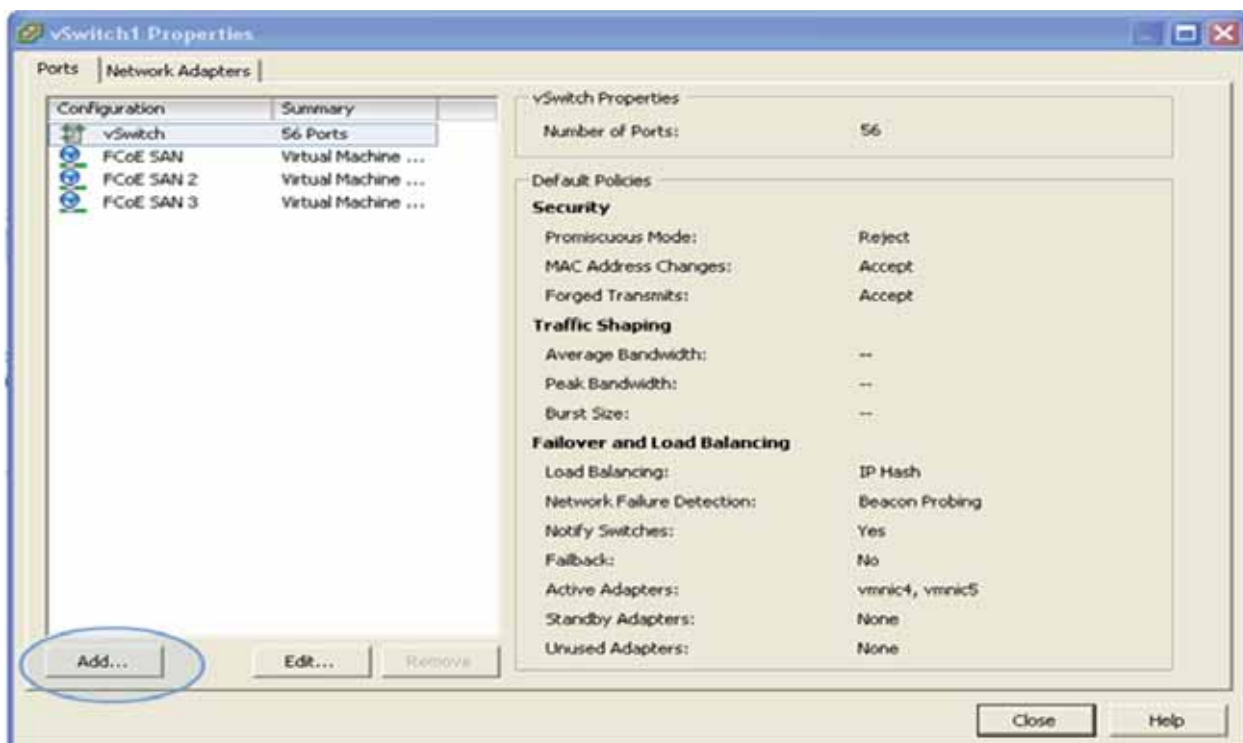
In VST mode, the VLAN tagging and untagging is performed by the Virtual Switch. After provisioning a port group for each VLAN, the Virtual Switch tags all outgoing Ethernet frames with the appropriate VLAN number and strips VLAN tags from incoming frames before forwarding to the appropriate VM. Unlike Virtual Guest Tagging, VST mode doesn't require installation of a VLAN-enabled driver. And unlike External Switch Tagging, VST mode supports multiple VLANs sharing the same pNIC. VST mode requires the link between the pNICs and external switch to be configured in “Trunk Mode”.

The following two steps should be performed to assign a VM to a VLAN. The first step consists of assigning a VLAN number to a given port group, and the second step consists of assigning the VM to a port group.

1. Click on the properties menu for the Virtual Switch.

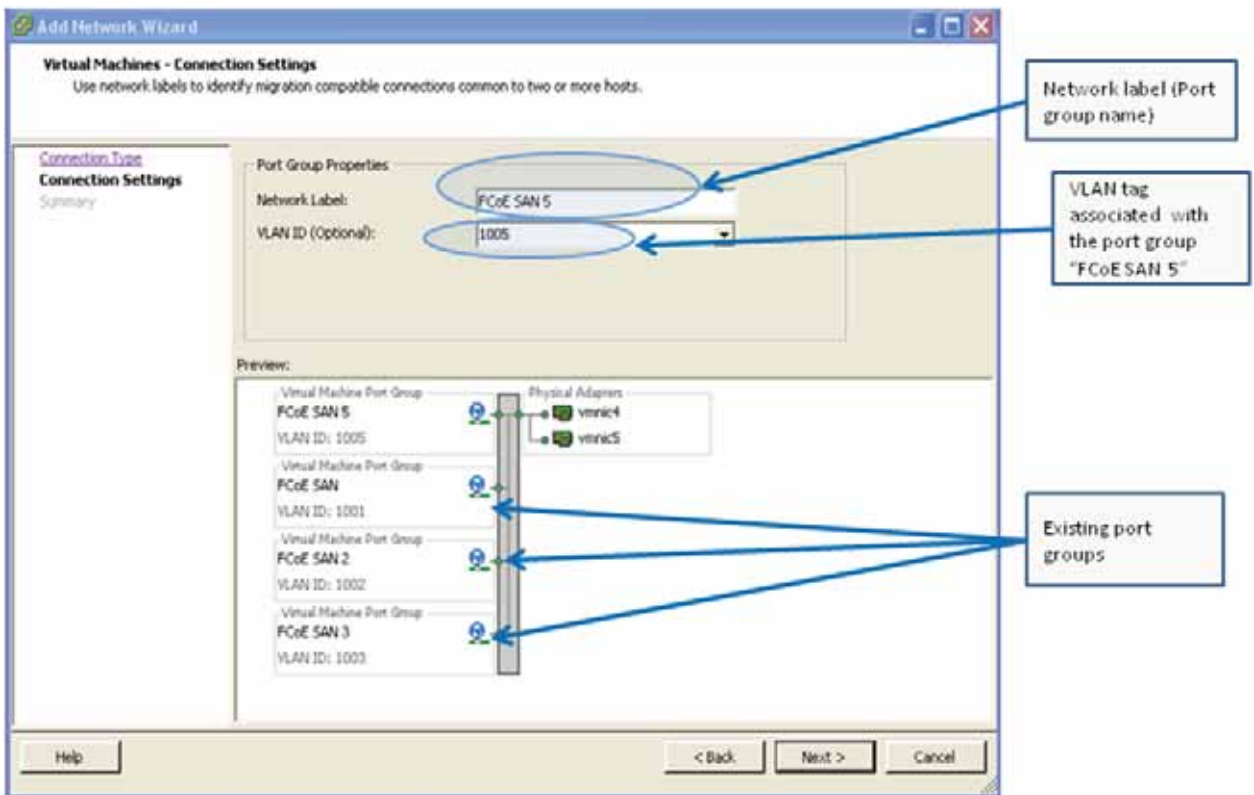


2. On the "Ports" tab of the vSwitch1 properties, click on the "Add" button. These will pop-up a new window with various options for connection type.



3. Click on the “Virtual Machine” connection type, followed by the “Next” button. The connection type “Virtual Machine” indicates a new port group creation. This will lead to the next screen for “Connection Setting” of the port group as shown below. For “Network Label” enter the name of the port group, which in this case is “**FCoE SAN 5**”, and associated VLAN tag, which in this case is **1005**. The screen also shows the existing port groups “**FCoE SAN**”, “**FCoE SAN 2**”, and “**FCoE SAN3**” in addition to the one being created. This step will assign VLAN tag to the port group.

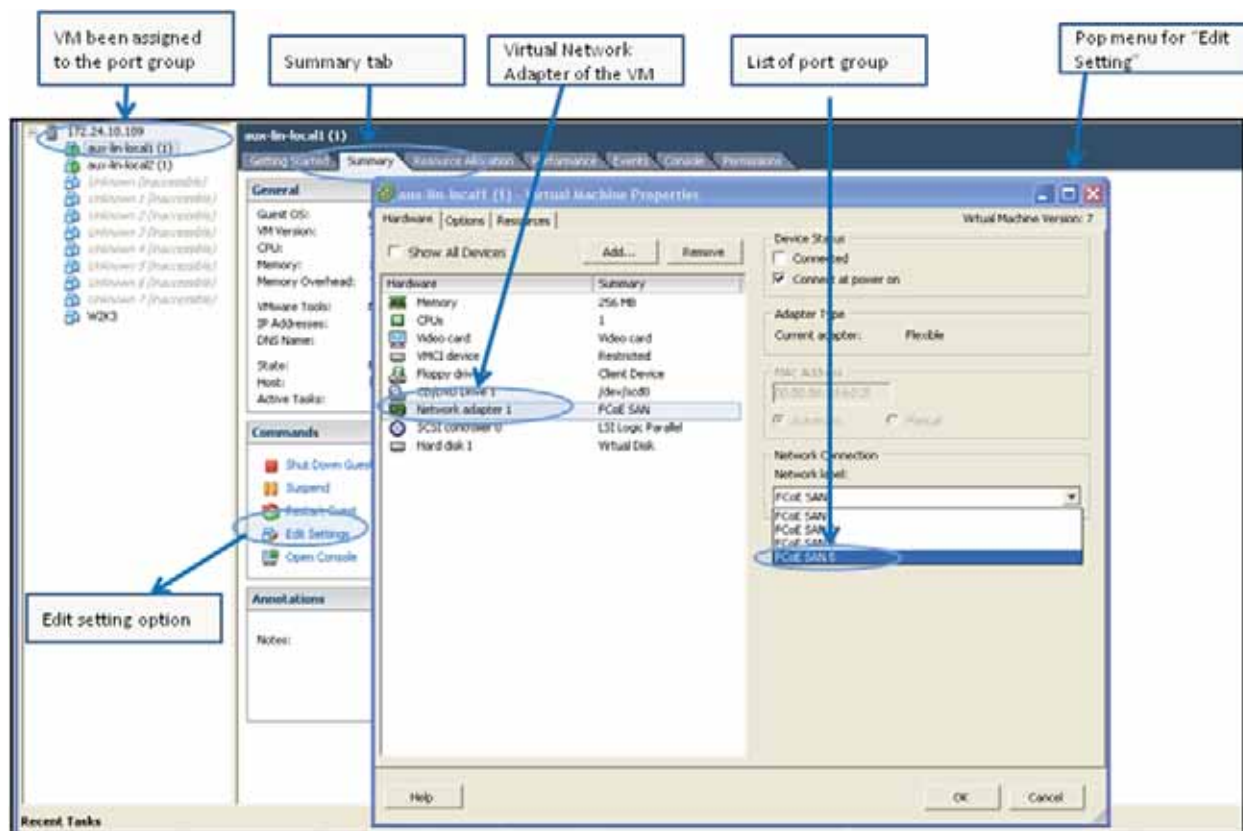
It is possible to update the VLAN tag or assign a VLAN tag to an existing port group if the port group was created without a VLAN tag. This can be achieved by visiting the switch attributes window and selecting the right port group under the “Ports” tab, followed by the “Edit” button. This will enable changing the attributes of a port group, including VLAN tag.



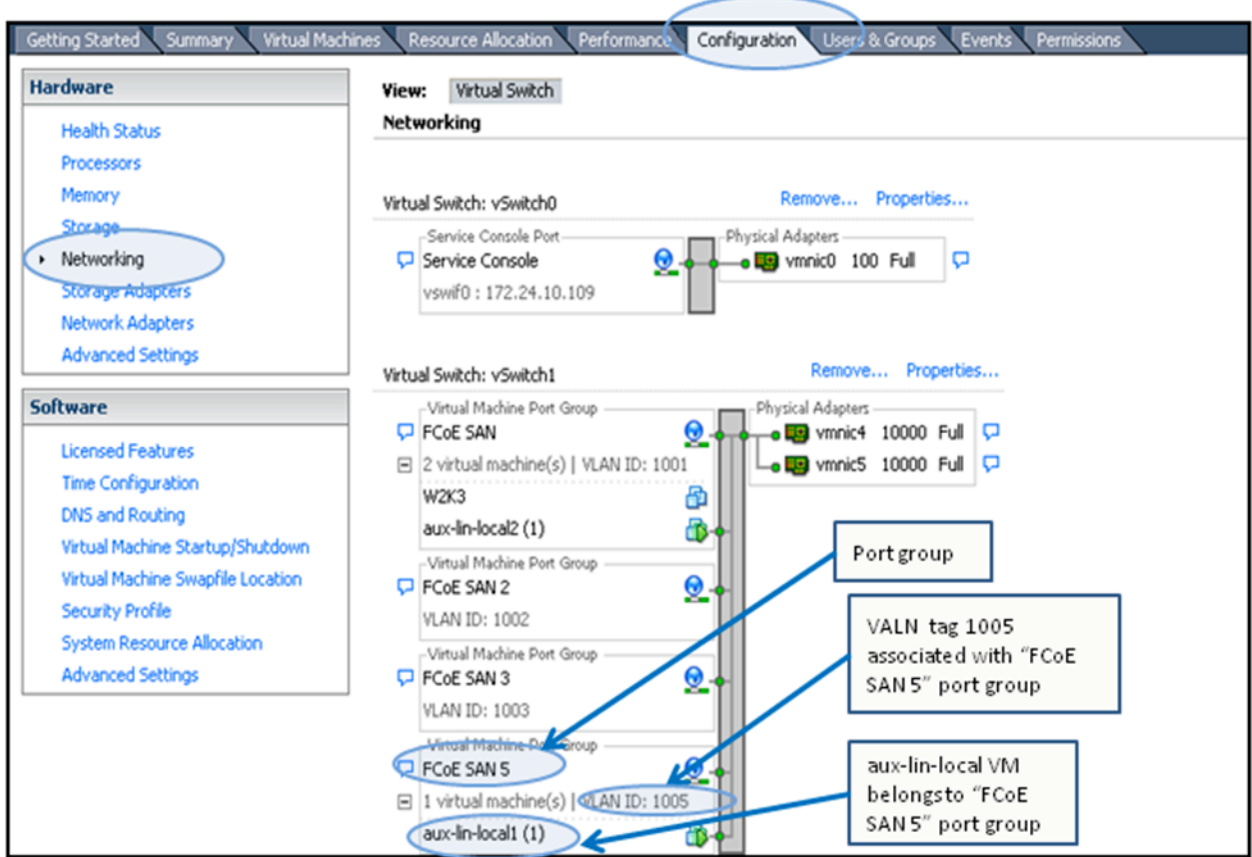
Assigning a Port Group to given VM:

1. Select the VM which needs to be assigned to a given port group. In this example, aux-lin-local1 was chosen. Under the Summary menu for the given VM, click on “Edit Setting”. This will pop-up the Virtual Machine Properties window with a list of hardware resources. Click on “Network Adapter 1” or the appropriate adapter if multiple virtual adapters are installed.

Upon selecting the “Network Adapter 1” virtual interface from the list of resources, the right hand side of the window will display a list of available “port group” under the “Network Connection” header. Select the appropriate port group (“FCoE SAN 5”) and assign it to the virtual network adapter.



1. Verify the port group assignment on the Virtual Switch. Click on the “Networking” option under the configuration menu of the server. As shown below for the “FCoE SAN 5” port group, the newly assigned “aux-lin-local” VM is one of the virtual machines belonging to the port group. Also, it shows the VLAN tag 1005 associated with the “FCoE SAN 5” port group. Basically, any LAN traffic from the “aux-lin-local” VM with a networking interface belonging to this port group will be marked with a VLAN tag of 1005.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, CCSI, Cisco Eos, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco-Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCPV, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumin, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLynx, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0910R)