



# Deploying Dell Networking MXL and PowerEdge I/O Aggregator FCoE (FIP Snooping Bridges) in a Cisco Nexus Environment

A Dell Deployment Guide

Dell Networking Development Solutions Engineering  
November 2013

©2013 Dell Inc., All rights reserved.

Except as stated below, no part of this document may be reproduced, distributed or transmitted in any form or by any means, without express permission of Dell.

You may distribute this document within your company or organization only, without alteration of its contents.

THIS DOCUMENT IS PROVIDED "AS-IS", AND WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED. IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE SPECIFICALLY DISCLAIMED. PRODUCT WARRANTIES APPLICABLE TO THE DELL PRODUCTS DESCRIBED IN THIS DOCUMENT MAY BE FOUND AT: <http://www.dell.com/learn/us/en/19/terms-of-sale-commercial-and-public-sector> Performance of network reference architectures discussed in this document may vary with differing deployment conditions, network loads, and the like. Third party products may be included in reference architectures for the convenience of the reader. Inclusion of such third party products does not necessarily constitute Dell's recommendation of those products. Please consult your Dell representative for additional information.

Trademarks used in this text:

Dell™, the Dell logo, Dell Boomi™, Dell Precision™, OptiPlex™, Latitude™, PowerEdge™, PowerVault™, Dell Networking™, PowerConnect™, OpenManage™, EqualLogic™, Compellent™, KACE™, FlexAddress™, Force10™ and Vostro™ are trademarks of Dell Inc. Other Dell trademarks may be used in this document. Cisco Nexus®, Cisco MDS®, Cisco NX-OS®, and other Cisco Catalyst® are registered trademarks of Cisco System Inc. EMC VNX®, and EMC Unisphere® is registered trademarks of EMC Corporation. Intel®, Pentium®, Xeon®, Core® and Celeron® are registered trademarks of Intel Corporation in the U.S. and other countries. AMD® is a registered trademark and AMD Opteron™, AMD Phenom™ and AMD Sempron™ are trademarks of Advanced Micro Devices, Inc. Microsoft®, Windows®, Windows Server®, Internet Explorer®, MS-DOS®, Windows Vista® and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat® and Red Hat® Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® and SUSE® are registered trademarks of Novell Inc. in the United States and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. Citrix®, Xen®, XenServer® and XenMotion® are either registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries. VMware®, Virtual SMP®, vMotion®, vCenter® and vSphere® are registered trademarks or trademarks of VMware, Inc. in the United States or other countries. IBM® is a registered trademark of International Business Machines Corporation. Broadcom® and NetXtreme® are registered trademarks of Broadcom Corporation. Qlogic is a registered trademark of QLogic Corporation. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and/or names or their products and are the property of their respective owners. Dell disclaims proprietary interest in the marks and names of others.



## Revisions

Date	Description	Authors
November 2013	Version 2.0 Release	Ed Blazek, Kevin Locklear



# Table of contents

Revisions .....	3
1 Introduction.....	7
1.1 Configuration Overview.....	7
1.2 Configuration Table .....	8
1.3 Component Information .....	9
1.4 Configuration 1 – Dell Networking MXL and PowerEdge M I/O Aggregator FSBs (FIP snooping bridges) in Nexus Fabric mode environment .....	10
1.5 Configuration 2 – NPV (N-Port Virtualization) mode with MDS 9148 and Dell modular FSBs (FIP snooping bridges) .....	11
2 M1000e I/O Modules (MXL and I/O Aggregator) .....	12
2.1 Flex IO Expansion Modules (External Ports) .....	12
2.2 I/O Module Internal Ports .....	13
2.3 Dell Networking MXL.....	14
2.4 PowerEdge I/O Aggregator.....	14
2.5 Factory Default Settings.....	15
2.5.1 The Dell Power I/O Aggregator.....	15
2.5.2 The Dell Networking MXL.....	15
2.6 FIP Snooping .....	15
3 M1000e Chassis Blade Servers and Network Adaptor .....	16
3.1 WWN/MAC Addresses .....	16
3.2 Configuring BCM57810 Network Adapter .....	18
3.2.1 Device Hardware Configuration .....	18
3.2.2 MBA Configuration.....	18
3.2.3 FCoE Boot Configuration .....	18
3.2.4 NIC Partition Configuration .....	19
3.2.5 Boot Sequence.....	20
3.2.6 Booting after NETWORK ADAPTOR configuration.....	20
3.3 Configuring QME8262 Network Adapter.....	21
4 Configuration 1 – Dell Networking MXL / PowerEdge M I/O Aggregator (FSBs) with Cisco Nexus fabric mode topology .....	22
4.1 Setup Cisco Nexus 5000 series switch (FCF) .....	23



4.2	Dell PowerEdge M I/O Aggregator and Networking MXL modular FSBs - Configuration in a fabric mode topology .....	27
5	Configuration 2 – Dell Networking MXL/PowerEdge M I/O Aggregator (FSBs) with Cisco Nexus NPV mode topology .....	29
5.1	Setup Cisco Nexus 5000 series and Cisco MDS 9100 series switches in an NPV mode topology .....	30
5.2	Dell Networking MXL (FSB) steps in Nexus NPV topology .....	34
6	EMC setup and assigning a LUN to Initiator .....	35
6.1	Broadcom CTRL+R .....	35
6.2	Initial initiator confirmation .....	35
6.3	Two ways to confirm the Host LUN ID info .....	37
6.4	Further explanation of key points: .....	38
7	Dell PowerEdge M I/O Aggregator (FSB) Validation .....	39
7.1	show interface status .....	39
7.2	show fip-snooping config .....	39
7.3	show fip-snooping fcf .....	40
7.4	show fip-snooping enode .....	40
7.5	show fip-snooping sessions .....	40
7.6	show fip-snooping statistics .....	41
7.7	show interface dcbx detail .....	42
7.8	show interfaces pfc detail .....	43
7.9	show interfaces ets detail .....	45
8	Cisco Nexus 5548UP and MDS 9148 Validation .....	49
8.1	Verifying Virtual Interfaces .....	49
8.1.1	show interface vfc vfc-id .....	49
8.1.2	show interface brief .....	49
8.2	Verifying NPV .....	50
8.2.1	show npv status .....	50
8.2.2	show npv flogi-table .....	51
8.2.3	show fcns database .....	51
8.3	show flogi database .....	52
8.4	show zoneset active .....	52
8.5	show spanning-tree summary .....	53
A	Basic Terminology .....	54



B   References..... 56



# 1 Introduction

This deployment guide provides configuration information to install Dell M1000e Blade server chassis I/O Modules "PowerEdge M I/O Aggregators" and "Dell Networking MXL" in an FCoE (FIP Snooping) topology. FCoE topologies exist in multiple network configurations and this document will focus on a select numbers of these. The configuration examples will have similar products to reduce the overlapping of content yet provide a high amount of customer coverage. Enterprise switch configuration items like security, Inter Switch Links (ISLs), Virtual port Channels (VPCs) should be accounted for.

## 1.1 Configuration Overview

The suggested configurations used to deploy solutions are presented in a sequential order for reading but as the below figure represents this is more of a simultaneous process. There are dependencies occurring during configuration that rely on other parts of the process, those dependencies will be explained during the deployment steps. Storage configurations are not covered in any depth due to the varied nature of those offerings in the general datacenter marketplace.

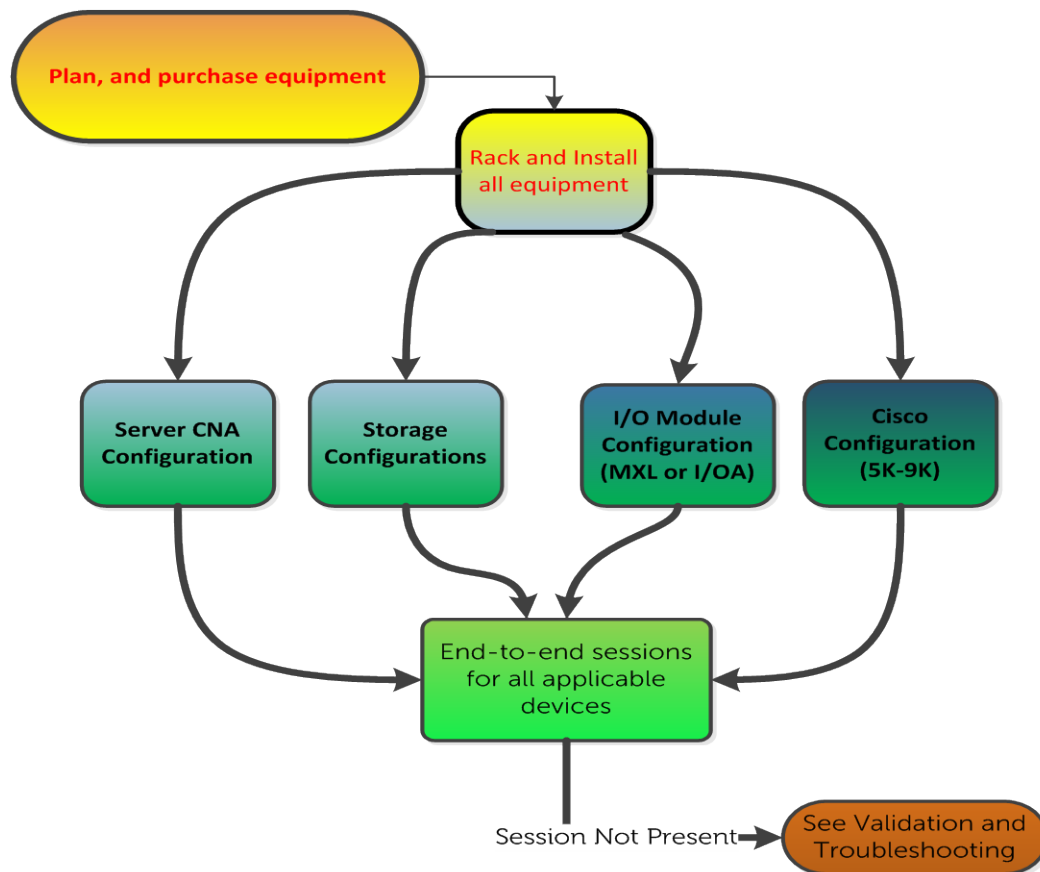


Figure 1 Configuration Overview



## 1.2 Configuration Table

The following table shows the information for each server and its applicable network adapter. Pre Planning is done to capture the FIP MAC address and World Wide Port Names (WWPN) for each network adapter port. You can then plan which VSAN each side of the fabric will use and the applicable chosen VFC numbers. The VSAN number should be easy to manage that facilitates troubleshooting (range 1 – 4094). Keep in mind that each side of the fabrics (SAN A and SAN B) will need different VSAN numbers.

		<b>SAN A</b>	<b>SAN B</b>
<b>Storage</b>	Storage Processor WWPN	50:06:01:6F:3E:E0:18:70	
	Boot LUN	0	0
<b>Server 1</b>	VSAN Number	2	3
	FCoE VLAN	1000	1001
	VFC Number	101	201
	Binding method	MAC	MAC
	Physical Port	1/1-2	1/1-2
	MAC Address of NETWORK ADAPTOR	20:01:5C:F9:DD:16:EF:07	20:01:5C:F9:DD:16:F0:10
	WWPN (20:01 + FIP MAC)	5C:F9:DD:16:EF:07	5C:F9:DD:16:F0:10
	FIP MAC	20:00:5C:F9:DD:16:EF:07	20:00:5C:F9:DD:16:F0:10
	WWNN (20:00 + FIP MAC)		
<b>Server 2</b>	FC interfaces	FC2/1 FC2/2	FC2/1 FC2/2
	VSAN Number	2	3
	FCoE VLAN	1000	1001
	VFC Number	102	202
	Binding method	MAC	MAC
	Physical Port	1/1-2	1/1-2
	MAC Address of NETWORK ADAPTOR	20:01:5C:F9:DD:16:EF:21	20:01:5C:F9:DD:16:F1:7E
	WWPN (20:01 + FIP MAC)	5C:F9:DD:16:EF:21	5C:F9:DD:16:F1:7E
<b>Cisco 5500</b>	FIP MAC	20:00:5C:F9:DD:16:EF:21	20:00:5C:F9:DD:16:F1:7E
	WWNN (20:00 + FIP MAC)		
<b>Cisco MDS</b>	FC interfaces	FC2/1 FC2/2	FC2/1 FC2/2
<b>Cisco 5500</b>	Physical Ports	FC 1/1-2 FC 2/1-2	FC 1/1-2 FC 2/1-2
<b>Cisco MDS</b>	Physical Ports	FC 1/13-14	FC 1/13-14

Table 1 Configuration Information





## 1.3 Component Information

In this table is the components used at the time this document was created and there various versions.

	Component	Version
<b>Chassis / Server</b>	M1000e Chassis Management Controller	4.50
	Dell PowerEdge M I/O Aggregator	9.2(0.0)
	Dell Networking MXL	9.2(0.0)
	Dell PowerEdge M620 Blade Server(s)	1.7.6 (BIOS)
	Lifecycle Controller	1.1.5.165
	Broadcom 10Gb 2P 57810S-k Mezzanine Card	7.6.15 (Firmware)
	QLogic 10Gb 2P QME8262-k Mezzanine Card	1.12.61 (Firmware)
<b>Storage Equipment</b>	EMC_3U VNX 5300	05.32.000.5.008
<b>Cisco Equipment</b>	Cisco Nexus 7004	6.1 (4)
	Cisco Nexus 5548UP	6.0(2)N1(2a)
	Cisco MDS 9148 (system and kickstart)	6.2(3)
<b>Cables</b>	SFP+ Optical Transceivers (SR or LR) with Fiber Cables	5 Meter Cable

Table 2 Component Information



## 1.4 Configuration 1 – Dell Networking MXL and PowerEdge M I/O Aggregator FSBs (FIP snooping bridges) in Nexus Fabric mode environment

In this configuration, a two-port connection between either the Dell Networking MXL or PowerEdge M I/O Aggregators (IOAs) acting as FIP snooping bridges (FSBs). The connection shows I/O Modules in B1 and B2 slots of the M1000e chassis. In a typical business environment, most configurations include several connections between servers, networking and storage devices.

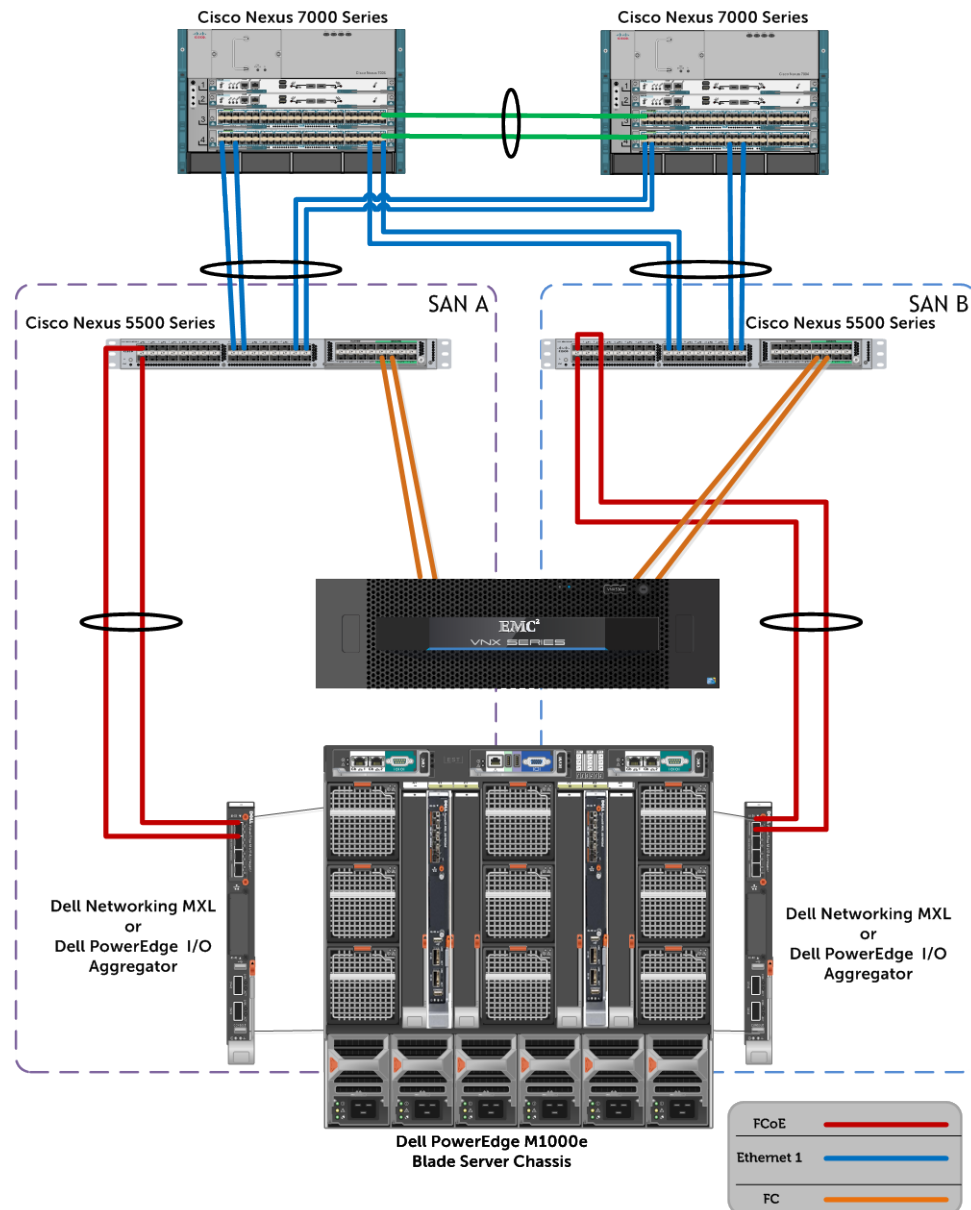


Figure 2 Configuration 1 – Fabric mode with FSBs (FIP Snooping Bridges)

## 1.5 Configuration 2 – NPV (N-Port Virtualization) mode with MDS 9148 and Dell modular FSBs (FIP snooping bridges)

In this configuration, we have two port connections from either the Dell Networking MXL or PowerEdge M I/O Aggregators similar to the previous example but this configuration is running NPV mode with MDS devices. As with Configuration 1, a typical business environment will include multiple connections between servers, networking devices and storage devices.

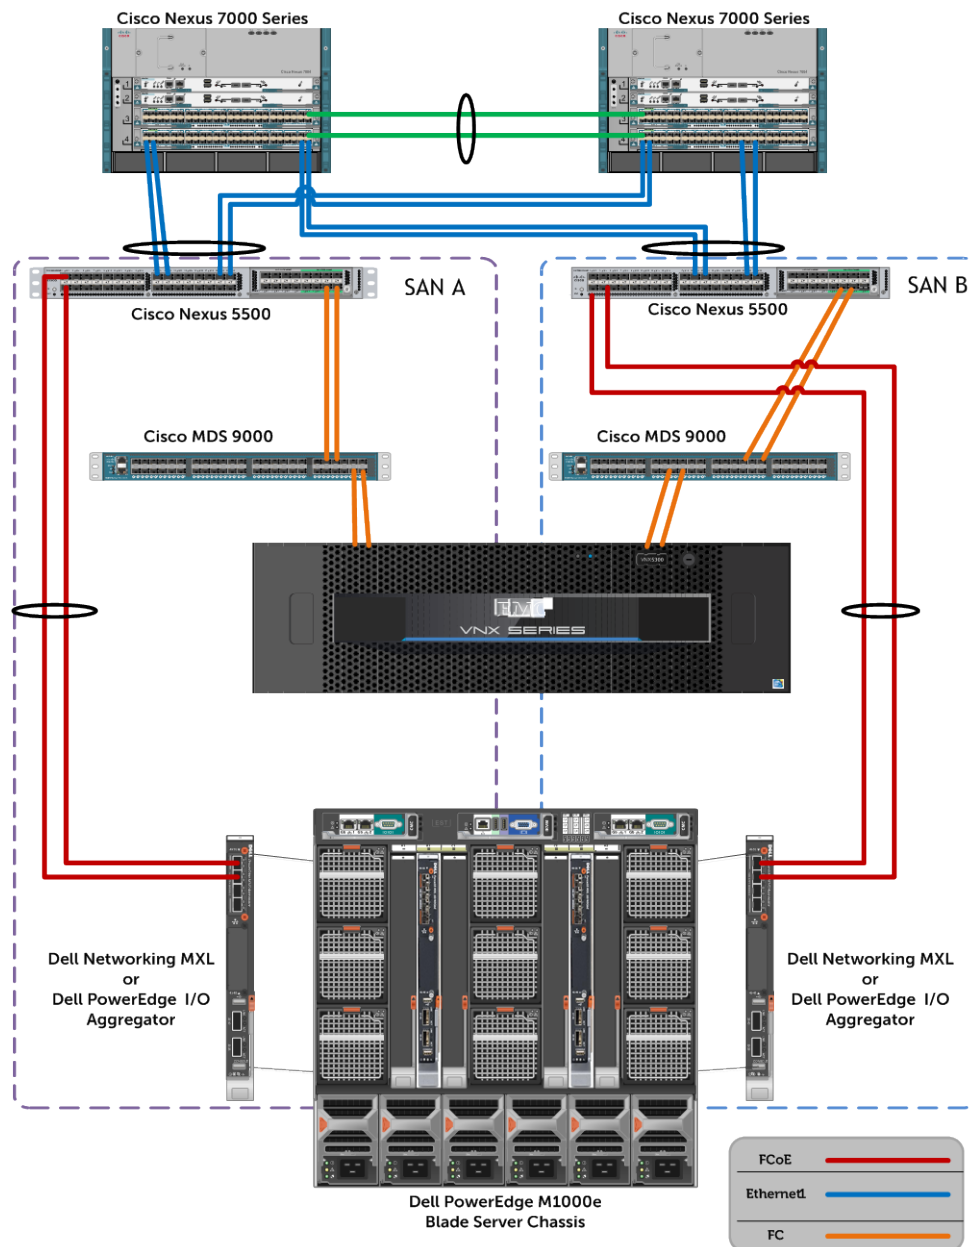


Figure 3 Configuration 2 - NPV with fabric services provided by MDS FC switch

## 2 M1000e I/O Modules (MXL and I/O Aggregator)

The Dell I/O Modules used in this document are the “Dell Networking MXL” and “PowerEdge M I/O Aggregator”. Both are designed with ease of use and common interchangeable Flex IO Expansion Modules.

### 2.1 Flex IO Expansion Modules (External Ports)

The Dell I/O Modules will support a combination of Flex I/O Modules. The four (4) different types of Flex IO expansion modules are:

- 4-port 10Gbase-T FlexIO module (only one 10Gbase-T module can be used)
- 4-port 10Gb SFP+ FlexIO module
- 2-port 40Gb QSFP+ FlexIO module
- 4-port Fiber Channel 8Gb module




Figure 4 FlexIO expansion modules

**NOTE:**

- The 4 Port Fiber Channel 8Gb module can only be used with the release of 9.3 firmware for the Dell modular FSB switches.
- Using the Fiber Channel, 8Gb module to provide non-FIP snooping FCoE configurations will be covered in a separate deployment guide.

## 2.2 I/O Module Internal Ports

The connections are 10 Gigabit Ethernet connections for basic Ethernet traffic, iSCSI storage traffic or FCoE storage traffic. In a typical M1000e configuration of 16 half-height blade servers' ports 1-16 are used and 17 -32 disabled. However if quad port adapters or quarter-height blade servers are used, ports 17-32 will be enabled.

Dell Networking MXL and Dell PowerEdge M I/O Aggregator – Port Mapping					
Expansion Slot 1	QSFP+ 2x40Gb	QSFP+ 8x10GB SFP+ (breakout)	SFP+ 4x10Gb	10G-BaseT 4x10Gb	FC8 x 4
		56			
		55			
		54			
	53	53			
		52	52	52	52
		51	51	51	51
		50	50	50	50
Expansion Slot 0	49	49	49	49	49
	QSFP+ 2 X 40Gb	QSFP+ 8 X 10GB SFP+ (breakout)	SFP+ 4 X10Gb	10G-BaseT 4 X 10Gb	FC8 x 4
		48			
		47			
		46			
	45	45			
		44	44	44	44
		43	43	43	43
Fixed QSFP Ports		42	42	42	42
	41	41	41	41	41
	QSFP+ 2 X 40Gb	QSFP+ 8 X 10GB SFP+ (breakout)	SFP+ 4 X10Gb	10G-BaseT 4 X 10Gb	FC8 x 4
		40			
		39			
		38			
	37	37			
		36	.	.	.
Internal 10/1 Gb		35	.	.	.
		34	.	.	.
	33	33	.	.	.
	Internal 10 / 1 GB interfaces				
	32	32	32	32	32
	31	31	31	31	31
	.	.	.	.	.
	.	.	.	.	.
.	.	.	.	.	
2	2	2	2	2	
1	1	1	1	1	

Table 3 Port Mapping Table



## 2.3 Dell Networking MXL

The MXL 10/40GbE Switch is a layer 2/3 blade switch with two fixed 40GbE ports on the base module and support for two optional plug-in modules. The switch operates in a PowerEdge M1000e Enclosure, which can support up to 32 servers and 6 MXL 10/40GbE Switches. The MXL 10/40GbE Switch runs the Dell Operating System (FTOS), providing switching, bridging, and routing functionality for transmitting data, storage, and server traffic.



Figure 5 Dell Networking MXL

## 2.4 PowerEdge I/O Aggregator

The I/O Aggregator is a zero-touch, layer 2-switch blade with two fixed 40 GB ports on the base module and support for two optional plug-in modules. The Aggregator operates in a Dell PowerEdge M1000e chassis, which can support up to 32 servers and 6 I/O Aggregator blade switches. The Aggregator runs the Dell Operating System (FTOS) and auto configures as an unmanaged switch with bridging and multiplexing functionality,

The IOA acts very similar to a pass-through module providing connectivity between Network adapters internally and external upstream network devices. Internally the IOA provides thirty-two (32) connections.

By default, all VLANs are allowed as well as any DCBx iSCSI or FCoE settings. In addition, the external ports are all part of the same LAG (link aggregation group) which obviates the need for Spanning-tree.

The document will not cover FCoE VLANs at Blade Servers and if required care must be taken to make sure that matching VLANs are configured correctly at upstream Top-Of-Rack switches.



Figure 6 Dell PowerEdge M I/O Aggregator



## 2.5 Factory Default Settings

### 2.5.1 The Dell Power I/O Aggregator

The Dell Power I/O Aggregator is installed in a Dell PowerEdge M1000e Enclosure. For information about how to restore the factory default settings of the Dell Power I/O Aggregator refer to the [Configuration Guide for the M I/O Aggregator Guide](#), located on the Dell Support website [www.support.dell.com](http://www.support.dell.com).

### 2.5.2 The Dell Networking MXL

The Dell Networking MXL is installed in a Dell PowerEdge M1000e Enclosure. For information about how to restore the factory default settings of the Dell Networking MXL refer to the [Configuration Guide for the MXL 10/40GbE Switch I/O Module](#), located on the Dell Support website [www.support.dell.com](http://www.support.dell.com)

CAUTION: Stacking is NOT recommended and will create highly undesirable problems in an FCoE environment with Dell Networking Switches. If the configuration is deployed in this manner proper congestion behavior, lossless Ethernet and reliability cannot be guaranteed.

## 2.6 FIP Snooping

This deployment guide focuses on enabling the data center Dell blade server environment for FCoE via FIP Snooping. The modular switches (MXL and IOA) can be termed as FIP snooping bridges when used in this manner. This provides the proper architecture so that DCB values (QoS values, ACL's, Etc.) are properly configured for iSCSI and/or FCoE to function optimally.

Using the FC FlexIO module that provides 8Gb Fibre Channel interfaces is not covered in this deployment guide. This guide covers the FIP Snooping topologies, which must reach the ToR FCF switch or the fabric switch in order to de-encapsulate Fibre Channel packets.





### 3 M1000e Chassis Blade Servers and Network Adaptor

The first step is to gather configuration information that is needed later in this document. The M1000e chassis contains M620 Blade Server examples. The information gathered in this section will be used throughout the document.




Figure 7 Blade Server - M620 Half Height Blade, with Broadcom Fabric B Mezzanine card




Figure 8 M1000e Chassis with two M I/O Aggregators in Fabric B1 and B2 slots

#### 3.1 WWN/MAC Addresses

At this time, we need to gather the Mac Address from the Mezzanine Cards in our Blade Servers.

1. Login to the Chassis Management Controller (CMC)
2. In the left windowpane, select "**Server Overview**" (See Figure below).
3. Once the Server Overview page populates select "**WWN / MAC**" in the top pane.



- The screen will show all servers MAC addresses. For our configuration, we need the MAC address from **Slot 1** (Blade 1). Scroll down to **Slot 1** (Blade 1) and set the **Filters**. In the **Filter** drop down, select "**Fabric B**" and in the next drop down select "**Fibrechannel**".
- Record the MAC addresses (Server-Assigned or Chassis-Assigned) listed with a check mark. In our example, the first **B1** MAC address is "**20:01:5C:F9:DD:16:EF:07**" and the first **B2** MAC address is "**20:01:5C:F9:DD:16:F0:10**".

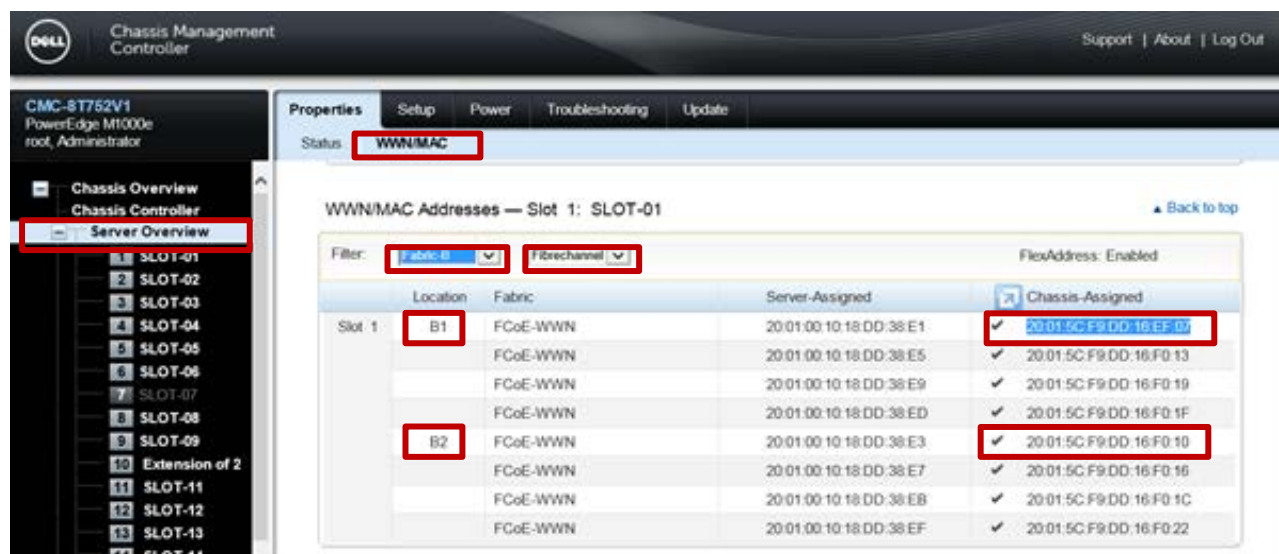


Figure 9 M1000e Chassis Management Controller WWN/MAC Screen (Server in Slot 1)

- The final step is to take the WWPN address drop the first two set of numbers (in bold) "**20:01**:5C:F9:DD:16:EF:07" to get you FIP Mac address (in bold). "**20:01:5C:F9:DD:16:EF:07**".
- The Table below shows the information gathered for two Servers.


		MAC Address B1 (SAN A)	MAC Address B2 (SAN B)
SERVER 1	WWPN (20:01 + FIP MAC)	20:01:5C:F9:DD:16:EF:07	20:01:5C:F9:DD:16:F0:10
	FIP MAC	5C:F9:DD:16:EF:07	5C:F9:DD:16:F0:10
	WWNN (20:00 + FIP MAC)	20:00:5C:F9:DD:16:EF:07	20:00:5C:F9:DD:16:F0:10
SERVER 2	WWPN (20:01 + FIP MAC)	20:01:5C:F9:DD:16:EF:21	20:01:5C:F9:DD:16:F1:7E
	FIP MAC	5C:F9:DD:16:EF:21	5C:F9:DD:16:F1:7E
	WWNN (20:00 + FIP MAC)	20:00:5C:F9:DD:16:EF:21	20:00:5C:F9:DD:16:F1:7E

Table 4 MAC Address Information for 2 servers



## 3.2 Configuring BCM57810 Network Adapter

To configure the Broadcom BCM57810 network adapter the following information is required storage. Storage Processor WWPN and Boot LUN (we used WWPN 50:06:01:6F:3E:E0:18:70 and Boot LUN = 0). In the "Broadcom NetXtreme II™ Network Adapter User Guide", contain detail information on the following steps. ([Link to Document](#)).

Click on the paper clip  icon on the left to see attached files with steps and screen shots.

1. Power On the server
2. Type "**Ctrl-s**" when the Broadcom adaptors Boot Code appears on screen.
3. Highlight the BCM57810 adaptor that will be used for FCoE and press the "**enter**" key
4. You are now in the "Main Menu" to configure the Broadcom adapter.

### 3.2.1 Device Hardware Configuration

In this section, we can configure NPAR, SR-IOV and DCB.

1. From the "**Main Menu**" choose "**Device Hardware Configuration**" and press the "**Enter**" key.
2. Change "Multi-Function Mode" to enable "NPAR" if desired. Note: NPAR (NIC Partitioning) is not required for FCoE.
3. Change "**DCB Protocol**" to "**Enabled**" (disabled by default).
4. Change "SR-IOV" to "Enabled" if desired. Note: SR-IOV is not required for FCoE
5. Press the "**Escape**" key to exit back to the Main Menu..

### 3.2.2 MBA Configuration

1. From the "**Main Menu**" choose "**MBA Configuration**" and press the "**Enter**" key.
2. "**Option ROM**" should be "**Enabled**"
3. "**Boot protocol**" should be set to "**FCoE**"
4. All other settings can be left at default
  - Boot Strap Type = Auto
  - Hide Setup Prompt = Disabled
  - Setup Key Stroke = Ctrl-S
  - Banner Message Timeout = 5 Seconds
  - Link Speed = AutoNeg
  - Pre-boot Wake On LAN = Disabled
  - VLAN Node = Disabled
  - VLAN ID = 1
  - Boot Retry Count = 0
5. Press the "**Escape**" key to exit back to the Main Menu

### 3.2.3 FCoE Boot Configuration

1. From the "**Main Menu**" choose "**FCoE Boot Configuration**" and press the "**Enter**" key.
2. From the "**FCoE Boot Main Menu**", highlight "**General Parameters**" and press the "**Enter**" key

3. Set the **"Boot to FCoE Target"** to **"One Time Disabled"** (See Note)
4. **"Target as First HDD"** should be **"Enabled"**.
5. All other settings can be left at default
  - Link Up Delay Time = 0
  - LUN Busy Retry Count = 0
  - Fabric Discovery Retry = 4
  - HBA Boot Mode = Enabled
6. Press the **"Escape"** key to exit back to the FCoE Boot Main Menu

Note: By setting one-time disabled, this allows the NETWORK ADAPTOR to bypass booting to the LUN (which has no OS installed) and allows the CD-ROM to be the boot device for loading the OS. Please understand that this setting will need re-enabled after each reboot if a further use of the CD-ROM is needed.

7. From the **"FCoE Boot Main Menu"**, highlight **"Target Information"** and press the **"Enter"** key
8. Highlight **"No. 1 Target"** and press the **"enter"** key
9. Change **"Connect"** to **"Enabled"**
10. Highlight **"WWPN"** and press the **"enter"** key to enter your Storage Processor WWPN information (we used 50:06:01:6F:3E:E0:18:70 in our lab configuration).
11. Set **"Boot LUN"** field to **"0"** (default for most arrays),
12. Press the escape key 2 times to exit back to the Main Menu

Note: The best way to determine the WWPN of Storage Processor is to look at the FLOGI database on the FC switch or top-of-rack FCF switch.

### 3.2.4 NIC Partition Configuration

1. From the **"Main Menu"** choose **"NIC Partition Configuration"** and press the **"Enter"** key if you enabled NPAR earlier during **"Device Hardware Configuration"**.
2. Flow Control= Auto
3. Select the first partition **"PF#0"** and press the **"Enter"** key
4. Ethernet protocol should be **"Enabled"**
5. iSCSI Offload Protocol should be **"Disabled"**
6. FCoE Offload Protocol should be **"Enabled"**
7. All other settings can be left at the default settings
8. Press the **"Escape"** key to exit back to the NIC Partition Configuration Menu
9. Enter each of the other partitions (I.E. PG#2, PF#4 and PF#6) and disable the protocols (Ethernet, iSCSI, and FCoE) for this initial setup.

Note: By disabling the other protocols this will ensure that, only one path will be presented to the Array. It will also ensure easier trouble shooting. However for the other partitions to function they must be reconfigured later for Ethernet or applicable protocol

10. Press the **"Escape"** key until you get back to the **"Exit Configuration Screen"** and select **"Exit and Save Configuration"**.



11. The Broadcom menu should reappear for any further configuration tasks
12. If no other menu configuration tasks are needed hit escape and the system will reboot

### 3.2.5 Boot Sequence

1. During POST press F2 to enter Bios
2. Select "System BIOS"
3. Select "Boot Settings"
4. Select "BIOS Boot Settings"
5. Select "Boot Sequence"
6. Highlight NIC and click on the plus sign button to move it to the top position
7. Highlight CD-ROM or virtual CD-ROM whichever will be used and click on plus button to move it to the second position below NIC
8. Select back and finish until a applet pops up to confirm exit
9. Click yes to exit
10. At this point, the boot sequence is set to be applicable for "Boot from SAN"


### 3.2.6 Booting after NETWORK ADAPTOR configuration

1. At this point, the Broadcom FCoE boot screen should appear and a succeeded should be seen for each step of the identification and login process.
2. The CD-ROM should be the next boot device found and the OS should start its initial load
3. Load latest Broadcom drivers during install; First driver is evbd.sys
4. Then load second driver fcoe.sys
5. This will allow windows to recognize remote LUN and start the loading of the OS



### 3.3 Configuring QME8262 Network Adapter

To configure the Qlogic "QME8262" network adapter the following information is required. 1) Storage Processor target WWPN = (we used 50:06:01:6F:3E:E0:18:70) and Boot LUN number (we used 0). We will configure the Network Adapter through the "Life Cycle Controller" pages. The attachments show the screen shots for this section.

Click on the paper clip  icon on the left to see attached files with steps and screen shots.

1. Power On the server
2. When you see the Dell Logo press "**F10**" to enter the "Life Cycle Controller"
3. On the Left side of the screen, highlight the "**System Setup**" then select "**Advanced Hardware Configuration**".
4. On the System Setup Main Menu select "**Device Settings**"
5. From the Device Setting Page select the FCoE QME8262 adaptor. In our configuration, we are selecting the "**NIC in Mezzanine 2B Port 1**" Device
6. From the Main Configuration Page select "**NIC Partitioning Configuration**".
7. In the NIC Partitioning Configuration page select "**Partition 4 Configuration**"
8. In the Partition 4 Configuration, "**Enable**" FCoE Mode.
9. Press the "**Back**" Button two times to get you back to the Main Page
10. From the Main Page you can now select "**FCoE Configuration**"
11. In the FCoE Configuration Page make sure "Connect" is "**Enabled**"
12. Enter in your "**Boot LUN**" (For our configurations we are using LUN 0)
13. Enter in your "**World Wide Port Name Target**" (We are using 50:06:01:6F:3E:E0:18:70)
14. You can now select the "**Back**" button to finish the configuration and save.
15. Finish and exit until the server restarts.



## 4 Configuration 1 – Dell Networking MXL / PowerEdge M I/O Aggregator (FSBs) with Cisco Nexus fabric mode topology

The Cisco Nexus Top of Rack is many times left with the default fabric mode, which allows the Nexus switch to perform as a fabric services provider and Fiber Channel switch. The following is a default FCoE with FIP-Snooping Bridge configured in a converged network environment. Configurations for both the Primary (SAN A) and Secondary (SAN B) are provided. The main differences are with FCoE VLANs, VSANs and ultimately connecting to the same LUN in a server multipath Environment.

Note: An LACP Port Channel must be configured on the Cisco Nexus connected ports for the link to become active on the Dell PowerEdge M I/O Aggregator even if it is just a single port connection.

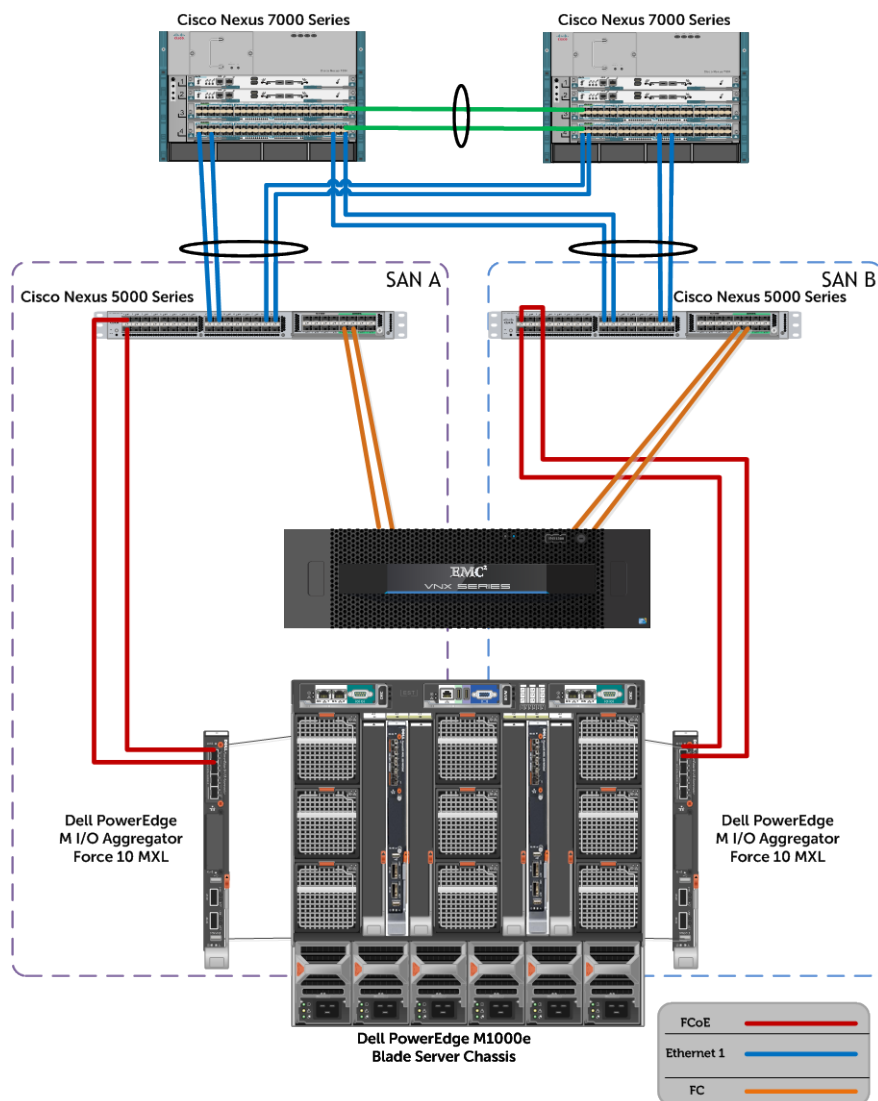


Figure 10 Configuration 1 - NPIV

## 4.1 Setup Cisco Nexus 5000 series switch (FCF)

The Cisco Nexus 5000 series is a family of datacenter “top-of-rack” switch/fabric network devices. In this document, the Cisco Nexus 5548UP switch is the “Top-of-Rack (ToR) switch used during the creation of this document.

The complexity involved with FCoE configurations will be quite confusing to many datacenters or businesses that wish to implement this type of storage solution. It is important to understand where and how the key parts of the solution are configured and why the configuration steps are done in this manner. Much of the configuration will rely on properly configuring the Cisco Nexus 5548UP switch, as it is the primary configuration point for the rest of the solution. The modular switches in the blade server chassis will pass DCB information from this “ToR” switch down to the server’s network adapters. The servers network adapter will take DCB values presented by the ToR “switch, adjust to match the values so FCoE will function optimally.

This document utilizes command line interface (CLI) commands to configure the devices. It should be noted that Cisco supplies various graphical interfaces for managing their equipment and this method may be easier for administrators.

The following tables/pages show the steps taken to configure the Nexus 5548UP “ToR” switch.



## Enable features, configure all pre-planned VSAN's, VLAN's, and VFC's (see planning matrix) - Side 1 (SAN A)

1. Enable "feature fcoe" to allow the FC ports to function.
2. Feature npiv allows multiple device logins through the same physical port
3. Feature lacp enables the LACP protocol for port-channels (LAGs)

```
feature fcoe
feature npiv
feature lacp
```

Command line interface

4. Create a new VSAN - instantiate it in the VSAN database.
5. Configure regular ethernet VLANs, and then the FCoE VLAN is created with an assignment to its respective VSAN
6. Instantiate but do not configure the upstream port-channel (LAG) to the core /aggregation switch
7. Instantiate but do not configure the downstream port-channel (LAG) to the IOA

```
vsan database
vsan 2
vlan 20,30-32, 88
vlan 1000
fcoe vsan 2
interface port-channel 1
interface port-channel 20
```

Command line interface

8. Create the VFC interface to bind to the servers CNA FIP MAC address. This can be located in the CMC WWN table or the IDRAC page for the server. (in this example 2 different server's CNA's are configured)

```
interface vfc101
bind mac-address 5C:F9:DD:16:EF:07
no shutdown
interface vfc102
bind mac-address 5C:F9:DD:16:EF:21
no shutdown
```

Command line interface

9. Move back into the VSAN database and create entries for the new VFC just created and create entries for the FC port(s) that will be used.

```
vsan database
vsan 2 interface vfc101
vsan 2 interface vfc102
vsan 2 interface fc2/1
vsan 2 interface fc2/2
```

Command line interface

## Enable features, configure all pre-planned VSAN's, VLAN's, and VFC's (see planning matrix) - Side 2 (SAN B)

1. Enable "feature fcoe" to allow the FC ports to function.
2. Feature npiv allows multiple device logins through the same physical port
3. Feature lacp enables the LACP protocol for port-channels (LAGs)

```
feature fcoe
feature npiv
feature lacp
```

Command line interface

4. Create a new VSAN - instantiate it in the VSAN database.
5. Configure regular ethernet VLANs, and then the FCoE VLAN is created with an assignment to its respective VSAN
6. Instantiate but do not configure the upstream port-channel (LAG) to the core /aggregation switch
7. Instantiate but do not configure the downstream port-channel (LAG) to the IOA

```
vsan database
vsan 3
vlan 21,30-32,88
vlan 1001
fcoe vsan 3
interface port-channel 2
interface port-channel 21
```

Command line interface

8. Create the VFC interface to bind to the servers CNA FIP MAC address. This can be located in the CMC WWN table or the IDRAC page for the server. (in this example 2 different server's CNA's are configured)

```
interface vfc201
bind mac-address 5C:F9:DD:16:F0:10
no shutdown
interface vfc202
bind mac-address 5C:F9:DD:16:F1:7E
no shutdown
```

Command line interface

9. Move back into the VSAN database and create entries for the new VFC just created and create entries for the FC port(s) that will be used.

```
vsan database
vsan 3 interface vfc201
vsan 3 interface vfc202
vsan 3 interface fc2/1
vsan 3 interface fc2/2
```

Command line interface

1<sup>st</sup> (SAN A) Cisco Nexus 5548 Configuration

2<sup>nd</sup> (SAN B) Cisco Nexus 5548 Configuration





**Side 1 (SAN A)**

10. Configure the needed port-channels/LAGs/Trunks.
- There will be one going upstream to the core/aggregation switch, and one going downstream to the IOA/MXL.

Command line interface

```
interface ethernet 1/1-2
channel-group 20 mode active
description FCoE_downlink_to_IOA-MXL
interface ethernet 1/9-10
channel-group 1 mode active
description Ethernet_uplink_to_7K
```

11. Configure the port-channels with the applicable settings which will then automatically apply to the individual ethernet interfaces if done in this order.
- VLAN 30-32, and 88 are example ethernet VLAN's just to show how these are applied in this configuration.
  - VLAN 20 is the native VLAN applicable to the port-channel that is being used for FCoE, however the default VLAN of the switch can be used in this case and this is not required to be configured.
  - VLAN 1000 is the FCoE VLAN and this must be configured for FCoE traffic to traverse from the FSB to the Nexus 5k and then to storage.

Command line interface

```
interface port-channel 1
description port-channel_eth9+10_to_7k
switchport mode trunk
switchport trunk allowed vlan 30-32,88

interface port-channel 20
description port-
channel_eth1+2_to_IOA-MXL
switchport mode trunk
switchport trunk native vlan 20
switchport trunk allowed vlan 20,1000
```

12. Turn on or enable the FC ports that are applicable to this configuration.

Command line interface

```
interface fc2/1-2
no shutdown
```

**Side 2 (SAN B)**

10. Configure the needed port-channels (LAGs)/Trunks.
- There will be one going upstream to the core/aggregation switch, and one going downstream to the IOA/MXL.

Command line interface

```
interface ethernet 1/1-2
channel-group 21 mode active
description FCoE_downlink_to_IOA-MXL
interface ethernet 1/9-10
channel-group 2 mode active
description Ethernet_uplink_to_7K
```

11. Configure the port-channels with the applicable settings which will then automatically apply to the individual ethernet interfaces if done in this order.
- VLAN 30-32, and 88 are example ethernet VLAN's just to show how these are applied in this configuration.
  - VLAN 21 is the native VLAN applicable to the port-channel that is being used for FCoE, however the default VLAN of the switch can be used in this case and this is not required to be configured.
  - VLAN 1001 is the FCoE VLAN and this must be configured for FCoE traffic to traverse from the FSB to the Nexus 5k and then to storage.

Command line interface

```
interface port-channel 2
description port-channel_eth9+10_to_7k
switchport mode trunk
switchport trunk allowed vlan 30-32,88

interface port-channel 21
description port-
channel_eth1+2_to_IOA-MXL
switchport mode trunk
switchport trunk native vlan 21
switchport trunk allowed vlan 21,1001
```

12. Turn on or enable the FC ports that are applicable to this configuration.

Command line interface

```
interface fc2/1-2
no shutdown
```



## Configure Zoning with Cisco Nexus in switch mode (NPIV)- Side 1 (SAN A)

**13.** Configure applicable zones that will match the blade server CNA's (2 are used in this example) to applicable FC port(s) going to storage.

- "zone1SAN\_A" is the example name being used here but this name is at the discretion of the administrator
- pwwn <hh...hh> is applicable to the PWWN of the CNA

```
Command line interface
zone name zone1SAN_A vsan 2
member pwwn <20:01:5c:f9:dd:16:ef:07>
member pwwn <20:01:5c:f9:dd:16:ef:21>
member interface fc2/1
member interface fc2/2
```

**14.** Create a zoneset that will group all the applicable zones together and allow for the group activation.

- "set1SAN\_A" is just an example name for this configuration

**15.** Activate the zoneset

```
Command line interface
zoneset name set1SAN_A vsan 2
member zone1SAN_A

zoneset activate name set1SAN_A vsan 2
```

- Check that the applicable zoneset and zones are now active

```
Command line interface
Show zoneset active
```

```
results
Zoneset name set1SAN_A vsan 2
zone name zone1SAN_A vsan 2
* fcid 0x850000 [interface pwwn
20:01:5c:f9:dd:16:ef:07]
* fcid 0x850100 [interface pwwn
20:01:5c:f9:dd:16:ef:21]
* fcid 0x850200 [interface fc2/1 swrn
20:41:54:7f:ee:53:3e:80]
* fcid 0x850300 [interface fc2/1 swrn
20:42:54:7f:ee:53:3e:80]
```

## Configure Zoning with Cisco Nexus in switch mode (NPIV)- Side 2 (SAN B)

**13.** Configure applicable zones that will match the blade server CNA's (2 are used in this example) to applicable FC port(s) going to storage.

- "zone1SAN\_B" is the example name being used here but this name is at the discretion of the administrator
- pwwn <hh...hh> is applicable to the PWWN of the CNA

```
Command line interface
zone name zone1SAN_B vsan 3
member pwwn <20:01:5c:f9:dd:16:f0:10>
member pwwn <20:01:5c:f9:dd:16:f1:7e>
member interface fc2/1
member interface fc2/2
```

**14.** Create a zoneset that will group all the applicable zones together and allow for the group activation.

- "set1SAN\_B" is just an example name for this configuration

**15.** Activate the zoneset

```
Command line interface
zoneset name set1SAN_B vsan 3
member zone1SAN_B

zoneset activate name set1SAN_B vsan 3
```

- Check that the applicable zoneset and zones are now active

```
Command line interface
Show zoneset active
```

```
results
Zoneset name set1SAN_B vsan 3
zone name zone1SAN_A vsan 3
* fcid 0x850000 [interface pwwn
20:01:5c:f9:dd:16:f0:10]
* fcid 0x850100 [interface pwwn
20:01:5c:f9:dd:16:f1:7e]
* fcid 0x850200 [interface fc2/1 swrn
20:41:54:7f:ee:56:55:40]
* fcid 0x850300 [interface fc2/1 swrn
20:42:54:7f:ee:56:55:40]
```

Figure 11 Cisco Nexus 5000 series configuration sequence – multiple links

## 4.2 Dell PowerEdge M I/O Aggregator and Networking MXL modular FSBs - Configuration in a fabric mode topology

The PowerEdge M I/O Aggregator (IOA) allows for an extremely easy addition to the network with simplified management and an increase to server bandwidth. The IOA enables easy, plug-and-play data center convergence with near zero configuration needed in most standard infrastructures. The MXL is a high-end, layer 2 / layer 3, highly configurable modular switch that can be setup to fit in all infrastructure needs. For ease of use in a standard environment, the IOA is the FSB of choice, but for those environments that need more manually configured advanced settings the MXL will surpass expectations.

It is necessary to understand that the IOA can have some key settings manually adjusted to meet a few of the non-standard needs but for the network administrator that wants full control of a modular switch the MXL should be used.

In the topology, scenarios captured in this document the modular switch discussed will be focused around the Dell Networking MXL with FSB support. This switch allows manual configuration of all converged networking settings. The IOA comes preconfigured to enable DCB and converged networking in the standard environments as previously discussed.



Figure 12 PowerEdge M I/O Aggregator modular 10/40 Gb port aggregator (FSB)



Figure 13 Dell Networking MXL modular 10/40 Gb switch (FSB)

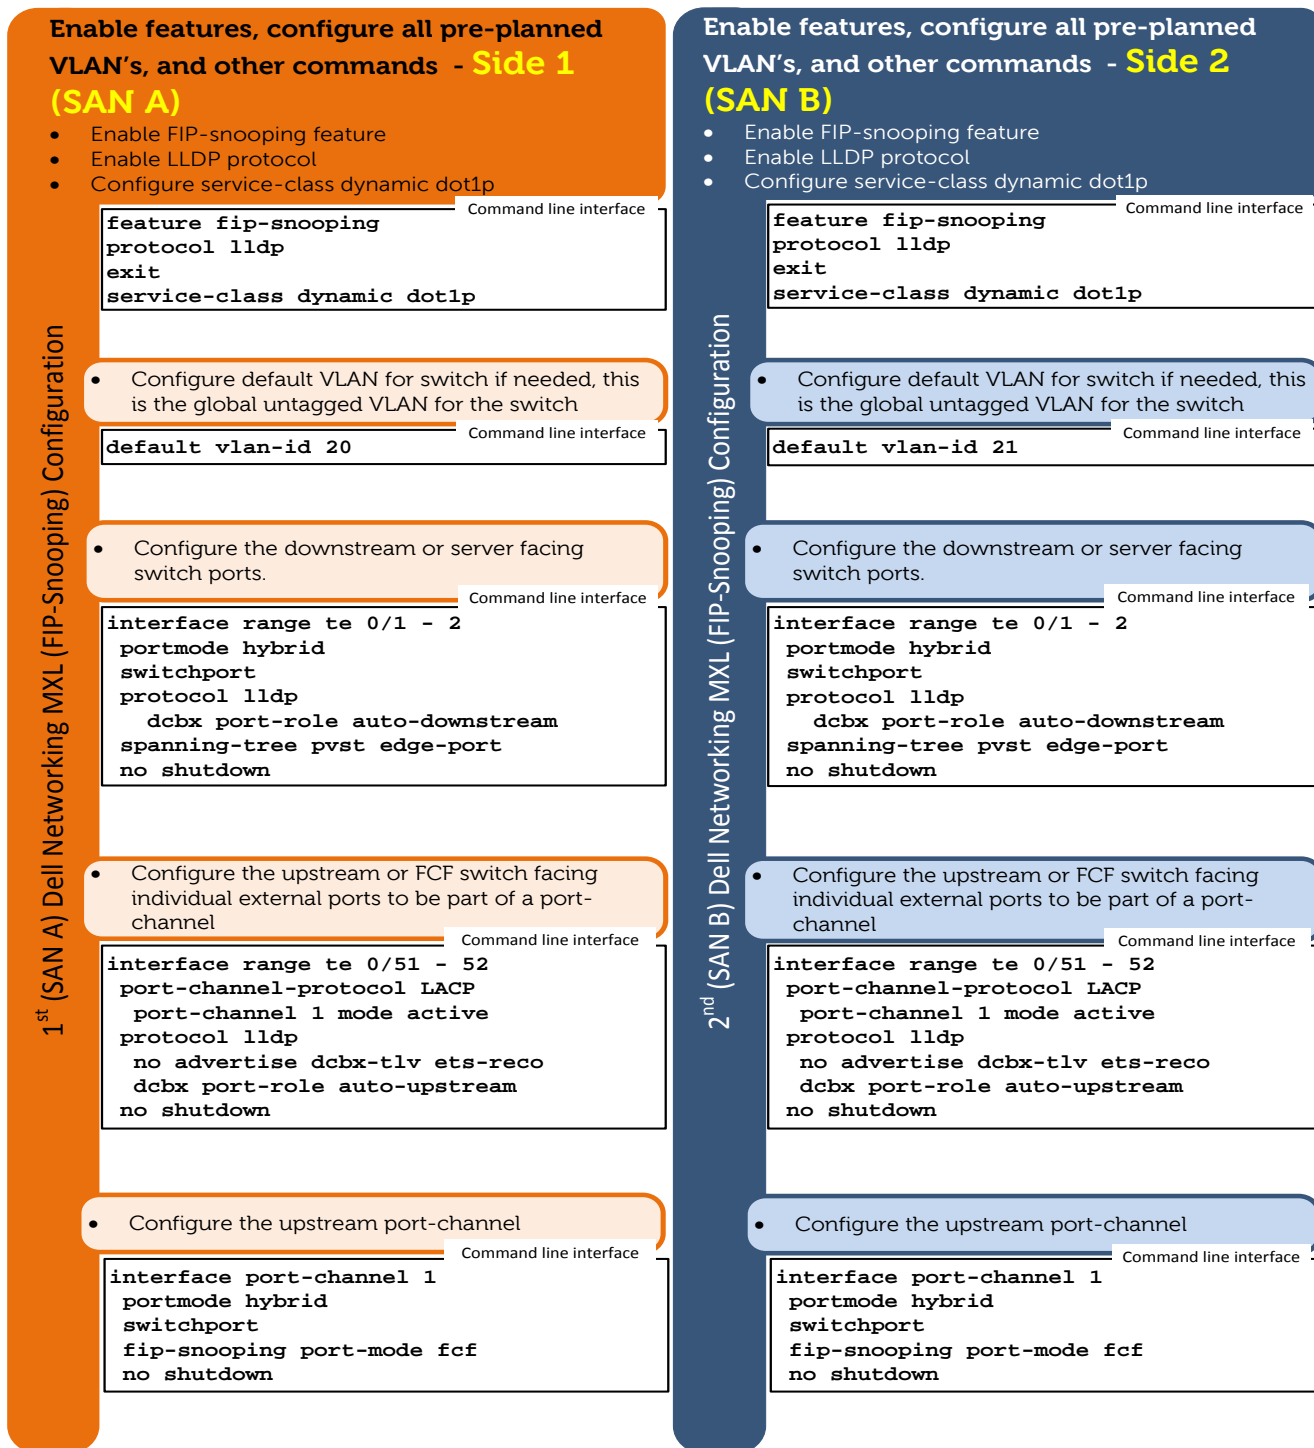


Figure 14 Dell Networking MXL modular switch configuration

## 5 Configuration 2 – Dell Networking MXL/PowerEdge M I/O Aggregator (FSBs) with Cisco Nexus NPV mode topology

The Cisco Nexus 5548UP Top of Rack switch is configured a majority of the time in an aggregation mode or NPV to simply pass Fiber Channel (FC) out to another terminating switch (in our example the Cisco MDS 9148). The following pictures and examples describe a two link LAG from an IOA to the Cisco 5548UP Top of Rack switch configured to NPV mode. The uplink configuration to further aggregation or core switches (such as the pictured Cisco Nexus 7004) will also be captured in the configuration

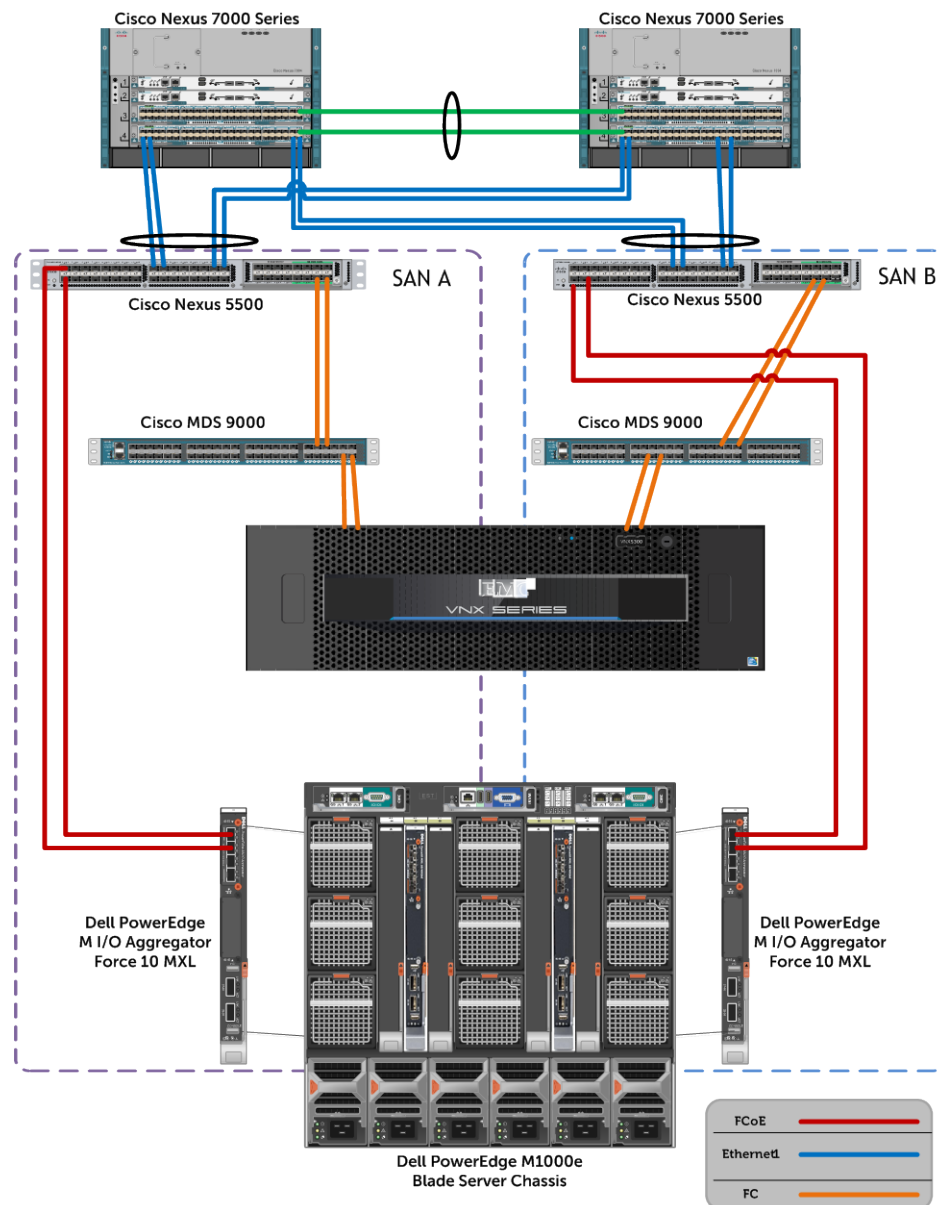


Figure 15 Configuration 2 – NPV with MDS

## 5.1 Setup Cisco Nexus 5000 series and Cisco MDS 9100 series switches in an NPV mode topology

The Cisco Nexus 5000 series switches operate by default in a normal FCF switching mode that is called "fabric mode" or in some instances depending on location of information "NPIV" (N port identifier virtualization) mode. There is a disadvantage to running the topology in this mode in a large datacenter with many edge FC switches. Each edge FC switch will need to consume a domain ID and the limit to the amount of domain ID's on the same SAN or VSAN is 239. This number can be kept down to a manageable number by having many of the edge devices act as aggregators or "NPV" (N Port virtualization) mode.

With the Cisco Nexus configured for "NPV" mode, the switch will not provide the essential fabric services yet it will pass these services through to end devices from an upstream fabric services core/aggregation device. Typically, in a Cisco environment this upstream device will be a Cisco MDS multilayer fabric switch operating in default fabric mode.



Enable features, configure all pre-planned VSAN's, VLAN's, and VFC's (see planning matrix) - **Side 1 (SAN A) - NPV**

1. For NPV configuration with FC SAN switches the feature NPV must be set and this will cause the configuration and switch to reload.

```
feature npv
```

Command line interface

2. Enable "feature fcoe" to allow the FC ports to function.
3. Configure "feature npiv" which allows multiple device logins through the same physical port.
4. Configure "feature lacp" which enables the LACP protocol for port-channels (LAGs)

```
feature fcoe
feature npiv
feature lacp
```

Command line interface

5. Create a new VSAN - instantiate it in the VSAN database.
6. Configure regular ethernet VLANs, and then the FCoE VLAN is created with an assignment to its respective VSAN
7. Instantiate but do not configure the upstream port-channel (LAG) to the core /aggregation switch
8. Instantiate but do not configure the downstream port-channel (LAG) to the IOA

```
vsan database
  vsan 2
vlan 20,30-32,88
vlan 1000
  fcoe vsan 2
interface port-channel 1
interface port-channel 20
```

Command line interface

9. Create the VFC interface to bind to the servers CNA FIP MAC address. This can be located in the CMC WWN table or the IDRAC page for the server. (in this example 2 different server's CNA's are configured)

```
interface vfc101
  bind mac-address 5C:F9:DD:16:EF:07
  no shutdown
interface vfc102
  bind mac-address 5C:F9:DD:16:EF:21
  no shutdown
```

Command line interface

Enable features, configure all pre-planned VSAN's, VLAN's, and VFC's (see planning matrix) - **Side 2 (SAN B) - NPV**

1. For NPV configuration with FC SAN switches the feature NPV must be set and this will cause the configuration and switch to reload.

```
feature npv
```

Command line interface

2. Enable "feature fcoe" to allow the FC ports to function.
3. Configure "feature npiv" which allows multiple device logins through the same physical port.
4. Configure "feature lacp" which enables the LACP protocol for port-channels (LAGs)

```
feature fcoe
feature npiv
feature lacp
```

Command line interface

5. Create a new VSAN - instantiate it in the VSAN database.
6. Configure regular ethernet VLANs, and then the FCoE VLAN is created with an assignment to its respective VSAN
7. Instantiate but do not configure the upstream port-channel (LAG) to the core /aggregation switch
8. Instantiate but do not configure the downstream port-channel (LAG) to the IOA

```
vsan database
  vsan 3
vlan 21,30-32,88
vlan 1001
  fcoe vsan 3
interface port-channel 2
interface port-channel 21
```

Command line interface

9. Create the VFC interface to bind to the servers CNA FIP MAC address. This can be located in the CMC WWN table or the IDRAC page for the server. (in this example 2 different server's CNA's are configured)

```
interface vfc201
  bind mac-address 5C:F9:DD:16:F0:10
  no shutdown
interface vfc202
  bind mac-address 5C:F9:DD:16:F1:7E
  no shutdown
```

Command line interface



**Side 1 (SAN A)**

**10.** Move back into the VSAN database and create entries for the new VFC just created and create entries for the FC port(s) that will be used.

```
vsan database
vsan 2 interface vfc101
vsan 2 interface vfc102
vsan 2 interface fc2/1
vsan 2 interface fc2/2
```

Command line interface

**11.** Configure the needed port-channels (LAGs)/Trunks. There will be one going upstream to the core/aggregation switch, and one going downstream to the IOA/MXL.

```
interface ethernet 1/1-2
channel-group 20 mode active
description FCoE_downlink_to_IOA-MXL1

interface ethernet 1/9-10
channel-group 1 mode active
description Ethernet_uplink_to_7K1
```

Command line interface

**12.** Configure the port-channels with the applicable settings which will then automatically apply to the individual ethernet interfaces if done in this order.

- VLAN 30-32, and 88 are example ethernet VLAN's
- VLAN 20 is the native VLAN
- VLAN 1000 is the FCoE VLAN and this must be configured for FCoE traffic to traverse from the FSB to the Nexus 5k and then to storage.

```
interface port-channel 1
description port-channel_eth9+10_to_7k
switchport mode trunk
switchport trunk allowed vlan 30-32,88

interface port-channel 20
description port-channel_eth1+2_to_IOA-MXL
switchport mode trunk
switchport trunk native vlan 20
switchport trunk allowed vlan 20,1000
```

Command line interface

- Turn on or enable the FC ports that are applicable to this configuration.

```
interface fc2/1-2
no shutdown
```

Command line interface

**Side 2 (SAN B)**

**10.** Move back into the VSAN database and create entries for the new VFC just created and create entries for the FC port(s) that will be used.

```
vsan database
vsan 3 interface vfc201
vsan 3 interface vfc202
vsan 3 interface fc2/1
vsan 3 interface fc2/2
```

Command line interface

**11.** Configure the needed port-channels (LAGs)/Trunks. There will be one going upstream to the core/aggregation switch, and one going downstream to the IOA/MXL.

```
interface ethernet 1/1-2
channel-group 21 mode active
description FCoE_downlink_to_IOA-MXL2

interface ethernet 1/9-10
channel-group 2 mode active
description Ethernet_uplink_to_7K2
```

Command line interface

**12.** Configure the port-channels with the applicable settings which will then automatically apply to the individual ethernet interfaces if done in this order.

- VLAN 30-32, and 88 are example ethernet VLAN's
- VLAN 21 is the native VLAN
- VLAN 1001 is the FCoE VLAN and this must be configured for FCoE traffic to traverse from the FSB to the Nexus 5k and then to storage.

```
interface port-channel 2
description port-channel_eth9+10_to_7k
switchport mode trunk
switchport trunk allowed vlan 30-32,88

interface port-channel 21
description port-channel_eth1+2_to_IOA-MXL
switchport mode trunk
switchport trunk native vlan 21
switchport trunk allowed vlan 21,1001
```

Command line interface

- Turn on or enable the FC ports that are applicable to this configuration.

```
interface fc2/1-2
no shutdown
```

Command line interface

Figure 16 Cisco Nexus 5548 NPV configuration sequence



## Configure Cisco MDS 9100 Series **Side 1** (SAN A)

### 1. Enable feature NPV

```
feature npiv
```

Command line interface

### 2. Create entries in the VSAN database to instantiate the VSAN and then add needed interfaces to that VSAN

```
vsan database
vsan 2
vsan 2 interface fc1/1-2
vsan 2 interface fc1/13-14
```

Command line interface

### 3. Configure zones

- The zone name used in this example is Blade1And2-SAN\_A. This can be any name desired that makes troubleshooting and management easy.

```
zone name Blade1And2-SAN_A vsan 2
member interface fc1/1-2
member interface fc1/13-14
```

Command line interface

### 4. Create zoneset and activate

- The zoneset is a larger container that can hold several zones to be activated as a group. This can be named any name desired for management and troubleshooting purposes. In this example set1-SAN\_A is the name used.

```
zoneset name set1-SAN_A vsan 2
member Blade1And2-SAN_A vsan 2

zoneset activate name set1-SAN_A vsan 2
```

Command line interface

- Check that the zoneset and zones are active.

```
show zoneset active
```

Command line interface

```
zoneset name set1-SAN_A vsan 2
zone name Blade1And2-SAN_A vsan 2
* fcid 0x850000 [interface fc1/1 swwn
20:00:54:7f:ee:50:45:e8]
* fcid 0x850001 [interface fc1/2 swwn
20:00:54:7f:ee:50:45:e8]
* fcid 0x850200 [interface fc1/13 swwn
20:00:54:7f:ee:50:45:e8]
* fcid 0x850201 [interface fc1/14 swwn
20:00:54:7f:ee:50:45:e8]
```

results

## Configure Cisco MDS 9100 Series **Side 2** (SAN B)

### 1. Enable feature NPV

```
feature npiv
```

Command line interface

### 2. Create entries in the VSAN database to instantiate the VSAN and then add needed interfaces to that VSAN

```
vsan database
vsan 3
vsan 3 interface fc1/1-2
vsan 3 interface fc1/13-14
```

Command line interface

### 3. Configure zones

- The zone name used in this example is Blade1And2-SAN\_B. This can be any name desired that makes troubleshooting and management easy.

```
zone name Blade1And2-SAN_B vsan 3
member interface fc1/1-2
member interface fc1/13-14
```

Command line interface

### 4. Create zoneset and activate

- The zoneset is a larger container that can hold several zones to be activated as a group. This can be named any name desired for management and troubleshooting purposes. In this example set1 is the name used.

```
zoneset name set1-SAN_B vsan 3
member Blade1And2-SAN_B vsan 3

zoneset activate name set1-SAN_B vsan 3
```

Command line interface

- Check that the zoneset and zones are active.

```
show zoneset active
```

Command line interface

```
zoneset name set1-SAN_B vsan 3
zone name Blade1And2-SAN_B vsan 3
* fcid 0x380000 [interface fc1/1 swwn
20:00:54:7f:ee:cb:0f:60]
* fcid 0x380001 [interface fc1/2 swwn
20:00:54:7f:ee:cb:0f:60]
* fcid 0x380100 [interface fc1/13 swwn
20:00:54:7f:ee:cb:0f:60]
* fcid 0x380101 [interface fc1/14 swwn
20:00:54:7f:ee:cb:0f:60]
```

results

1<sup>st</sup> (SAN A) Cisco MDS 9000 Series Configuration

2<sup>nd</sup> (SAN B) Cisco MDS 9000 Series Configuration

Figure 17 Cisco MDS 9148 configuration



## 5.2 Dell Networking MXL (FSB) steps in Nexus NPV topology

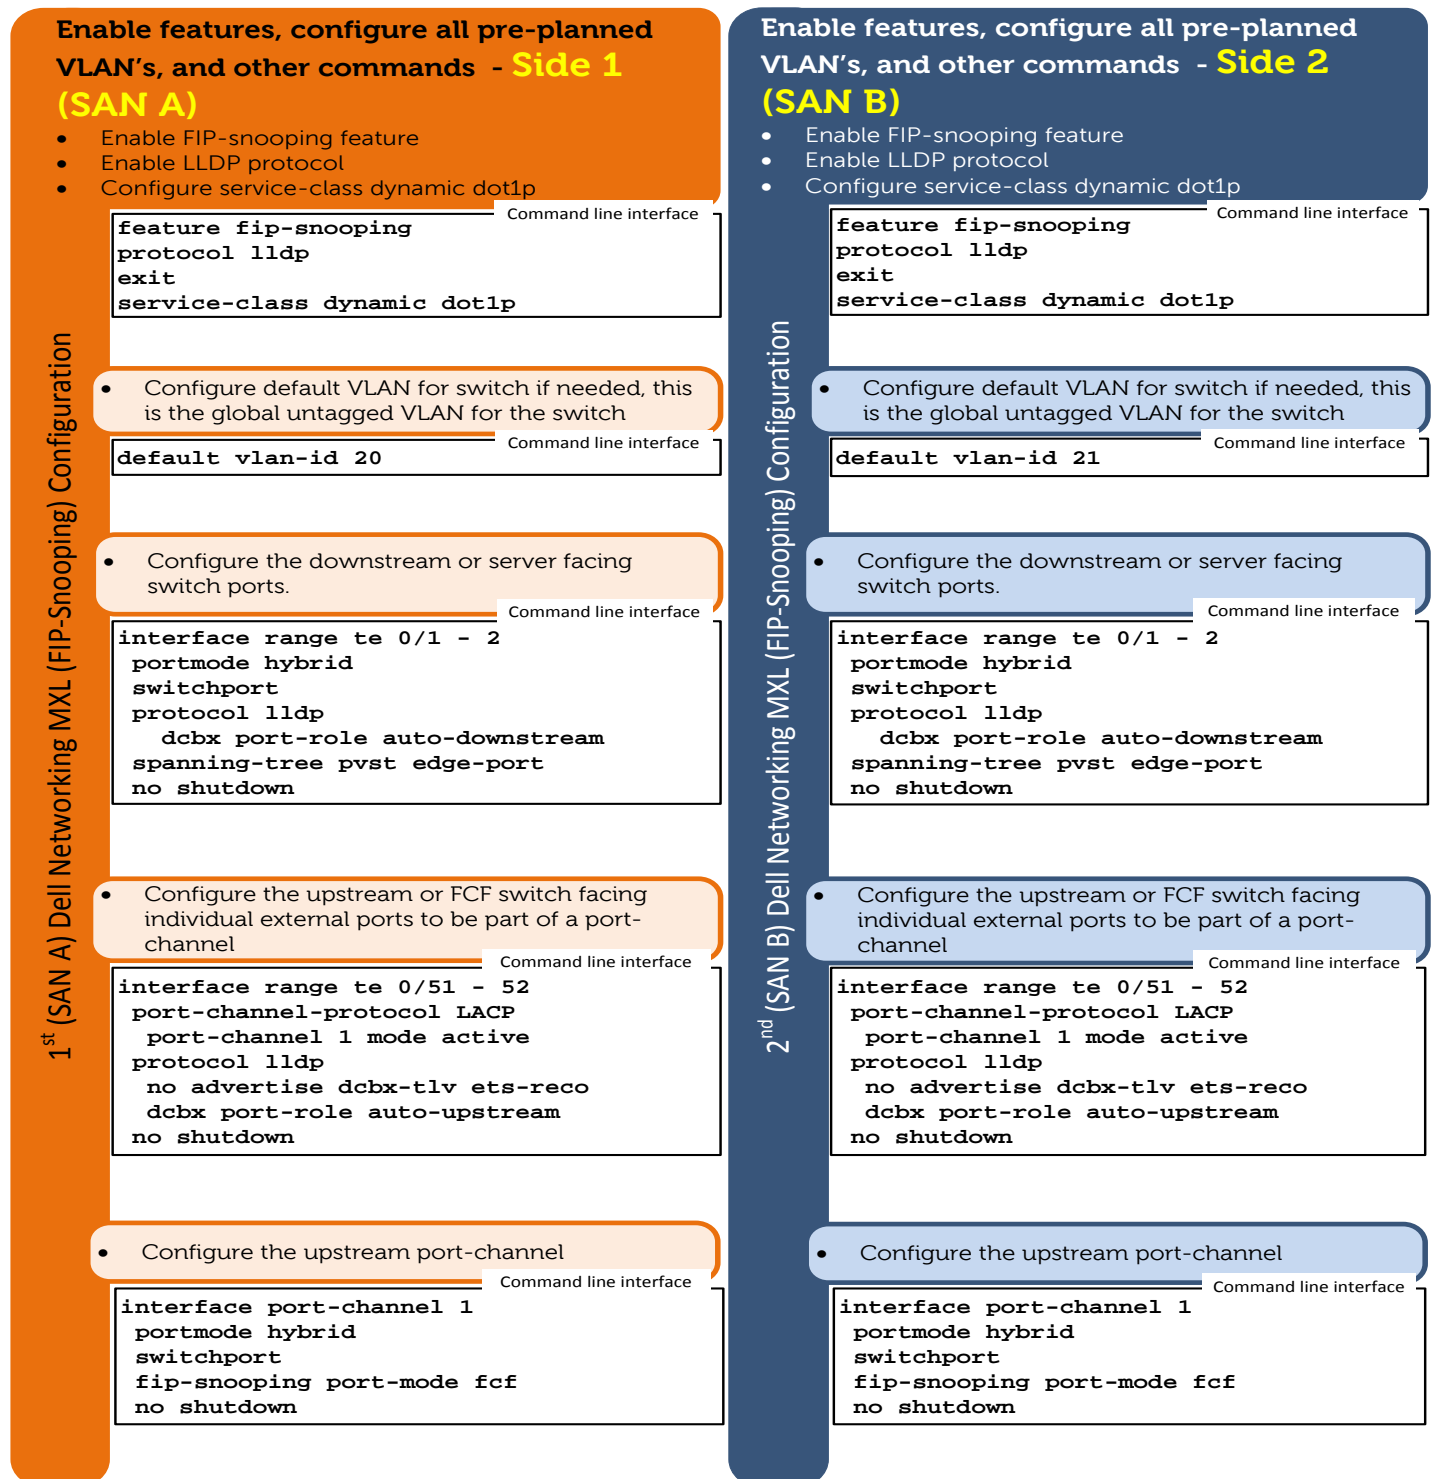


Figure 18 Dell Networking MXL modular switch configuration steps



## 6 EMC setup and assigning a LUN to Initiator

The following steps will configure the EMC storage device (EMC VNX 5300 for this example)

### 6.1 Broadcom CTRL+R

The "CTRL+R" method allows you to use the boot initiator to bring up the link and login into all available fabrics and targets. Using this method, you can ensure that the initiator was logged into the fabric/target before making provisioning changes, and as such, can provision without manually typing in WWPNs.

1. Configure at least one boot target through CCM (CTRL-S) as described above.
2. Allow the system to attempt to boot through the selected initiator.
3. Once the initiator boot starts, it will commence with DCBX sync, FIP Discovery, Fabric Login, Target Login, and LUN readiness checks. As each of these phases completes, if the initiator is unable to proceed to the next phase, MBA will present the option to press CTRL+R.
4. Once "CTRL+R" has been activated, the boot initiator will maintain a link in whatever phase has most recently succeeded and allow you time to make the necessary provisioning corrections to proceed to the next phase.
5. If the initiator logs into the fabric, but is unable to log into the target, a "CTRL+R" will pause the boot process and allow you to configure fabric zoning.
6. Once zoning is complete, the initiator will automatically log into all visible targets. If the initiator is unable to discover the designated LUN on the designated target as provisioned in step 1, "CTRL+R" will pause the boot process and allow you to configure selective LUN presentation.
7. The boot initiator will periodically poll the LUN for readiness, and once the user has provisioned access to the LUN, the boot process will automatically proceed.

NOTE: This does not preclude the need to put the boot initiator into one-time disabled mode

### 6.2 Initial initiator confirmation

1. Confirm EMC storage see the initiators.
2. Create new Host and Register initiator with storage device, typically done same time on first initiator.
3. Create required number of LUNs
4. Next Create storage group, this will be listed under the Host Menu
5. Add Host and LUNs to new storage group.
  - a. Confirm EMC storage sees the initiators
  - b. From tab selection Menu , select "System List", EMC VNX 5300 link
  - c. A set of tab selections will be presented, select Host tab > initiators
  - d. Look for initiator in list, should show as not registered. Select "Not Registered" from Connection status to help filter initiators. It is also possible to use the last four digits of the WWPN of the network adapter.

Note: The order of the following steps is just one suggest manner of implementation, there are several other methods for achieving the same end configuration,



6. Create New Host and register initiator.
  - a. Select the row with desired initiator
  - b. On bottom left corner of screen, select Register
  - c. A New window will show up, use the following suggested settings
    - i. Initiator type= Clarion VNX (for Windows/VMware)
    - ii. Fail over Mode=Active-Active (default)
    - iii. Select New Host
    - iv. Enter host name , it may be preferable to include Rack\_Server\_Fabric\_SAN info (just to help with later identification)
    - v. IP address is needed even if not using host agent or PowerPath, IP address can be a fake address in order to create/register the initiator.
    - vi. Select OK
  - d. There will be several prompts to continue/ confirm your request.
  - e. At this point the final success box will appear, select OK when prompted
  - f. Continue registering the other initiators if multiple paths are available.
  - g. Select "existing host" (the one just created)
    - i. Browse Host...
    - ii. In the filter box, host name can be typed in partially to make the search easier if there are several hosts created.
    - iii. Select the host desired and select OK
  - h. Create desired number of LUNs.
    - i. From top menu tabs, select Storage>LUNs
    - ii. At the bottom left select "create" to open the Create LUNs dialog box and use the following suggested settings
      1. Storage type: Leave at Pool
      2. RAID type: Select default this example uses RAID 6
      3. Select Storage Pool for new LUN, this example uses "Pool 0"
      4. LUN Properties
        - a. User Capacity: select size, for ESX host at least 20GB minimum
        - b. LUN ID : allow system to assign, no selection required
      5. Number of LUNS to Create: Enter a number based on the infrastructure storage needs. (Note LUNs assigned to SPA and SPB in round robin) So the storage should be configured with access to SPA and SPB.
      6. LUN Name: leave at default, which automatically assigns LUN IDS.
      7. Select "Apply"
7. Create storage group (listed under the Host Menu).
  - a. From top menu tabs, select Host > Storage Groups
  - b. Mid to Bottom left , select "Create"
    - i. Create Storage box is presented
      1. Enter Storage Name (Storage admins will typically have protocols for this naming convention)
      2. Select OK, (this will show success then step through to add LUNs and Host)



8. Add Host and LUNs to new storage group
  - a. From LUNs tab, Select LUNs to add, note SPA and SPB should show the applicable LUNs
  - b. Select "add", LUN moved to selected LUNs
  - c. Select "Hosts" tab
    - i. Select Host from Available Hosts
    - ii. Select right pointing arrow to move Host to "Hosts to be Connected"
    - iii. Select "OK"

## 6.3 Two ways to confirm the Host LUN ID info

- This can be viewed through the storage group and note the Host LUN ID
  - o From Storage group list: Host > Storage group
  - o Select Storage group desired
  - o In the Details window, bottom of web page, select LUNs tab
  - o View the LUNs assigned to the storage group
  - o The far right column will list the "Host LUN ID"
- From Host > Host List
  - o Select Host from List
  - o In the Details pane, far right column will show Host LUN ID.

With the Cisco Nexus 5K series in NPV mode verify that negotiations has happened properly between the FCF and the network adapters end devices, in this case the Cisco Nexus 55548 is the FCF. Type: *show npv flogi-table*

```
NX5548K#show npv flogi-table
```

SERVER INTERFACE	VSAN	FCID	PORT NAME	NODE NAME	EXTERNAL INTERFACE
vfc101	2	0x850002	20:01:5c:f9:dd:16:ef:07	20:00:5c:f9:dd:16:ef:07	fc2/1
vfc102	2	0x850201	20:01:5c:f9:dd:16:ef:21	20:00:5c:f9:dd:16:ef:21	fc2/2

Figure 19 "show npv flogi-table" command results

The Cisco MDS FC switch is now the device providing fabric services

```
MDS9148# show flogi database
```

INTERFACE	VSAN	FCID	PORT NAME	NODE NAME
fc1/1	2	0x850000	20:1d:54:7f:ee:56:55:40	20:02:54:7f:ee:56:55:41
fc1/1	2	0x85001d	20:01:e0:db:55:1c:82:91	20:00:e0:db:55:1c:82:91
fc1/2	2	0x850100	20:1e:54:7f:ee:56:55:40	20:02:54:7f:ee:56:55:41
fc1/25	2	0x850f00	50:06:01:69:3e:e0:18:70	50:06:01:60:be:e0:18:70
fc1/26	2	0x850e00	50:06:01:68:3e:e0:18:70	50:06:01:60:be:e0:18:70

Total number of flogi = 9.



Figure 20 "show flogi database from the Cisco MDS 9148"

NOTE: SAN Port Channels is a preferred method of configuration for the links between the Cisco Nexus 5000 Series and the Cisco MDS switches. This topic will not be covered in the document, please refer to this document for detailed information around this configuration:

[http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/san\\_switching/Cisco\\_Nexus\\_5000\\_Series\\_NX-OS\\_SAN\\_Switching\\_Configuration\\_Guide\\_chapter7.html](http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/san_switching/Cisco_Nexus_5000_Series_NX-OS_SAN_Switching_Configuration_Guide_chapter7.html) (Cisco Systems Inc., 2009)

## 6.4 Further explanation of key points:

In order for devices to communicate end-to-end they must participate in the same VSAN. In this document the VFCs are bound to FCoE FIP MAC Address of the Network Adapter in the blade server or a physical interface that would be a connection from the modular switch (in this case the IOA) connecting the blade servers. The recommended configuration for this is to bind to the FCoE FIP MAC address of the NETWORK ADAPTOR, which enables better specific troubleshooting but many large deployments will chose to use binding to the logical port channel (LAG)

As an easy initial setup step, it is possible to bind a VFC to the physical interface or port-channel in order to determine which FCoE FIP MAC addresses will present from the network adapter. These can be noted and then matched with the M1000e' s Chassis Management Controller - CMC -> Server Overview -> WWN/MAC Summary. This enables a simple method of copy pasting the appropriate FCoE FIP MAC addresses into place for VFC bind configuration.



## 7 Dell PowerEdge M I/O Aggregator (FSB) Validation

### 7.1 show interface status

With all cables in place between the switches, the FC/FCoE SAN configured, and all switch settings configured use the following commands to validate this newly created configuration.

First check general status of the ports, and links by using: *show interfaces status*

```
IOA-B1#show interfaces status

Port      Description  Status Speed      Duplex Vlan
Te 0/1                    Down  Auto       Auto    --
Te 0/2                    Down  Auto       Auto    --
Te 0/3                    Up    10000 Mbit Full  20,1000
- section removed for sizing
Te 0/32                   Down  Auto       Auto    --
Fo 0/33                   Down  40000 Mbit Auto    --
Fo 0/37                   Down  40000 Mbit Auto    --
Te 0/41                   Down  Auto       Auto    --
- section removed for sizing
Te 0/52                    Up    10000 Mbit Full  20,1000
Te 0/53                   Down  Auto       Auto    --
Te 0/54                   Down  Auto       Auto    --
Te 0/55                   Down  Auto       Auto    --
Te 0/56                   Down  Auto       Auto    --
```

Figure 21 "show interfaces status" command

### 7.2 show fip-snooping config

On the PowerEdge M I/O Aggregator, the fip-snooping configuration can be checked by using: *show fip-snooping config*. This will show that the FIP-snooping feature is enabled and global configuration is set. It will also display the current default or configured FC-MAP.

```
IOA-B1 #show fip-snooping config

FIP Snooping Feature enabled Status: Enabled
FIP Snooping Global enabled Status: Enabled
Global FC-MAP Value: 0X0EFC00
Maximum Sessions Per Enode Mac: 32
```

Figure 22 "show fip-snooping config" command



## 7.3 show fip-snooping fcf

The IOA - FIP-snooping FCF connection information can be checked by using:

**show fip-snooping fcf**. This command will display the FCF's MAC address, the connecting interface and applicable VLAN configured for the connection.

```
IOA-B1# show fip-snooping fcf
```

FCF MAC	FCF Interface	VLAN	FC-MAP	FKA_ADV_PERIOD	No.of Enodes
00:05:73:dc:04:88	Po 128	1000	0e:fc:00	8000	1
00:05:73:dc:04:89	Po 128	1000	0e:fc:00	8000	1

Figure 23 "show fip-snooping fcf" command

## 7.4 show fip-snooping enode

The FCoE enabled connection to an enode or NETWORK ADAPTOR in a host server can be checked by using:

**show fip-snooping enode**. This command will display the Enode MAC address, which port the Enode is connected with, the FCF's MAC address, the applicable VLAN configured for the connection and the FC-ID.

```
IOA-B1# show fip-snooping enode
```

Enode MAC	Enode Interface	FCF MAC	VLAN	FC-ID
5c:f9:dd:16:ef:07	Te 0/1	00:05:73:dc:04:89	1000	85:02:01
5c:f9:dd:16:ef:21	Te 0/3	00:05:73:dc:04:88	1000	85:00:01

Figure 24 "show fip-snooping enode" command

## 7.5 show fip-snooping sessions

Using the Dell PE IOA, the fip-snooping sessions can be validated using:

**show fip-snooping sessions**. This command will display the following items that are in an FCoE session: Enode MAC address, Enode interface, and FCF MAC address, FCF connecting interface, FCoE VLAN being used, FCoE MAC address, FC-ID, Port WWPN, and Port WWNN being used for the session. Not all of the commands output is show (Port WWPN and WWNN).

```
IOA-B1# show fip-snooping sessions
```

Enode MAC	Enode Intf	FCF MAC	FCF Intf	VLAN	FCoE MAC	FC-ID
5c:f9:dd:16:ef:21	Te 0/3	00:05:73:dc:04:88	Po 128	1000	0e:fc:00:85:00:01	85:00:01
5c:f9:dd:16:ef:07	Te 0/1	00:05:73:dc:04:89	Po 128	1000	0e:fc:00:85:02:01	85:02:01

Figure 25 "show fip-snooping sessions" results





## 7.6 show fip-snooping statistics

Using the IOA, fip-snooping statistics can be validated using:

**show fip-snooping statistic** command. This command will display the various counters applicable to the fip-snooping feature and you can specify specifics like vlan ids and ports.

```
IOA-B1# show fip-snooping statistics interface vlan 1000
Number of Vlan Requests                               :0
Number of Vlan Notifications                           :0
Number of Multicast Discovery Solicits                  :7
Number of Unicast Discovery Solicits                    :0
Number of FLOGI                                         :5
Number of FDISC                                         :0
Number of FLOGO                                         :2
Number of Enode Keep Alive                             :30963
Number of VN Port Keep Alive                           :2757
Number of Multicast Discovery Advertisement             :39420
Number of Unicast Discovery Advertisement              :5
Number of FLOGI Accepts                                :5
Number of FLOGI Rejects                                :0
Number of FDISC Accepts                                :0
Number of FDISC Rejects                                :0
Number of FLOGO Accepts                                :2
Number of FLOGO Rejects                                :0
Number of CVL                                           :1
Number of FCF Discovery Timeouts                       :0
Number of Enode Mac Timeouts                           :2
Number of VN Port Session Timeouts                     :0
Number of Session failures due to Hardware Config      :0

IOA-B1# show fip-snooping statistics interface tengigabitethernet 0/1
Number of Vlan Requests                               :3
Number of Vlan Notifications                           :0
Number of Multicast Discovery Solicits                  :4
Number of Unicast Discovery Solicits                    :0
Number of FLOGI                                         :3
Number of FDISC                                         :0
Number of FLOGO                                         :2
Number of Enode Keep Alive                             :27197
Number of VN Port Keep Alive                           :2421
Number of Multicast Discovery Advertisement             :0
Number of Unicast Discovery Advertisement              :0
Number of FLOGI Accepts                                :0
Number of FLOGI Rejects                                :0
Number of FDISC Accepts                                :0
Number of FDISC Rejects                                :0
Number of FLOGO Accepts                                :0
Number of FLOGO Rejects                                :0
Number of CVL                                           :0
Number of FCF Discovery Timeouts                       :0
Number of Enode Mac Timeouts                           :0
Number of VN Port Session Timeouts                     :0
Number of Session failures due to Hardware Config      :0
```



Figure 26 “show fip-snooping statistics” command

## 7.7 show interface dcbx detail

The following command will show vital information on the DCBX configuration. Type: **show interfaces dcbx detail**. This command will show specific detail information about the configuration that has been negotiated between the devices. In this example interface, TenGigabitEthernet 0/3 is used as a downlink to a host server and NETWORK ADAPTOR installed in that system. Interface TenGigabitEthernet 0/52 is the uplink into the top-of-rack FCF switch.

Some key items in these results will be “true or false” for *is configuration source*, *Local DCBX Compatibility mode is*; and *Peer Operating version is*. These results will be important in understanding the negotiations happen between the switches and network adaptor.

```
IOA-B1# show interfaces dcbx detail

E-ETS Configuration TLV enabled          e-ETS Configuration TLV disabled
R-ETS Recommendation TLV enabled         r-ETS Recommendation TLV disabled
P-PFC Configuration TLV enabled          p-PFC Configuration TLV disabled
F-Application priority for FCOE enabled  f-Application Priority for FCOE disabled
I-Application priority for iSCSI enabled i-Application Priority for iSCSI disabled

-----

Interface TenGigabitEthernet 0/3
  Remote Mac Address 5c:f9:dd:16:ef:1f
  Port Role is Auto-Downstream
  DCBX Operational Status is Enabled
  Is Configuration Source? FALSE
  Local DCBX Compatibility mode is CEE
  Local DCBX Configured mode is AUTO
  Peer Operating version is CEE
  Local DCBX TLVs Transmitted: ErPFI

Interface TenGigabitEthernet 0/52
  Remote Mac Address 54:7f:ee:53:3e:88
  Port Role is Auto-Upstream
  DCBX Operational Status is Enabled
  Is Configuration Source? FALSE
  Local DCBX Compatibility mode is CIN
  Local DCBX Configured mode is AUTO
  Peer Operating version is CIN
  Local DCBX TLVs Transmitted: ErPFI
```

Figure 27 “show interfaces dcbx detail”



## 7.8 show interfaces pfc detail

The next command will show vital information on the DCBX established PFC configuration. Type: *show interfaces pfc detail*. This command will show specific detail information about the configuration that is negotiated between the devices. In this example interface, TenGigabitEthernet 0/3 is used as a downlink to a host server and NETWORK ADAPTOR installed in that system. Interface TenGigabitEthernet 0/52 is the uplink into the top-of-rack FCF switch.

Some key items in these results will be Priority list is three, Remote Willing status, (this should be willing enabled for the Host NETWORK ADAPTOR, and willing is disabled for FCF Nexus switch). Other important items is PFC DCBX Oper status, FCOE TLV TX Status is enabled.



IOA-B1# **show interfaces pfc detail**

Interface TenGigabitEthernet 0/3

Admin mode is on  
Admin is enabled  
Remote is enabled, Priority list is 3  
Remote Willing Status is enabled  
Local is enabled, Priority list is 3  
Oper status is internally propagated  
PFC DCBX Oper status is Up  
State Machine Type is Symmetric  
TLV Tx Status is enabled  
PFC Link Delay 45556 pause quntams  
Application Priority TLV Parameters :

-----  
FCOE TLV Tx Status is enabled  
ISCSI TLV Tx Status is disabled  
Local FCOE PriorityMap is 0x8  
Local ISCSI PriorityMap is 0x0  
Remote FCOE PriorityMap is 0x8  
Remote ISCSI PriorityMap is 0x10

11 Input TLV pkts, 28 Output TLV pkts, 5 Error pkts, 1202 Pause Tx pkts, 0  
Pause Rx pkts

Interface TenGigabitEthernet 0/52

Admin mode is on  
Admin is enabled  
Remote is enabled, Priority list is 3  
Remote Willing Status is disabled  
Local is enabled, Priority list is 3  
Oper status is recommended  
PFC DCBX Oper status is Up  
State Machine Type is Feature  
TLV Tx Status is enabled  
PFC Link Delay 45556 pause quntams  
Application Priority TLV Parameters :

-----  
FCOE TLV Tx Status is enabled  
Local FCOE PriorityMap is 0x8  
Remote FCOE PriorityMap is 0x8

0 Input TLV pkts, 1 Output TLV pkts, 0 Error pkts, 0 Pause Tx pkts, 12748  
Pause Rx pkts

Figure 28 "show interfaces pfc detail"



## 7.9 show interfaces ets detail

The following command will show vital information on the ETS settings in the configuration. Type: **show interfaces ets detail**. This command will show specific detail information about the ETS configuration that has been negotiated between the devices. In this example interface, TenGigabitEthernet 0/3 is used as a downlink to a host server and NETWORK ADAPTOR installed in that system. Interface TenGigabitEthernet 0/52 is the uplink into the top-of-rack FCF switch.

Some key items in these results will be TC-grp, Priority#, Bandwidth, and TSA. Note that only these two interfaces are seen typically this command will list them all.

```
IOA-B1# show interfaces ets detail

Interface TenGigabitEthernet 0/3
Max Supported TC Groups is 4
Number of Traffic Classes is 8
Admin mode is on

Admin Parameters:
-----
Admin is enabled

TC-grp      Priority#      Bandwidth      TSA
-----
0           0,1,2,3,4,5,6,7  100%           ETS
1           -                -              -
2           -                -              -

Remote Parameters :
-----
Remote is enabled

TC-grp      Priority#      Bandwidth      TSA
-----
0           0,1,2,4,5,6,7  50 %           ETS
1           3              50 %           ETS
2           -                -              -
3           -                -              -
```

Figure 29 "show interfaces ets detail"



```

Remote Willing Status is enabled
Local Parameters:
-----
Local is enabled

TC-grp      Priority#      Bandwidth      TSA
-----
0           0,1,2,4,5,6,7      50 %           ETS
1           3              50 %           ETS
2           -              -              -
3           -              -              -
4           -              -              -

Oper status is internally propagated
ETS DCBX Oper status is Up
State Machine Type is Feature
Conf TLV Tx Status is enabled

1 Input Conf TLV Pkts, 1 Output Conf TLV Pkts, 1 Error Conf TLV Pkts

Interface TenGigabitEthernet 0/52
Max Supported TC Groups is 4
Number of Traffic Classes is 8
Admin mode is on

Admin Parameters :
-----
Admin is enabled

TC-grp      Priority#      Bandwidth      TSA
-----
0           0,1,2,3,4,5,6,7      100%           -
1           -              -              -
2           -              -              -
3           -              -              -

```

Figure 30 "show interfaces ets detail" - continued



```

Remote Parameters :
-----
Remote is enabled

TC-grp      Priority#      Bandwidth      TSA
-----
0           0,1,2,4,5,6,7      50 %           ETS
1           3              50 %           ETS
2           -              -              -
3           -              -              -
4           -              -              -
5           -              -              -
6           -              -              -
7           -              -              -
15          -              -              -

Remote Willing Status is disabled
Local Parameters :
-----
Local is enabled

TC-grp      Priority#      Bandwidth      TSA
-----
0           0,1,2,3,4,5,6,7      100%           ETS
1           -              -              -
2           -              -              -
3           -              -              -
4           -              -              -
5           -              -              -
6           -              -              -
7           -              -              -
15          -              -              -

Oper status is init
ETS DCBX Oper status is Up with Mismatch
State Machine Type is Feature
Conf TLV Tx Status is enabled

23 Input Conf TLV Pkts, 7 Output Conf TLV Pkts, 0 Error Conf TLV Pkts

```

Figure 31 "show interfaces ets detail" - continued

After these validation steps, go into the disk management interface of the server. Verify the SAN is configured appropriately. The server will have an available LUN to use for storage.



```

Interface TenGigabitEthernet 0/3
  mtu 12000
  portmode hybrid
  switchport
  auto vlan
  flowcontrol rx on tx off
  dcb-policy input smux-dcb-in
  dcb-policy output smux-dcb-out
!
  protocol lldp
    advertise management-tlv system-name
    dcbx pot-role auto-downstream
  no shutdown

```

Figure 32 This is the typical downlink (server facing port) configured for FCoE

```

Interface TenGigabitEthernet 0/3
  mtu 12000
  portmode hybrid
  switchport
  auto vlan
  flowcontrol rx on tx off
  dcb-policy input smux-dcb-in
  dcb-policy output smux-dcb-out
!
  protocol lldp
    advertise management-tlv system-name
    dcbx pot-role auto-downstream
    dcbx version cee
  no shutdown

```

Figure 33 Addition steps for using Intel x520 NETWORK ADAPTOR

This additional command line "**dcbx version CEE**" will configure the port to only use the CEE version of DCBX and not negotiate other values. This setting is currently necessary for the Intel x520 NETWORK ADAPTOR with FCoE and applicable storage.



## 8 Cisco Nexus 5548UP and MDS 9148 Validation

In order to have a fully supported configuration the cables or SFP+ transceivers must only be Cisco-branded products for the Cisco Nexus. Many examples have been shortened from the actual results to reduce content.

### 8.1 Verifying Virtual Interfaces

You can verify information about each virtual interface by performing a few commands on the Nexus 5500 switch.

#### 8.1.1 `show interface vfc vfc-id`

To show information about a virtual Fiver Channel Interfaces use the *`show interface vfc vfc-id`* command. We had multiple vfs (101, 102, 201, and 202) on multiple switches in this document. Below is one example for vfc101.

```
NX5548K#show interface vfc 101
vfc101 is trunking
  Bound MAC is 5c:f9:dd:16:ef:07
  Hardware is Ethernet
  Port WWN is 20:64:54:7f:ee:53:3e:bf
  Admin port mode is F, trunk mode is on
  snmp link state traps are enabled
  Port mode is TF
  Port vsan is 2
  Trunk vsans (admin allowed and active) (1-2)
  Trunk vsans (up) (2)
  Trunk vsans (isolated) ( )
  Trunk vsans (initializing) (1)
  1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    0 frames input, 0 bytes
    0 discards, 0 errors
    0 frames output, 0 bytes
    0 discards, 0 errors
```

Figure 34 “show interface vfc101” command

#### 8.1.2 `show interface brief`

To display the results of active ports and VFC interfaces. Verify that the ports, which are expected to have links, show up correctly. Type: *`show interface brief`*



NX5548K#show interface brief									
<hr/>									
Interface	Vsan	Admin Mode	Admin Trunk Mode	Status	SFP	Oper Mode	Oper Speed (Gbps)	Port Channel	
<hr/>									
fc2/1	2	NP	off	up	sw1	NP	8	--	
fc2/2	2	NP	off	up	sw1	NP	8	--	
<hr/>									
Ethernet Interface	VLAN	Type	Mode	Status	Reason		Speed	Port Ch #	
<hr/>									
Eth1/1	20	eth	trunk	up	none		10G(D)	20	
Eth1/2	20	eth	trunk	up	none		10G(D)	20	
Eth1/3	1	eth	trunk	down	SFP not inserted		10G(D)	1	
*rest of the ports removed for sizing									
<hr/>									
Port-channel Interface	VLAN	Type	Mode	Status	Reason		Speed	Protocol	
<hr/>									
Po1	1	eth	trunk	down	No operational members		auto(D)	lACP	
Po20	20	eth	trunk	up	none		a-10G(D)	lACP	
<hr/>									
Interface	Vsan	Admin Mode	Admin Trunk Mode	Status	Bind Info		Oper Mode	Oper Speed (Gbps)	
<hr/>									
vfc101	2	F	on	trunking	5c:f9:dd:16:ef:07		TF	auto	
vfc102	2	F	on	trunking	5c:f9:dd:16:ef:21		TF	auto	

Figure 35 "show interface brief" command

## 8.2 Verifying NPV

A few commands to verify NPV information is listed below

### 8.2.1 show npv status

To display the NP uplink interfaces, enter the *show npv status* command.



```

NX5548K#show npv status
npiv is enabled
disruptive load balancing is disabled
External Interfaces:
=====
Interface:  fc2/1, VSAN:      2, FCID: 0x850000, State: Up
Interface:  fc2/2, VSAN:      2, FCID: 0x850200, State: Up
Number of External Interfaces: 2

Server Interfaces:
=====
Interface:  vfc101, VSAN:      2, State: Up
Interface:  vfc102, VSAN:      2, State: Up
Number of Server Interfaces: 2

```

Figure 36 "show npv status" command

## 8.2.2 show npv flogi-table

The devices, server interfaces and their NP uplinks can be displayed with the *show npv flogi-table* command. The command shows that vsan 2 has two servers logged in.

```

NX5548K# show npv flogi-table vsan 2
-----
SERVER
INTERFACE VSAN FCID          PORT NAME          NODE NAME          EXTERNAL
-----
vfc101     2      0x850201 20:01:5c:f9:dd:16:ef:07 20:00:5c:f9:dd:16:ef:07 fc2/2
vfc102     2      0x850001 20:01:5c:f9:dd:16:ef:21 20:00:5c:f9:dd:16:ef:21 fc2/1

Total number of flogi = 2.

```

Figure 37 "show npv flogi -table" command

## 8.2.3 show fcns database

The devices, server interfaces and their NP uplinks can be displayed with the *show fcns database* command. The command shows that vsan 2 has two servers logged in (scsi-fcp: init).



```
MDS9148# show fcns database
```

```
VSAN 2:
```

FCID	TYPE	PWWN	(VENDOR)	FC4-TYPE:FEATURE
0x850000	N	20:41:54:7f:ee:53:3e:80	(Cisco)	npv
0x850001	N	20:01:5c:f9:dd:16:ef:21		scsi-fcp:init
0x850100	N	50:06:01:6f:3e:e0:18:70	(Clariion)	scsi-fcp
0x850200	N	20:42:54:7f:ee:53:3e:80	(Cisco)	npv
0x850201	N	20:01:5c:f9:dd:16:ef:07		scsi-fcp:init

Total number of entries = 5

```
VSAN 2222:
```

FCID	TYPE	PWWN	(VENDOR)	FC4-TYPE:FEATURE
0x300000	N	20:2a:00:1e:c9:f1:03:de		
0x300001	N	20:01:14:fe:b5:8e:5b:fe		scsi-fcp:init
0x300200	N	50:06:01:69:3e:e0:18:70	(Clariion)	scsi-fcp

Total number of entries = 3

Figure 38 "show fcns database" command

## 8.3 show flogi database

Next, verify that negotiations have happened properly between the FCF and end devices. In this case, the Cisco Nexus 5548 is the FCF. Type: *show flogi database*

```
NX5548K# show flogi database
```

INTERFACE	VSAN	FCID	PORT NAME	NODE NAME
fc1/29	2	0xc6000c	20:04:00:20:c2:10:ab:cd	10:00:00:20:c2:10:ab:cd
vfc2	2	0xc60009	20:01:14:fe:b5:07:12:34	

Figure 39 "show flogi database" command

At this point, the VFC and FC interfaces should be populated in the FLOGI database. This command is showing the devices that have done a valid FLOGI (fabric login) to the Cisco Nexus switch. The VFC should show the expected port and node WWN of the NETWORK ADAPTOR being used in the server.

## 8.4 show zoneset active

The following command will show the status of the zones that have been put into place. Type: *show zoneset active*. This command will show the current activated zoneset and all the participating zones with their individual members.



```

NX5548K# show zoneset active
zoneset name set1 vsan 2
zone name blade1 vsan 2
* fcid 0xc6000c [interface fc1/29 swwn 20:00:00:05:73:ab:12:34]
* fcid 0xc60009 [pwwn 20:01:14:fe:b5:07:12.34]

```

Figure 40 Example of “show zoneset” active command

**NOTE:** A key item to understand in looking at these validation/debug commands is that both configurations need to be in place for many of these commands to show the expected results. In other words if the Cisco Nexus 5548 has been configured based on the CLI examples in this document but the Dell PowerEdge M I/O Aggregator has not been connected properly many of the results shown here will not match. The storage device and SAN switches will also need to be configured if using NPV mode.

## 8.5 show spanning-tree summary

Check the spanning tree configuration. Blocking ports should be understood and when they show in this entry, it should be the number expected. Otherwise it could be that an unintentional cable loop has been created that will need to be resolved. Type: *show spanning-tree summary*

```

NX5K# show spanning-tree summary
Switch is in rapid-pvst mode
Root bridge for: VLAN0020, VLAN0088, VLAN1000
Port Type Default is disable
Edge Port [PortFast] BPDU Guard Default is disabled
Edge Port [PortFast] BPDU Filter Default is disabled
Bridge Assurance is enabled
Loopguard Default is disabled
Pathcost method used is short
STP-Lite is enabled

```

Name	Blocking	Listening	Learning	Forwarding	STP Active
VLAN0020	0	0	0	1	1
VLAN0088	0	0	0	1	1
VLAN1000	0	0	0	1	1
3 vlans	0	0	0	3	3

Figure 41 “show spanning-tree summary” command



## A Basic Terminology

### CLI

Command line interface (CLI). This is the text-based telnet, SSH, or serial type interface that are used for entering commands into the Dell Networking MXL model switch. At release, the MXL is only configurable via CLI.

### CMC

Chassis management controller (CMC) is the module, which controls the Dell PowerEdge™ M1000e blade server chassis. Through this controller a telnet, SSH or serial based connection can be used to manage the MXL switch.

### ETS

Enhanced Transmission Selection (ETS) is defined in the IEEE 802.1Qaz standard (IEEE, 2011). ETS supports allocation of bandwidth amongst traffic classes. It then allows for sharing of bandwidth when a particular traffic class does not fully utilize the allocated bandwidth. The management of the bandwidth allocations is done with bandwidth-allocation priorities, which coexist with strict priorities (IEEE, 2011).

### FIP-Snooping

With FIP-Snooping enabled on the Dell Networking MXL model switch, FIP logins, solicitations, and advertisements are monitored. In this monitoring or snooping process, the switch gathers information pertaining to the ENode and FCF addresses. With this information, the switch will then place filters that only allow access to ENode devices that have logged-in successfully. This enables the FCoE VLAN to deny all other traffic except this lossless FCoE storage traffic.

The filtering process also secures the end-to-end path between the ENode device and the FCF. The ENode will only be able to talk with the FCF in which it has logged into.

### FIP-Snooping bridge (FSB)

With a switch configured to performing FIP-Snooping the industry, term for this switch is FSB or FIP-Snooping Bridge. It is performing FIP-Snooping as described in the previous term.

### FCF

FCoE forwarders (FCFs) act as an Ethernet and FC switch combined. All typical termination functions that would occur on a FC switch occur on the FCF. FCF's give VF\_Ports and VE\_Ports for their virtual FC interfaces.

### IOM

IO module (IOM) refers to the modules on the rear of the Dell PowerEdge M1000e chassis that will receive and transmit I/O (Ethernet, FC, InfiniBand, etc.) from the blade servers. The Dell Networking MXL switch is as an IOM for the M1000e blade server chassis.



## MAC Address

Media Access Control Address (MAC Address) is a layer-2 node identifier. In Ethernet bridging, MAC addresses are used for source and destination identification. They can also be used as system identifiers since vendor-assigned (or burned-in) MAC addresses are globally unique. An Ethernet MAC address is 48 bits long and generally written in groupings of two hexadecimal digits often separated by colons or hyphens like this: 00:1e:c9:00:cb:01. But, are sometimes written in groupings of four hexadecimal digits separated by periods like this: 001e.c900.cb01

## NPIV

N-port identifier virtualization, which enables multiple N-port fabric logins at the same time on the same physical FC link (Cisco Systems, Inc., 2011). This term, is in reference to the Cisco Nexus 5000 series switches implementation of NPIV. NPIV must be enabled to share multiple logins across a single port/link or a port-channel/multiple-line link.

## NPV

N-port Virtualizer is a FC aggregation method, which passes traffic through to end devices, while eliminating the need to use a domain ID for this device (Cisco Systems, Inc., 2011). This term is also in reference to configuration settings on the Cisco Nexus 5000 series switches.

## PFC

Priority Flow Control (PFC) or Per-Priority Pause is defined in the IEEE 802.1Qbb standard. PFC is flow control based on priority settings and adds additional information to the standard pause frame. The additional fields added to the pause frame allow devices to pause traffic on a specific priority instead of pausing all traffic. (IEEE, 2009) Pause frames will be initiated by the FCF in most cases when its receive buffers are starting to reach a congested point. With PFC, traffic is paused instead of dropped and retransmitted. This provides the lossless network behavior necessary for FC packets to be encapsulated and passed along the Ethernet paths.

## ToR

Top of Rack (ToR) is a term for a switch that is actually positioned at the top of a server rack in a data center.

## VLAN

Virtual Local Area Network (VLAN) is a single layer-2 network (also called a broadcast domain, as broadcast traffic does not escape a VLAN on its own). Multiple VLANs can be passed between switches using switchport trunk interfaces. When passed across trunk links, frames in a VLAN are prefixed with the number of the VLAN that they belong to—a twelve-bit value that allows just over 4000 differently numbered VLANs.

## VSAN

Virtual SAN is a logical partitioning of physical connections to provide for fabric or SAN separation. VSAN is a term that is particular to the Cisco Nexus series switches.



## B References

- Cisco Systems, Inc. (2011). *Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide, Release 5.1(3)N1(1)*. San Jose: Cisco Systems, Inc  
[http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/san\\_switching/513\\_n1\\_1/b\\_Cisco\\_n5k\\_nxos\\_sanswitching\\_config\\_guide\\_rel513\\_n1\\_1.html](http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/san_switching/513_n1_1/b_Cisco_n5k_nxos_sanswitching_config_guide_rel513_n1_1.html)
- Cisco Systems, Inc. (2010, 01). *Cisco 4000 Series Design Guide*. Retrieved 06 19, 2012, from Cisco.com:  
[http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps10596/deployment\\_guide\\_c07-574724.html](http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps10596/deployment_guide_c07-574724.html)
- Fibre Channel over Ethernet Initialization Protocol, Cisco,  
[http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/white\\_paper\\_c11-560403.html](http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/white_paper_c11-560403.html)
- IEEE. (2008, November 24). 802.1Qaz/D0.2. Draft Standard for Local and Metropolitan Area Networks – Virtual.
- IEEE. (2009, Feb 9). 802.1Qbb/D1.0. Draft Standard for Local and Metropolitan Area Networks – Virtual.
- IEEE. (2011, Sep 15). *802.1Qaz - Enhanced Transmission Selection*. Retrieved May 30, 2012, from IEEE802.org:  
<http://www.ieee802.org/1/pages/802.1az.html>
- T11/08-264v0 FCoE: Increasing FCoE Robustness using FIP-Snooping and FPMA, T11,  
<http://www.t11.org/ftp/t11/pub/fc/bb-5/08-264v0.pdf>
- T11/09-291v0 FIP VLAN discovery updates, T11,  
<http://www.t11.org/ftp/t11/pub/fc/bb-5/09-291v1.pdf>
- Cisco Systems Inc. (2009). *Configuring SAN Port Channel*. Retrieved 7 15, 2013, from Cisco.com - Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide:  
[http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/san\\_switching/Cisco\\_Nexus\\_5000\\_Series\\_NX-OS\\_SAN\\_Switching\\_Configuration\\_Guide\\_chapter7.html](http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/san_switching/Cisco_Nexus_5000_Series_NX-OS_SAN_Switching_Configuration_Guide_chapter7.html)
- Cisco Systems, Inc. (2009, 07). *Virtual PortChannel Quick Configuration Guide*. Retrieved Jun 18, 2012, from Cisco.com:  
[http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/configuration\\_guide\\_c07-543563.html](http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/configuration_guide_c07-543563.html)
- IEEE. (2011, Sep 15). *802.1Qaz - Enhanced Transmission Selection*. Retrieved May 30, 2012, from IEEE802.org:  
<http://www.ieee802.org/1/pages/802.1az.html>
- Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide, Release 5.2(1)N1(1)  
[http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/san\\_switching/521n11/b\\_5k\\_SAN\\_Switching\\_Config\\_521N11.html](http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/san_switching/521n11/b_5k_SAN_Switching_Config_521N11.html)

