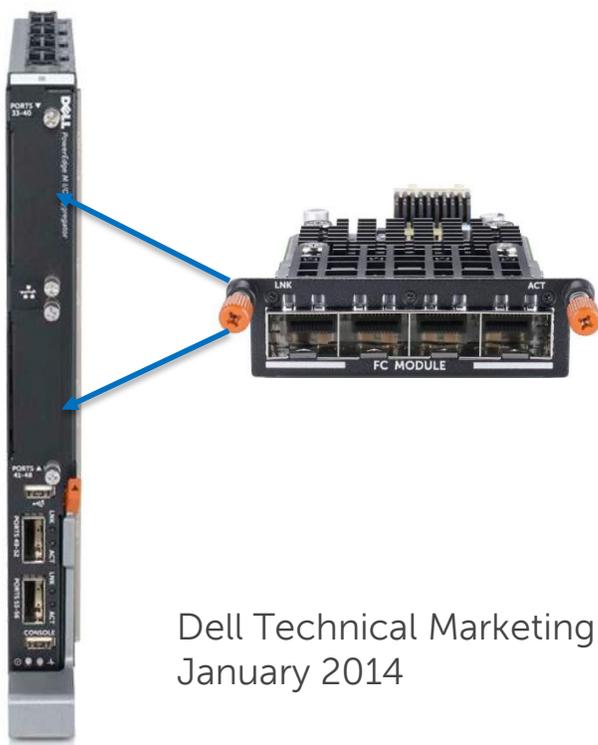# Dell Networking FC Flex IOM:

Deployment of FCoE with Dell FC Flex IOM, Brocade FC switches, and Dell Compellent Storage Array

A Dell Deployment/Configuration Guide

Dell Technical Marketing – Data Center Networking
January 2014

2   Dell Networking FC Flex IOM: Deployment of FCoE with Dell FC Flex IOM, Brocade FC switches, and Dell Compellent Storage Array

Deployment/Configuration Guide

Revisions

Date	Description	Authors
01/16/2014	Version 1 - Initial Document	Humair Ahmed, Neal Beard

Contents

1 Executive Summary

In the "Dell Networking FC Flex IOM: Infrastructure & Network Convergence w/ FCoE" whitepaper we demonstrated and explained the movement from a traditional non-converged LAN/SAN network to a converged LAN/SAN infrastructure and network and how the Dell MXL/IOA with Dell FC Flex IOM is an ideal solution for this transition. In addition to covering the industry trend towards IO Consolidation and network convergence, we covered the many benefits offered by FC Flex IOM which offers not only a converged networking solution but also a converged infrastructure where the FCoE-FC bridging functionality is implemented directly within the blade switch.

The Dell FC Flex IO module transforms a Dell MXL/IOA blade switch from a Ethernet-only switch to a converged switch capable of bridging between Ethernet and Fibre Channel. The FC Flex IOM moves the convergence layer from a typical dedicated ToR switch down to the blade switch via extensible IO module providing the benefits of IO consolidation such as less infrastructure hardware, less maintenance, and considerable cost savings. With its unique modular design FC Flex IOM allows end users to migrate to a converged solution and introduce FCoE functionality to the MXL/IOA blade switch at their own pace without replacing the entire switch. This benefit is unmatched in the industry. In this whitepaper we cover detailed topology and configuration examples.
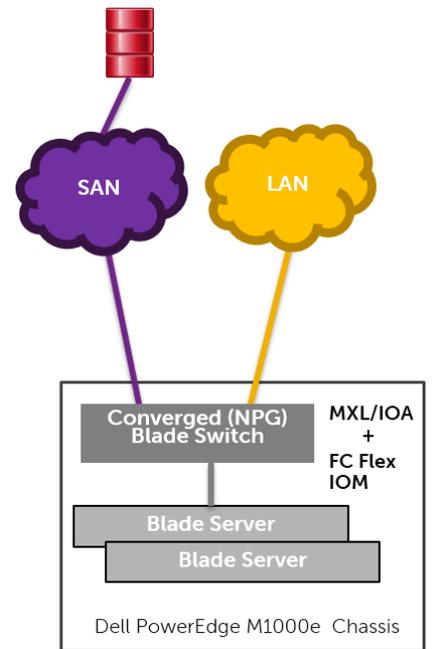
Figure 1   Converged Infrastructure/Network

2 Dell PowerEdge M1000e Overview

The PowerEdge M1000e Modular Server Enclosure solution supports up to (32) server modules, and (6) network I/O modules. The M1000e contains a high performance and highly available passive midplane that connects server modules to the infrastructure components, power supplies, fans, integrated KVM and Chassis Management Controllers (CMC). The PowerEdge M1000e uses redundant and hot-pluggable components throughout to provide maximum uptime. The chassis has the ability to house 6 x I/O modules allowing for a greater diversity of roles for all of the enclosed blade servers.

The (6) I/O slots in the back of the chassis are classified as 3 separate fabrics with each fabric containing 2 slots (A1/A2, B1/B2, C1/C2); these fabric I/O slots relate to the ports found on the server side network adaptors. The I/O modules can be used independently of each other, and each I/O module must contain the same technology. For example, fabric A is hardwired to the 2 network adaptors on the blade server mainboards, which means the I/O modules in fabric A must support Ethernet; Fabrics B and C can be used for Ethernet, Fibre Channel, or InfiniBand. Figure 2 below exemplifies the I/O mappings between the server side dual/quad port networking adaptors and the I/O modules.

> Note: The networking adaptors in Fabric A have also been described as LOM's (LAN on Motherboards), and bNDC's (blade Network Daughter Card's). All of these terms describe the same device: A network adaptor that performs Ethernet/iSCSI/FCoE tasks on behalf of the Server and its operating system.

Figure 2   M1000e Front and Back View
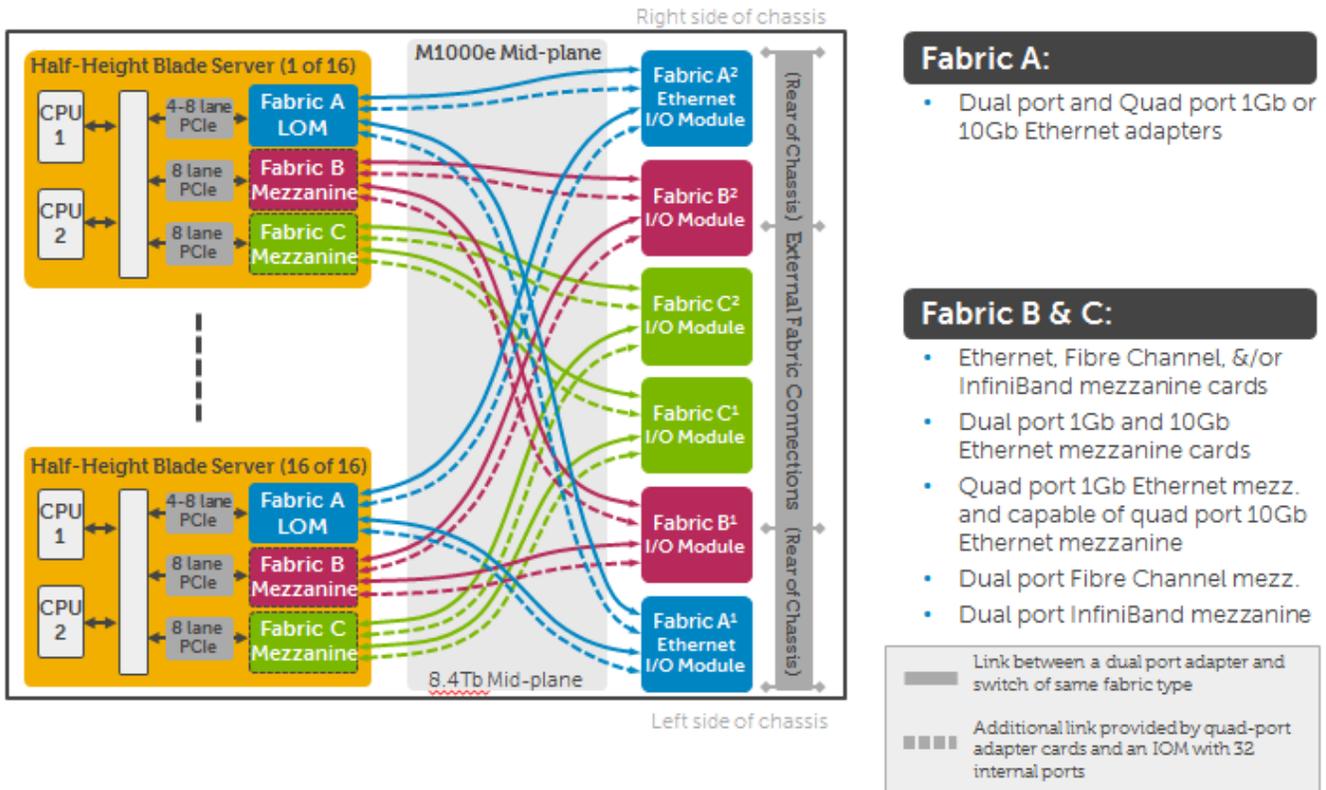
Figure 3   M1000e Midplane Dual/Quad Port Network Adaptor I/O Mappings

2.1   Flex IO Expansion Modules (External Ports)

The Dell I/O Modules will support a combination of Flex IO 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

4 Port Fibre Channel 8Gb Module

4 Port 10G-BaseT Module

4 Port 10Gb SFP+ Module

2 Port 40Gb QSFP+ Module
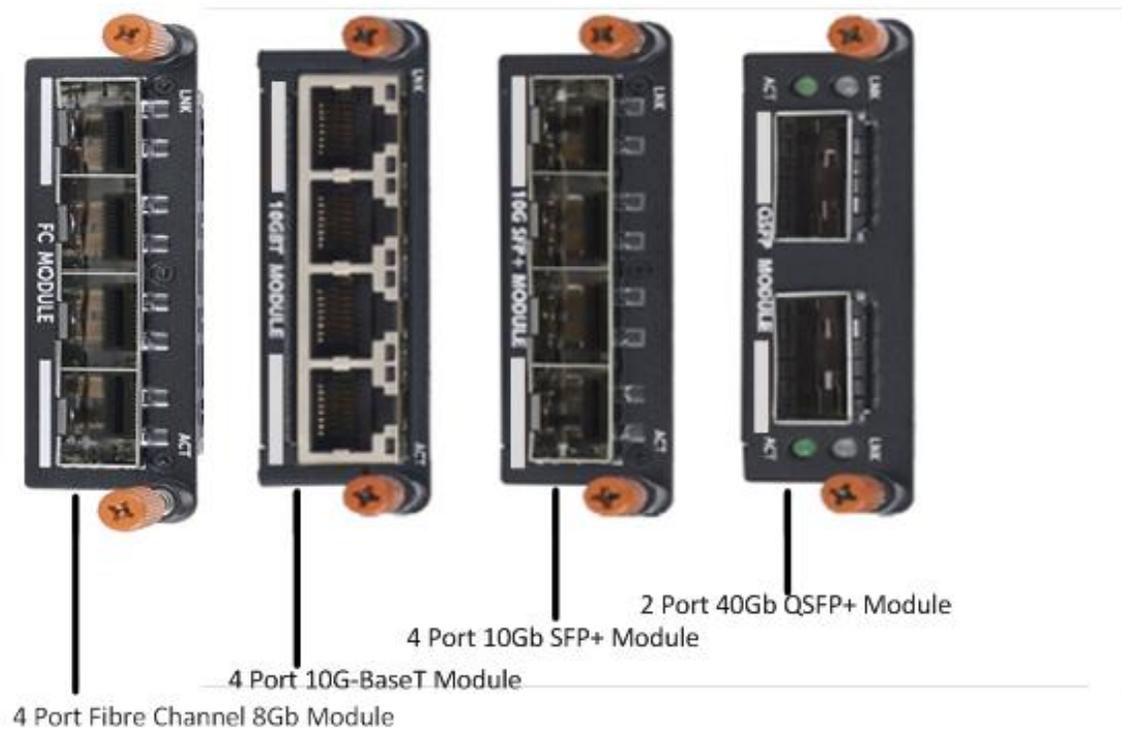
Figure 4   FlexIO expansion modules

NOTE:
The 4 Port Fibre Channel 8Gb module can only be used with the release of 9.3 firmware for the Dell MXL/IOA modular switches.

3 Dell FC Flex IOM and Convergence on Dell MXL/IOA Overview
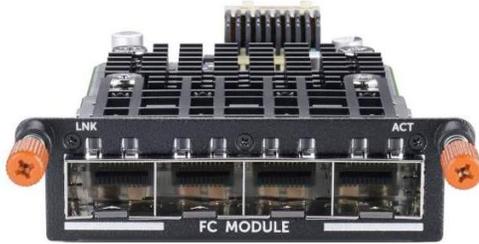
The Dell FC Flex IO module provides the additional functionality to a Dell MXL/IOA blade switch to enable it to act as a NPIV Proxy Gateway (NPG) capable of bridging between Ethernet and FC. FC Flex IOM takes the convergence from the ToR down to the blade level consolidating infrastructure while still providing the benefits of network convergence and leveraging the backend FC SAN.

In NPG mode, the Dell MXL/IOA w/ FC Flex IOM does not consume a fabric Domain ID or become part of the switched fabric but instead simply acts as a gateway to the fabric by de-encapsulating FC from FCoE and forwarding the frames to the existing backend FC SAN.

NPG technology in the FC FlexIO module transparently performs a proxy fabric login (FLOGI) on behalf of the CNAs in the M1000e blade servers. With NPIV (N_Port ID Virtualization) support on the Fibre Channel switch connected to the FC Flex IOM, the respective FC Flex IOM port acts as a N_Port and is able to obtain multiple fabric addresses (FC_IDs) per request by end nodes. A FLOGi is done initially by the FC Flex IOM to the FC switch providing fabric services and then the FC Flex IOM is able to obtain additional addresses over the same link by converting FLOGis from the server to FDISCs (Fabric Discoveries). The FC Flex IOM initially performs a FLOGi and then it obtains additional addresses for server nodes via FDISC (Fabric Discoveries).

Fibre Channel fabric services are maintained by the FC switch that the FC FlexIO module is connected to. The features that the FC FlexIO module provides are:

1. Manages the FLOGI and FDISC conversion process
2. Manages PFC (Priority Flow Control)
3. Manages ETS (Enhanced Transmission Selection)
4. Manages FIP keep alives

4 Converged Network Solution – Dell PowerEdge M1000e, Dell Compellent Storage Array, Brocade FC Switch, and Dell MXL/IOA w/ FC Flex IO as NPIV Proxy Gateway

This solution demonstrates network and infrastructure convergence via Dell MXL blade switch w/ FC Flex IOM. The Dell PowerEdge M1000e chassis consists of one Dell PowerEdge M620 Blade Server with a 2-port Broadcom 57810 CNA and two Dell MXL blade switches each containing a FC Flex IO module. The gray links in Figure 5 below represent the converged traffic. The convergence takes place within the chassis and the MXL switch w/ FC Flex IOM breaks out the LAN and SAN traffic to the respective backend LAN and SAN networks. Two Dell S6000s are used for upstream LAN switches and two Brocade 6505 switches are used for the backend SAN. A Dell Compellent Storage Array is used for FC storage and consists of one SC220 storage enclosure and two SC8000s with 4 x FC ports each.

In this example 40 GbE active fiber DACs are used from the 40GbE ports on the MXL to the 40 GbE ports on the S6000, but 40 GbE copper DACs could also have been used. OM3 fiber and 8G FC optics are used to connect the FC ports from the FC Flex IOM to the FC ports on the Brocade FC switches.

This solution provides redundancy for both the LAN and the SAN. The server sees two LAN connections and two SAN connections via the 2-port CNA adapter. The Broadcom 57810 CNA is NIC teamed via 'Smart Load Balancing and Failover' and one port is made active while the other is made standby. In this setup, it is also possible to add an additional 40 GbE or 10 GbE module to the MXLs and create a VLTi link between them. Although a VLT LAG cannot be utilized down to the server CNA as FCoE is not supported over a VLT lag, VLT can still be utilized up to the S6000s if desired. A VLTi link between the MXLs can help speed-up failover as MAC/ARP tables are synced and addresses do not have to be learned upon NIC failover; however, this is not required and will consume additional ports requiring the need for an additional Flex I/O module. Also, it's important to note that the FCoE traffic is unaffected by the NIC team, and, in this setup, the mutipathing policy on the server is set to **Round Robin** so both fabric A and fabric B are active and utilized.
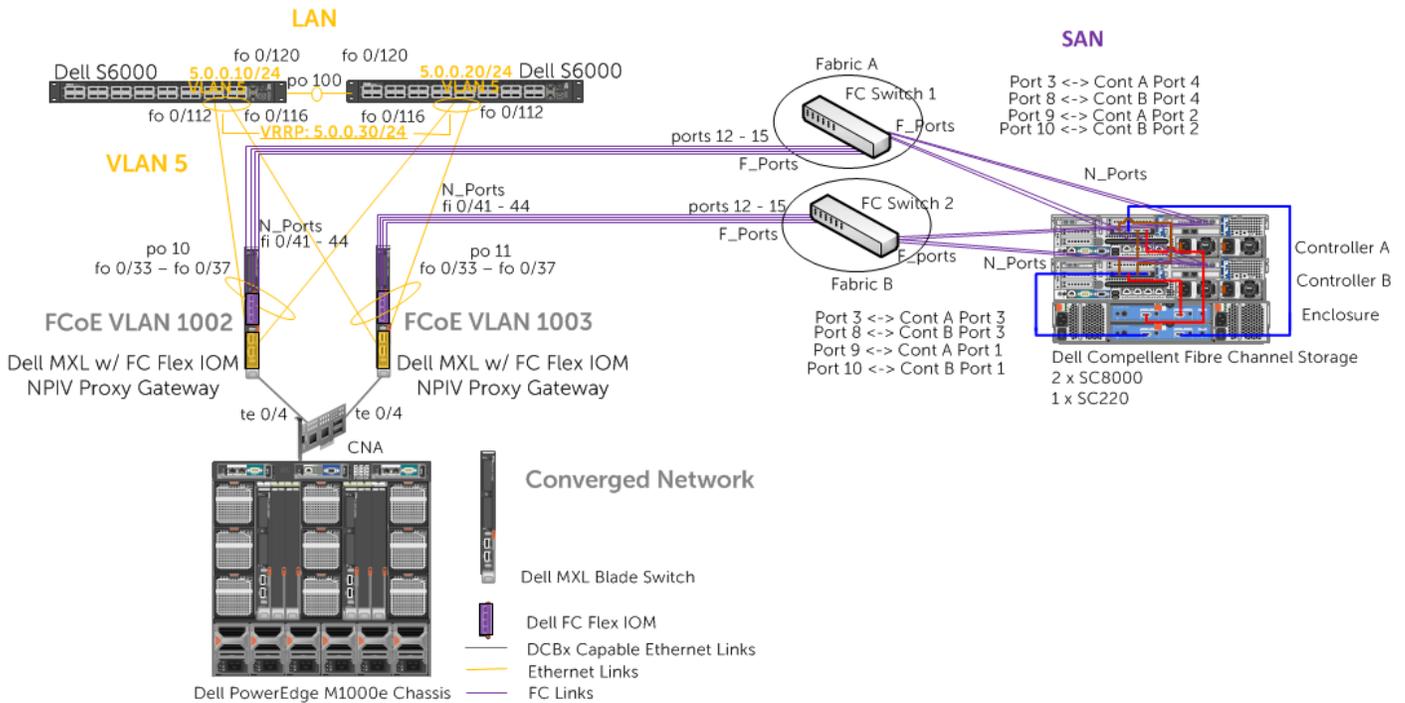
Figure 5   Dell MXL w/ FC Flex IOM acting as a NPIV Proxy Gateway and allowing for a converged
            infrastructure

This setup involves configuration on the CNA, server, MXL blade switches w/ FC Flex IOMs, LAN Dell
S6000 switches, Brocade 6505 FC switches, and the Dell Compellent storage array. This document
steps through the configuration one node at a time. Since the Dell S6000s are just providing the
connectivity to the backend LAN and there is no specific configuration involved in regards to
convergence, the configuration for these switches is standard LAN configuration and will not be shown.
The configuration of the Broadcom 57810 CNA, Dell PowerEdge M620 server, Dell MXL blade switch w/
FC Flex IOM, Brocade 6505 FC switch, and Dell Compellent Storage array is shown below. This example
demonstrates the configuration of the Broadcom 57810 CNA. To see how to configure a QLogic CNA
see Section 5 of this document. Additionally, this example uses Brocade FC switches and a Dell
Compellent Storage array. To see an example of Dell MXL/IOA w/ FC Flex IOM utilized with Cisco MDS
switches and a EMC VNX 5300 storage array, see the "Deploying Dell Networking MXL and PowerEdge
M I/O Aggregator with the FC FlexIO Module in a Cisco MDS Environment" guide.

4.1   Broadcom 57810 CNA Configuration

Broadcom offers the Broadcom BCM57810S in three formats for Dell servers: standard PCI Express, mezzanine card for Dell blade servers, and Network Daughter Card (NDC) for Dell blade servers. The Broadcom BCM57810S allows for NIC partitioning (NPAR) with up to four partitions per physical port and eight partitions total per 2-port adapter. A partition can be looked upon as a virtual port.

This example will use a Dell PowerEdge M620 blade server with a Broadcom BCM57810S CNA NDC and Microsoft Windows Server 2008 R2 Enterprise installed. By default, NPAR is not enabled and only the NIC functionality is enabled. FCoE must be manually enabled on the CNA for the virtual FCoE ports to be identified in Windows. This example uses the default 'Single Function' mode where NPAR is not utilized.

1.   Install the Broadcom BCM57810 drivers and Broadcom Advanced Control Suite 4.

2.   Double click the **Broadcom Advanced Control Suite 4** icon/shortcut in Windows. The Broadcom Advanced Control Suite 4 (Figure 6) opens. In this setup, Adapter 1 is being utilized. Note the small red **x** which appears next to the adapter ports. This is because the Dell MXL switches which the ports are connecting to have not yet been configured.
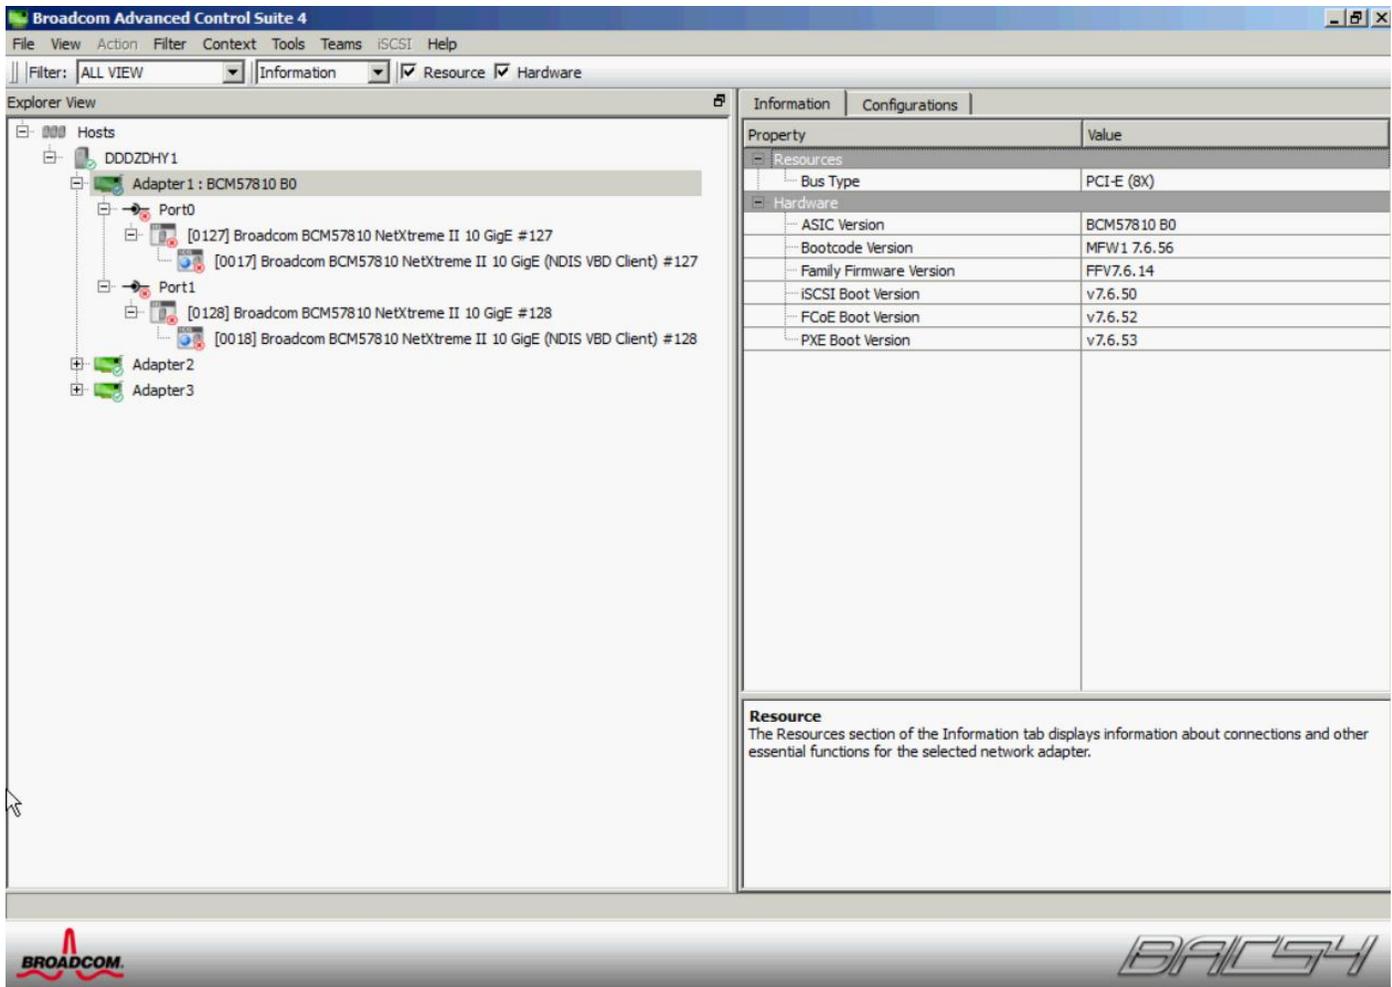
Figure 6   Broadcom 57810 CNA adapter view in Broadcom Advanced Control Suite 4

Note: In the current configuration (Figure 6), there are two ports and only one partition per port. This is because the adapter is in Single Function mode and NPAR is not being utilized.

If NPAR is desired, configure it by clicking the **Adapter BCM57810** label, clicking the **+** symbol to the right of **Multi-function**, clicking the **Configure** button and selecting **NIC Partition** from the drop down list (Figure 7).
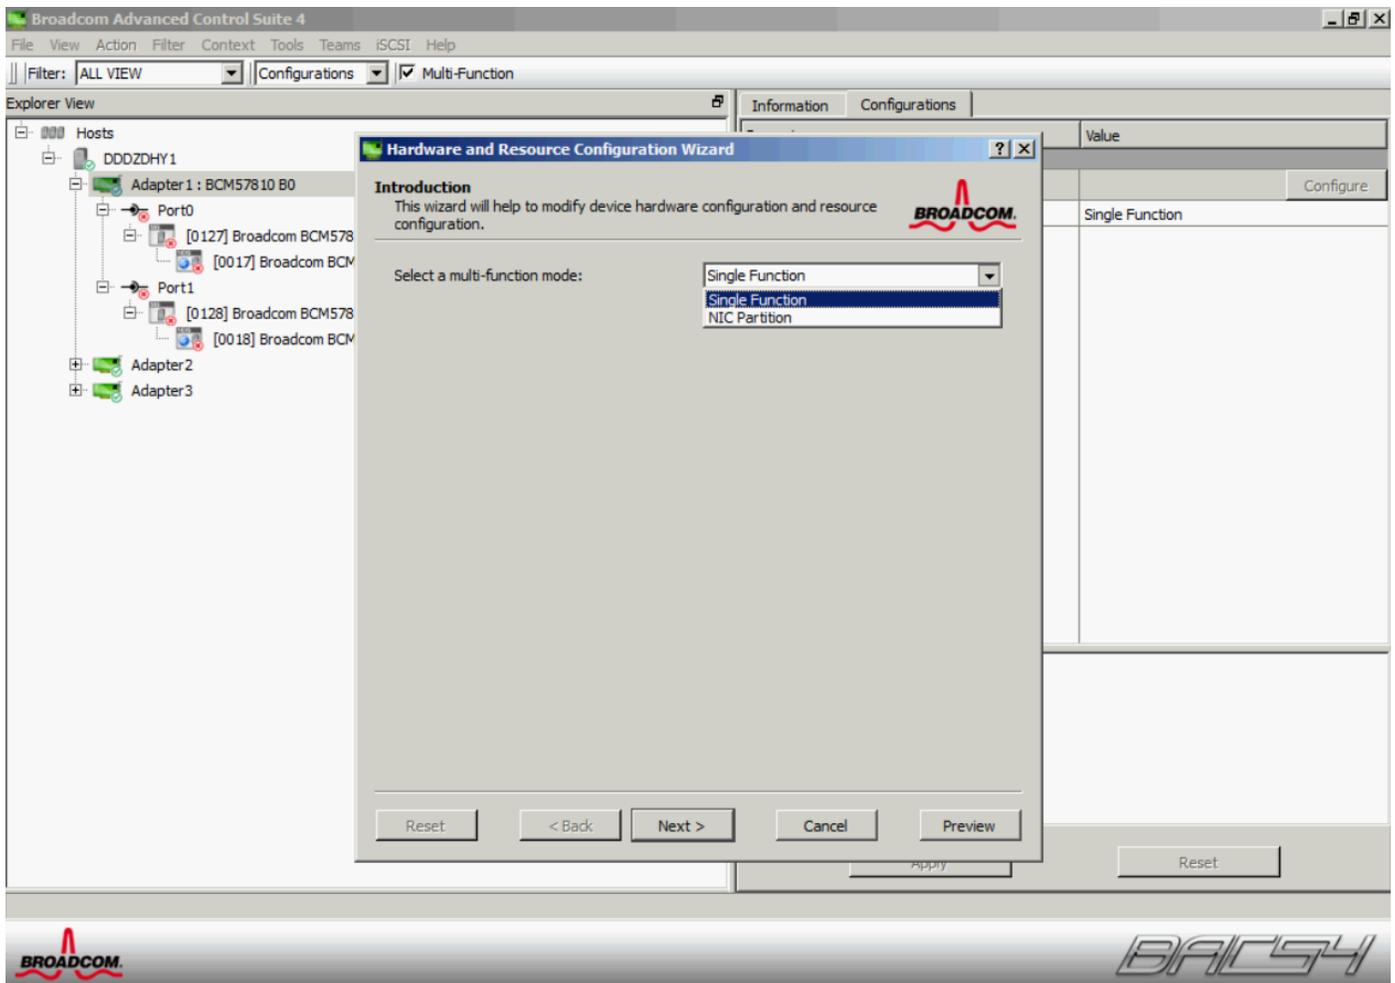
Figure 7   Default 'Single Function' mode setting on Broadcom 57810 CNA

3. Go to **Control Panel** > **Network and Internet** > **Network and Sharing Center** > **Change adapter settings**, The Broadcom Network Adapters are displayed. The two Broadcom 57810 CNA ports being utilized in this setup are highlighted (Figure 8).
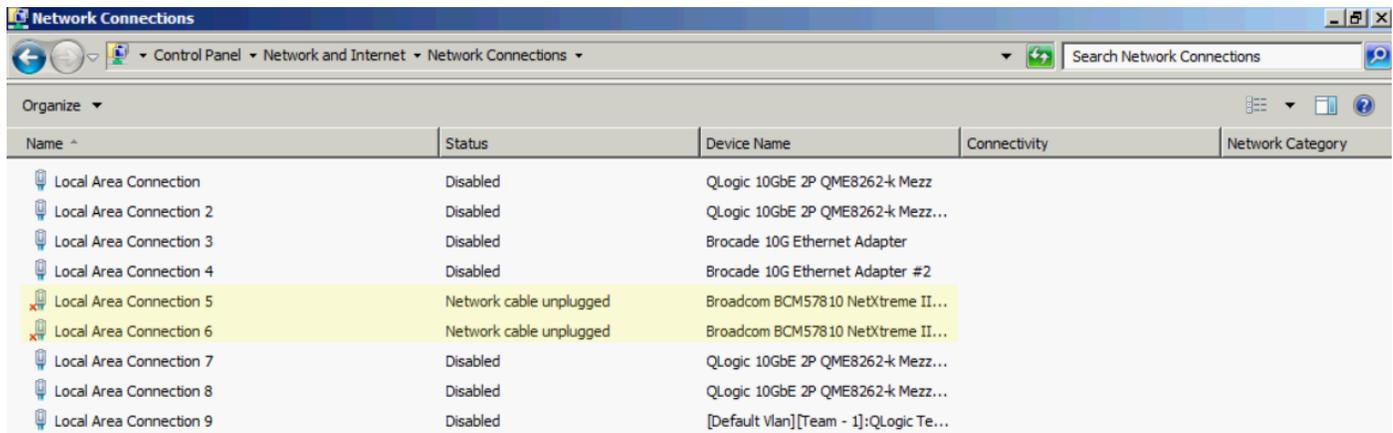
Figure 8   'Network Connections' view in Windows Server 2008 R2 Enterprise

4. Use the **Broadcom Advanced Control Suite 4** to enable FCoE by selecting the single partition under each port and clicking the **Configurations** tab on the right. Click the **+** symbol next to **Resource Reservations** and click the **Configure** button. Check the checkbox next to **FCoE** (Figure 9).

5. Click **Next** and then **Apply**. Repeat this step to enable FCoE on the second port.
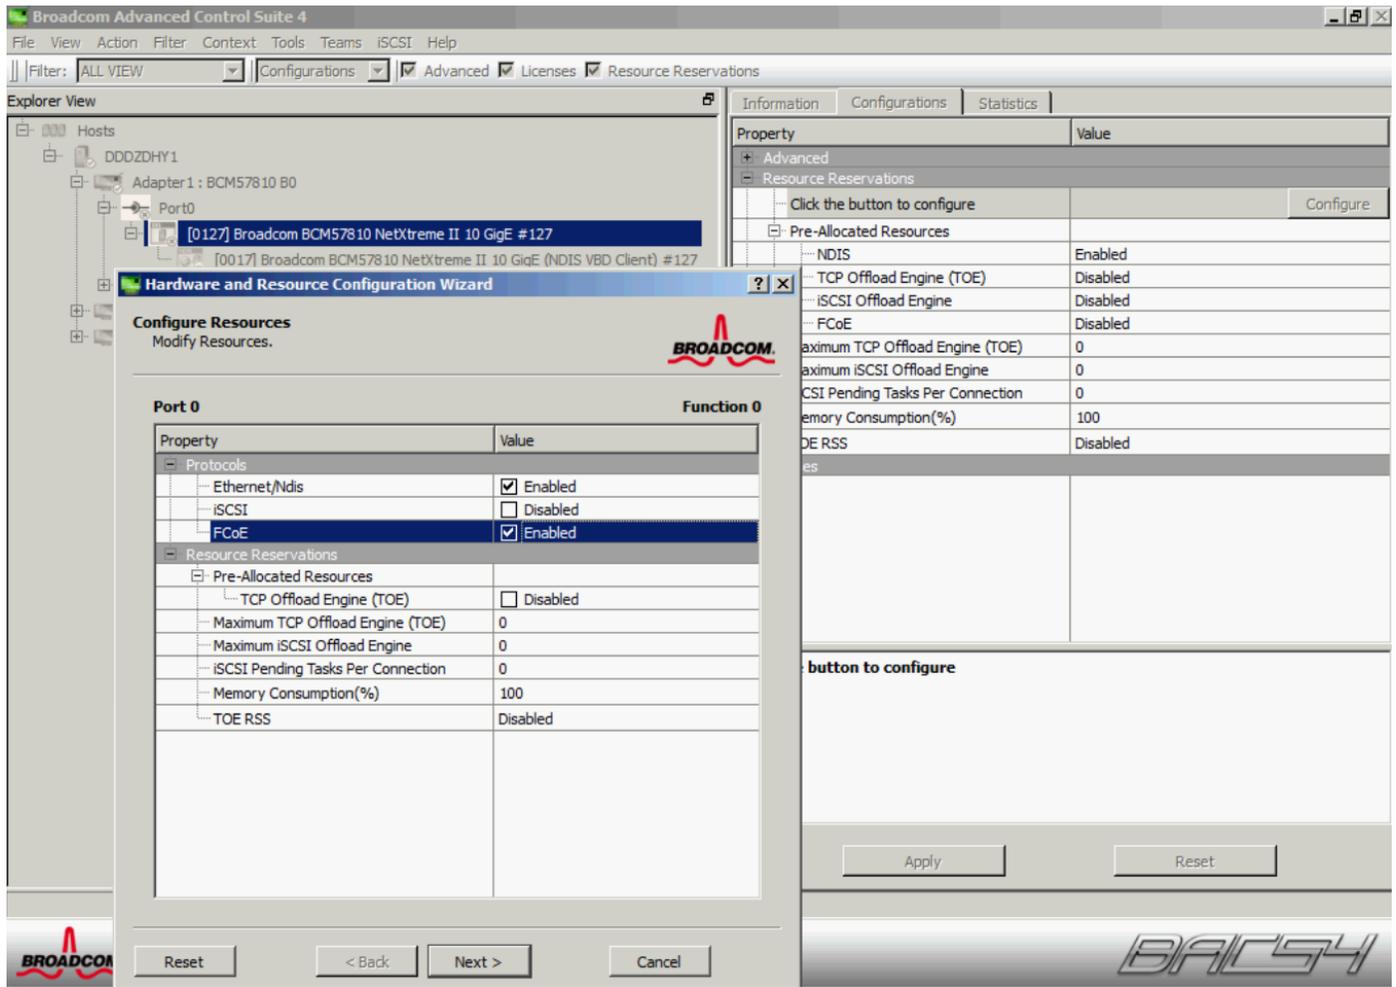
Figure 9   Enabling FCoE on Broadcom 57810 CAN

Note: As the FCoE adapter is being enabled, a message indicating the device driver was successfully installed may be displayed (Figure 10). Once complete, an virtual FCoE port will be displayed under each physical port (Figure 11).

Dell Networking FC Flex IOM: Deployment of FCoE with Dell FC Flex IOM, Brocade FC switches, and Dell Compellent Storage Array
Deployment/Configuration Guide

Figure 10   Windows Server 2008 R2 Enterprise detects Broadcom 57810 FCoE Adapter
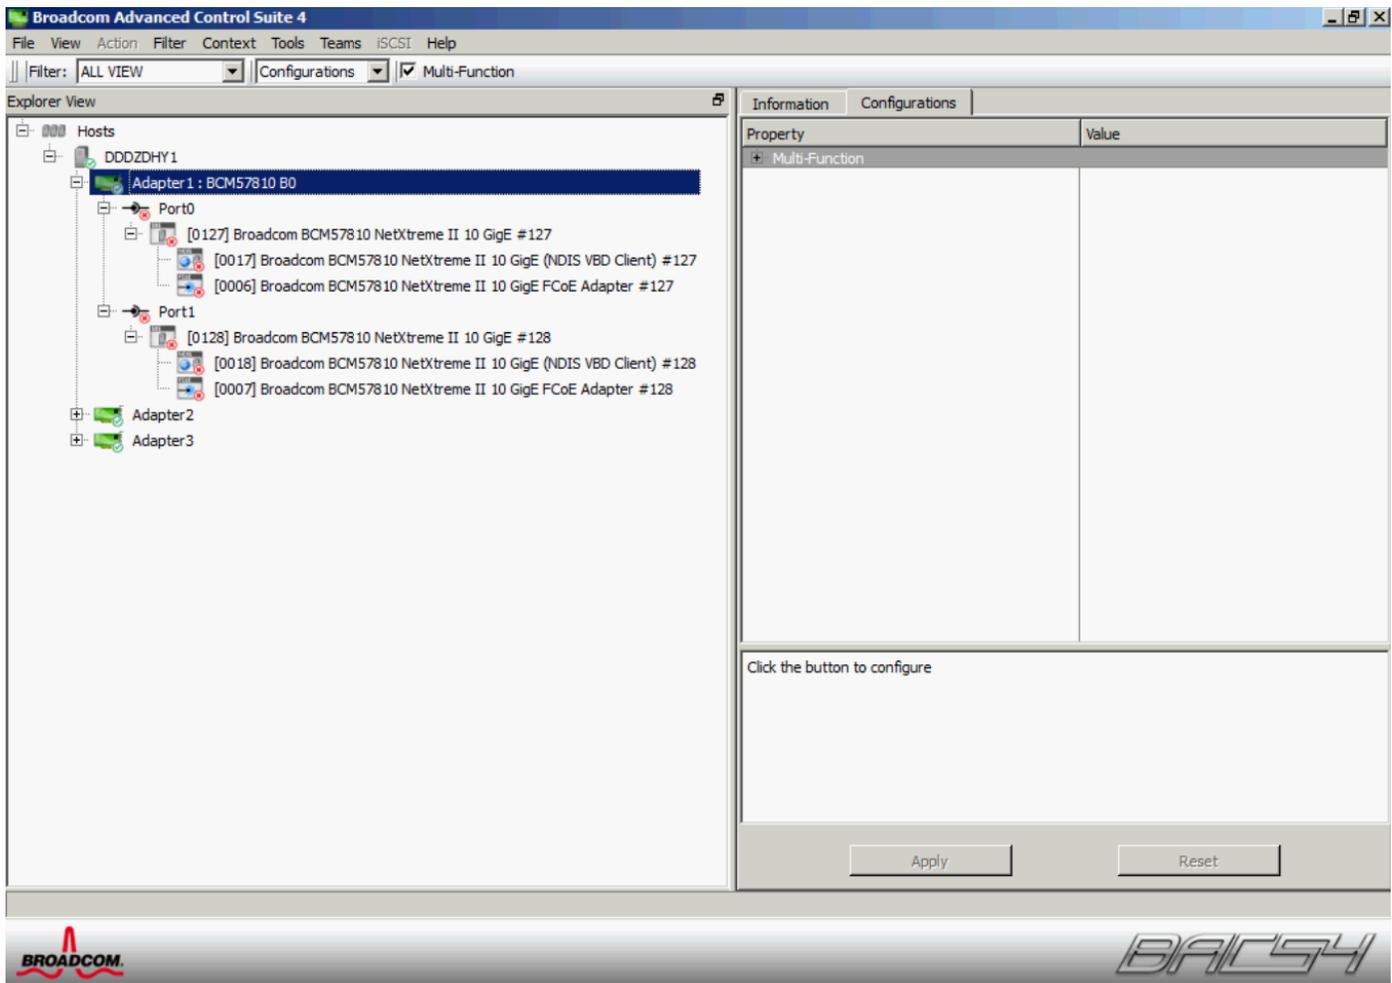
Figure 11   Broadcom 57810 CNA configured for FCoE

Since the respective FCoE settings will be pushed down from the Dell MXL switch with FC Flex IOM via

Data Center Bridging Capability Exchange protocol (DCBX), the FCoE side requires no further configuration.

6. To configure the NIC for LAN traffic, select **TEAM VIEW** from the Filter drop down box (Figure 12).
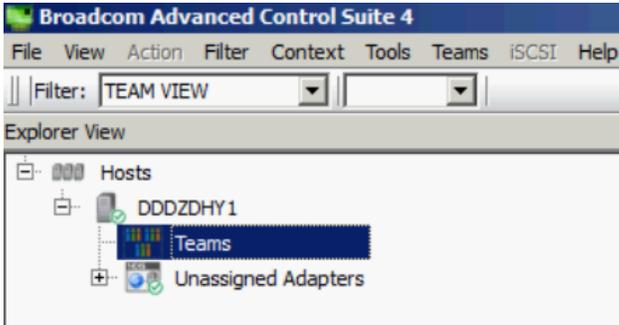
Figure 12   'Team View' in Broadcom Control Suite 4

7. Right click **Teams** and select **Create Team**. The Broadcom Teaming Wizard (Figure 13) will start.
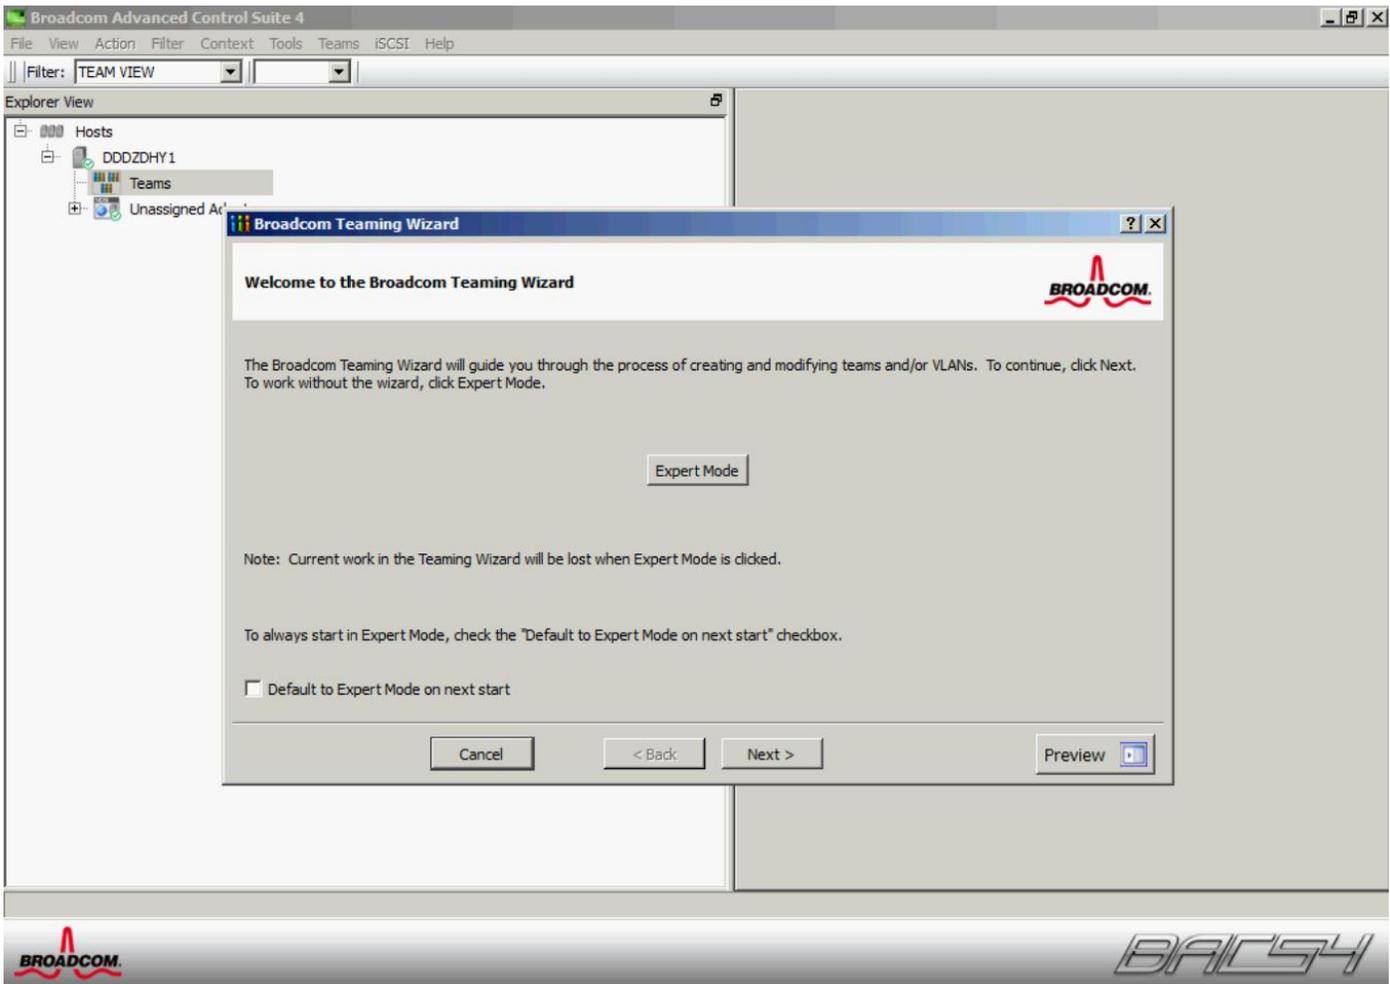
Figure 13   Creating a NIC team with Broadcom 57810 CNA

8. Click **Next**, keep the default team name (Figure 14) and click **Next** again.
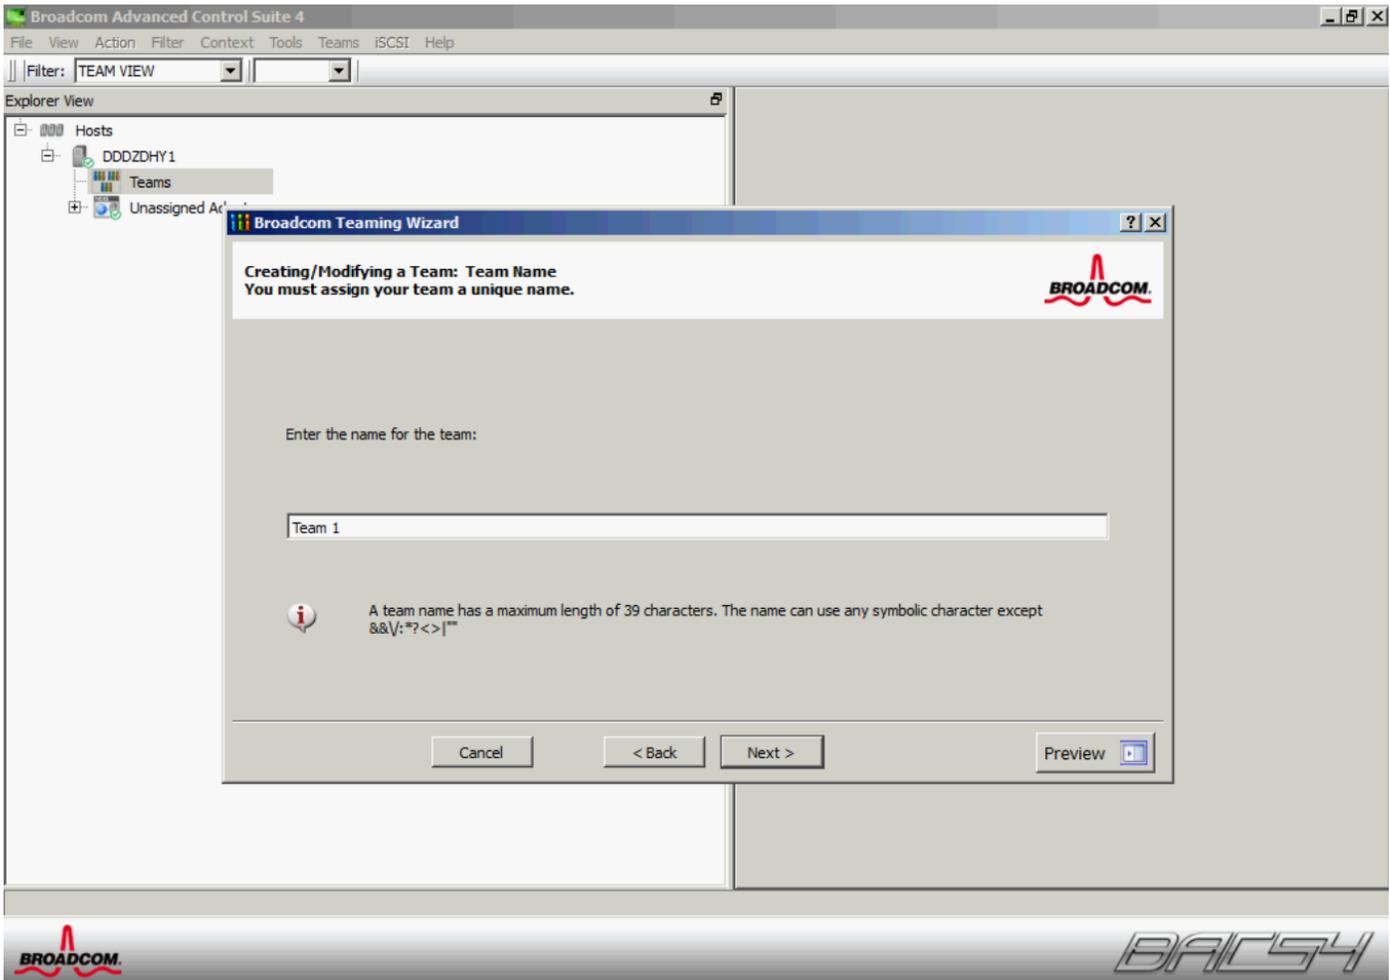
Figure 14   Naming the NIC team on Broadcom 57810 CNA

9. Click the radio button next to Smart Load Balancing™ and Failover (SLB) (Figure 15), and click next.

Note: The switch will not be aware of the NIC team and no LAG configuration will be required on upstream switches.
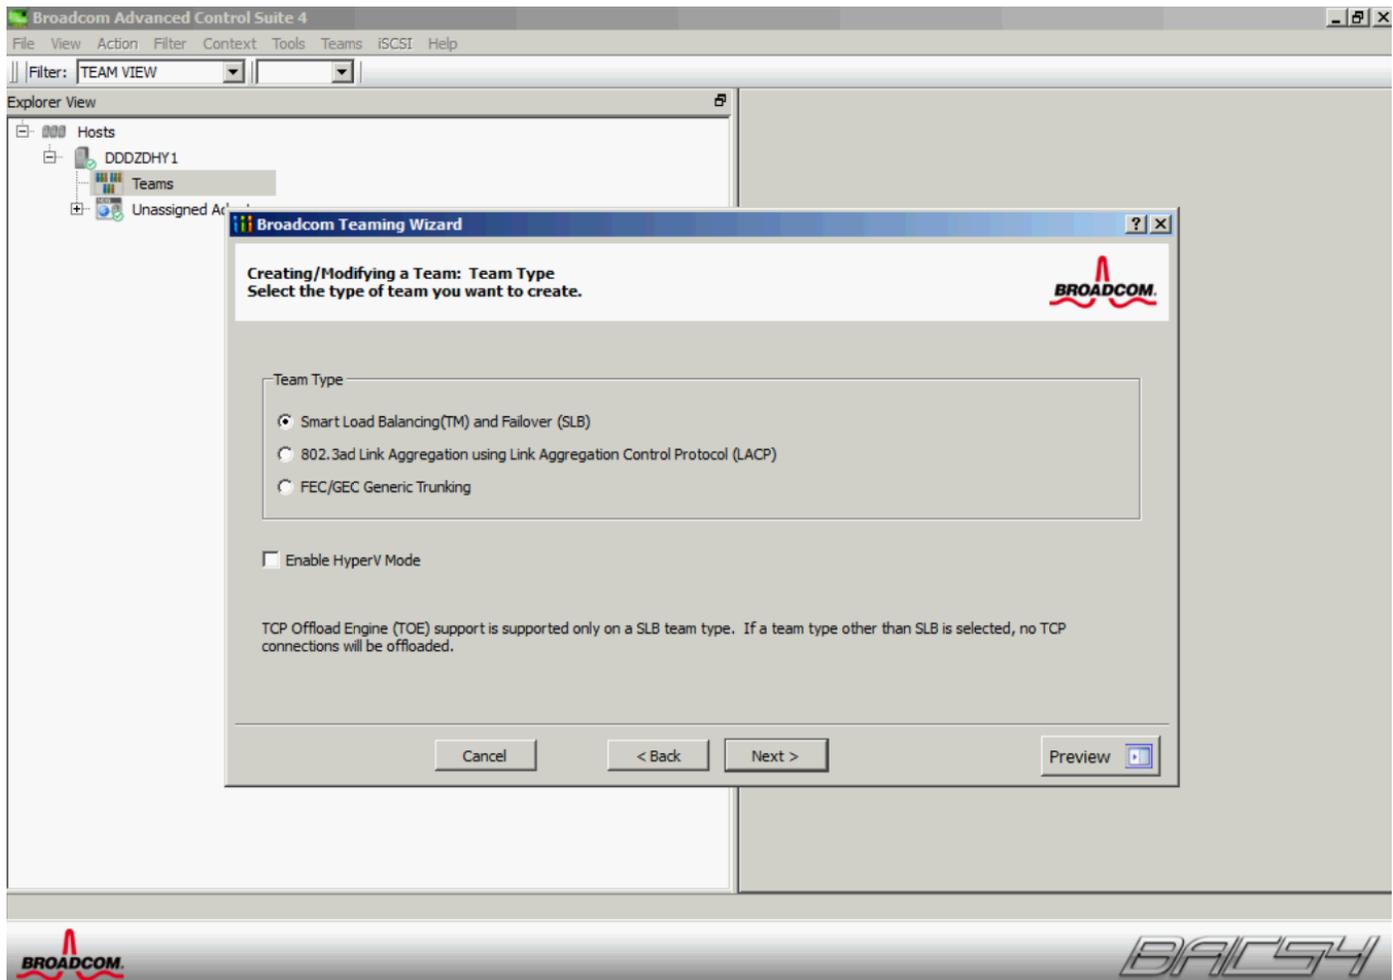
Figure 15   Selecting 'Smart Load Balancing and Failover' option for NIC teaming

10. Select one of the adapter ports to NIC team and click the **Add** button, repeat this for the second port. Both ports should now be listed under **Team Members** (Figure 16). Click **Next**.
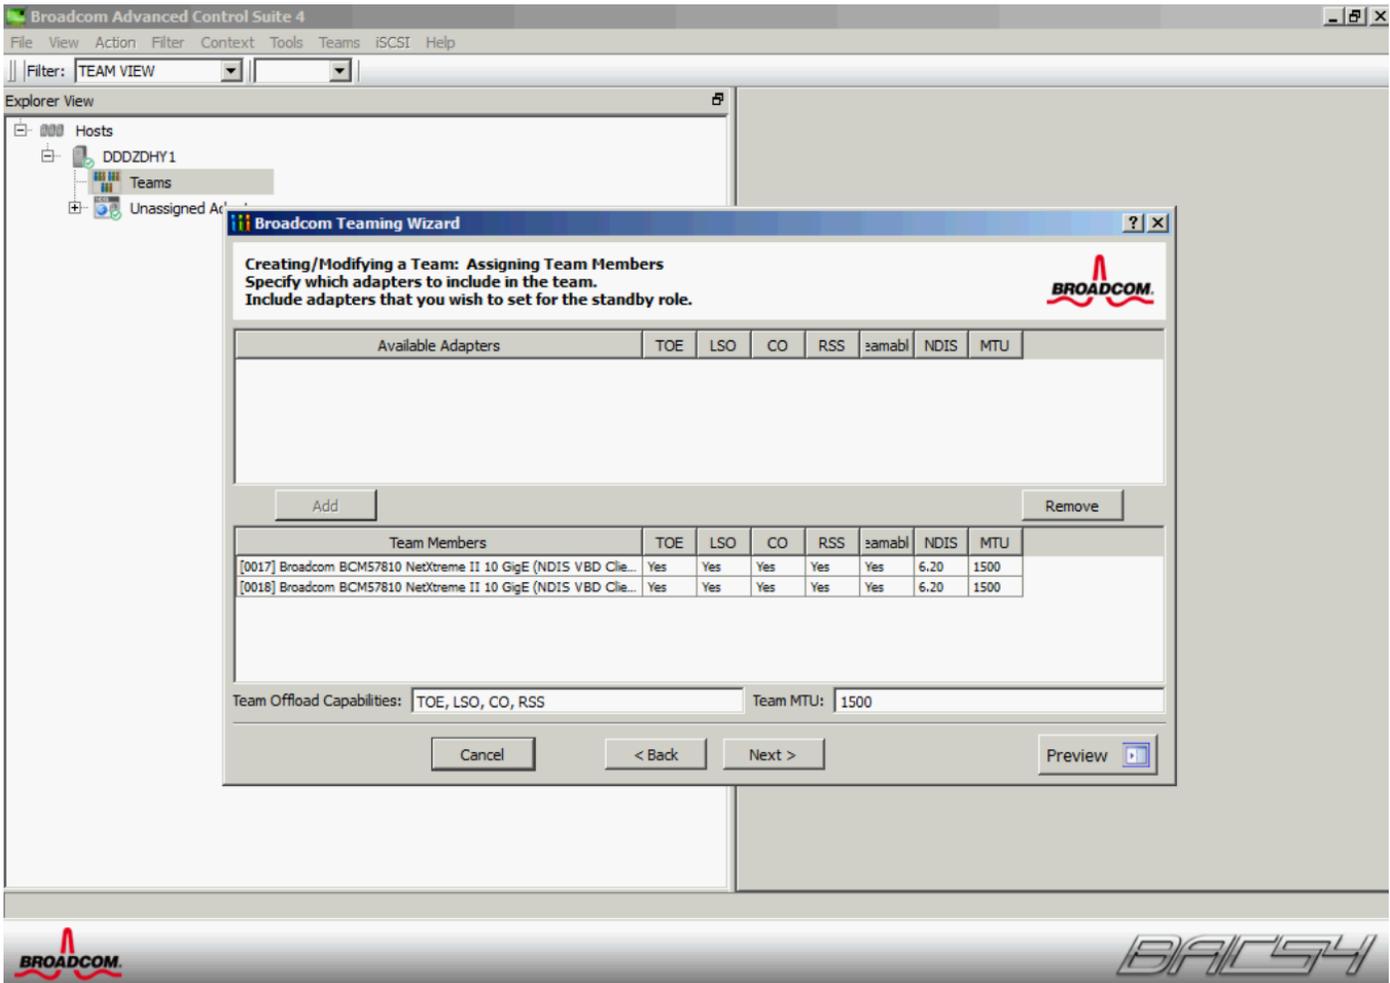
Figure 16   Broadcom teaming wizard port selection dialog for NIC teaming

11. The NIC team is completed as active-standby (Figure 17), so upon failover, the standby port will become active. Click **Next**.
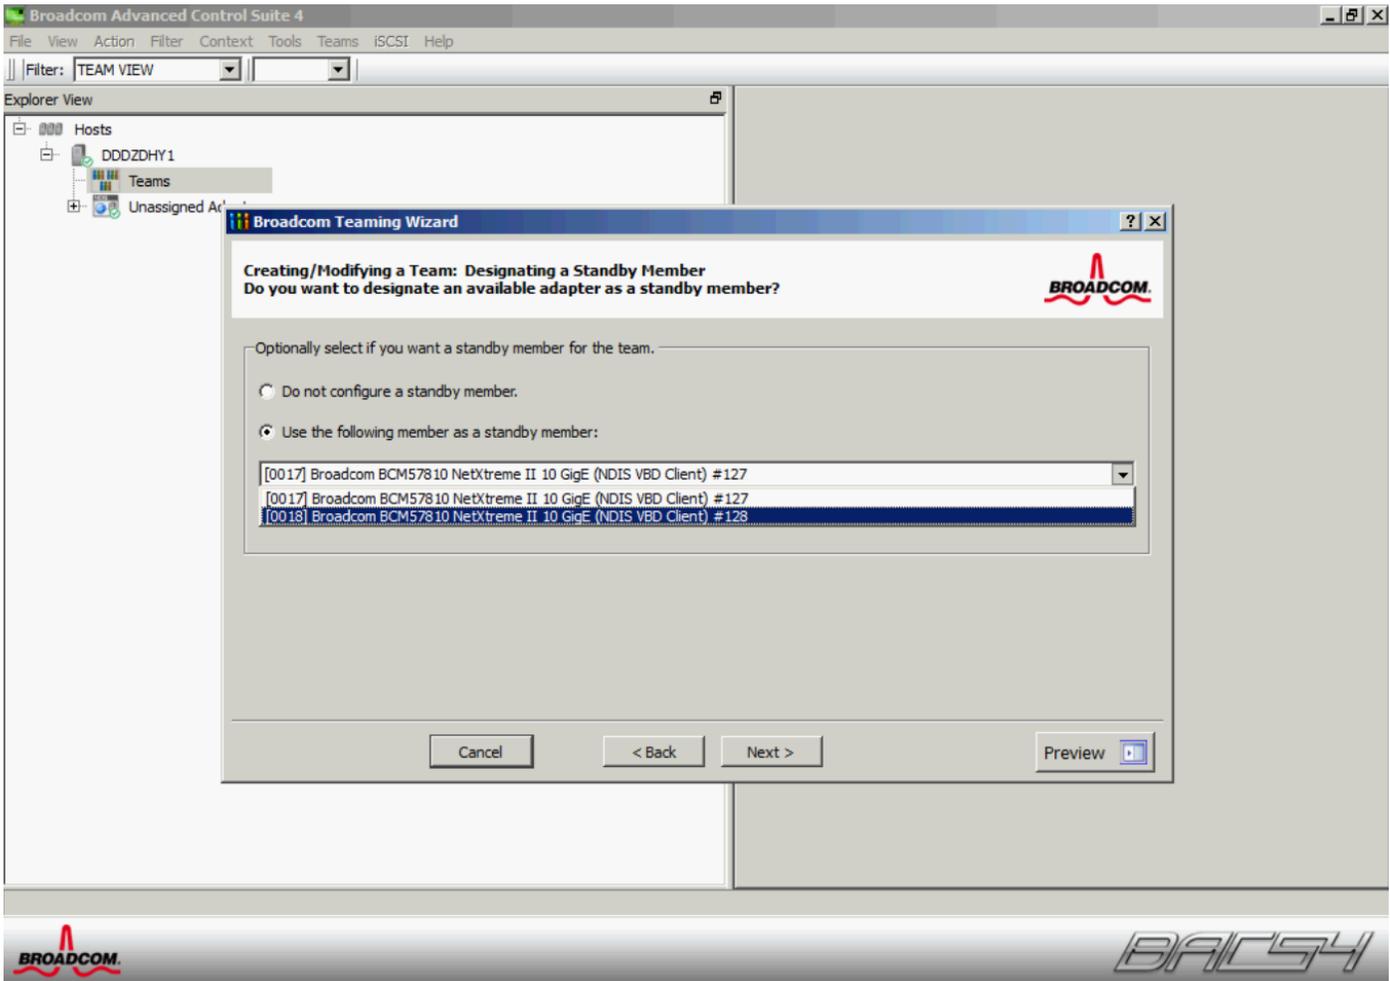
Figure 17   Selecting a 'stand-by' member of the NIC team

12. Broadcom LiveLink is a feature that minimizes any downtime due to spanning tree loop determination when failing-over. In this setup, the ports connected to the server will be configured as edge ports and not participate in spanning tree so this option is left at the default of disabled (Figure 18). Click **Next**.
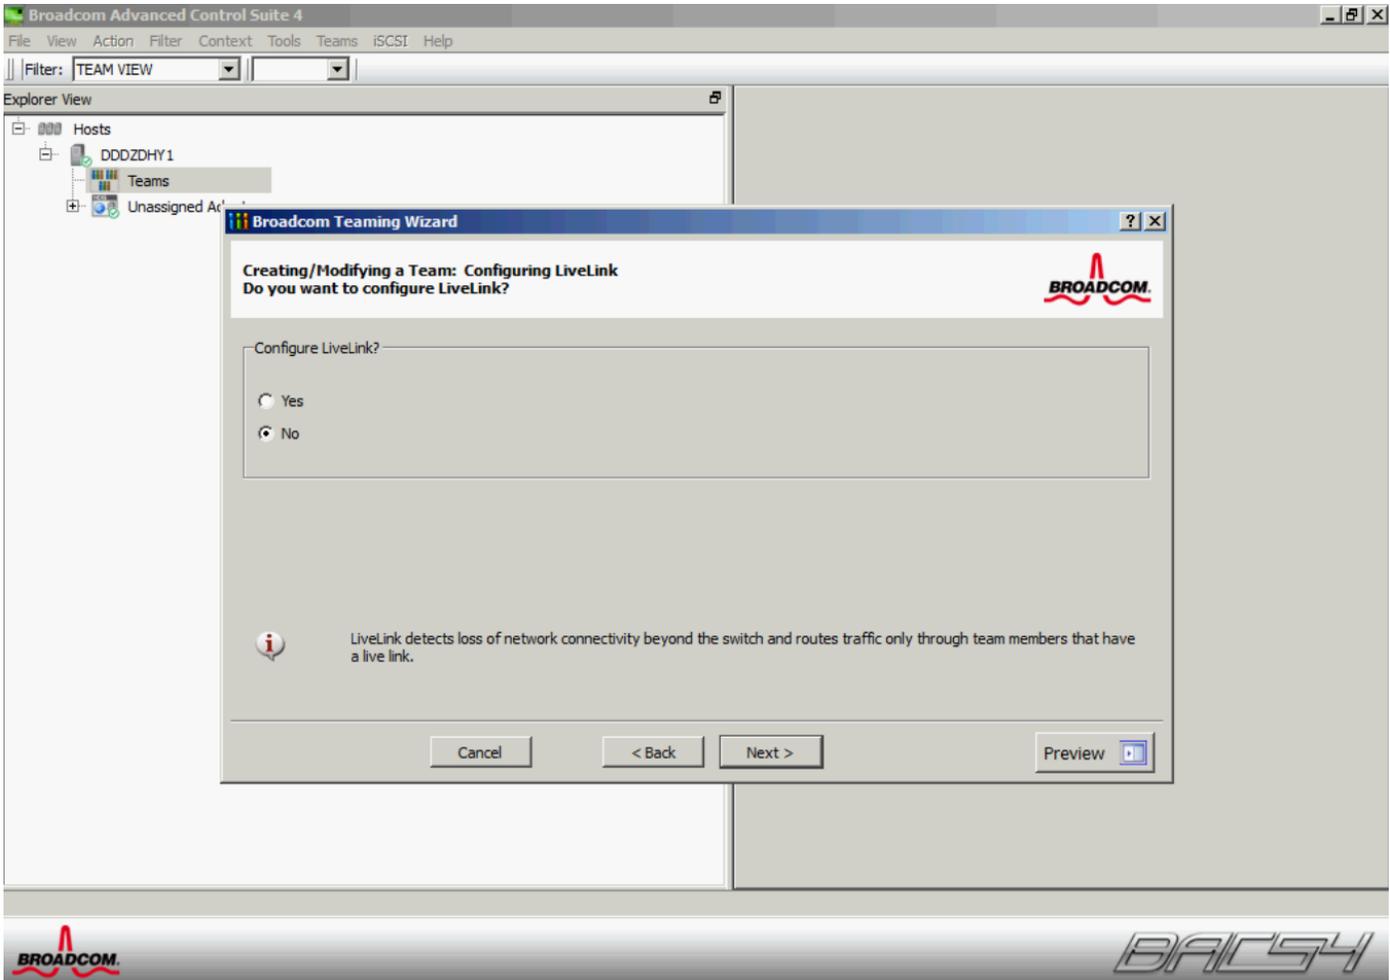
Figure 18   Selecting 'No' for Broadcom LiveLink feature

13. Select the **Add VLAN** option (Figure 19) and click **Next**.
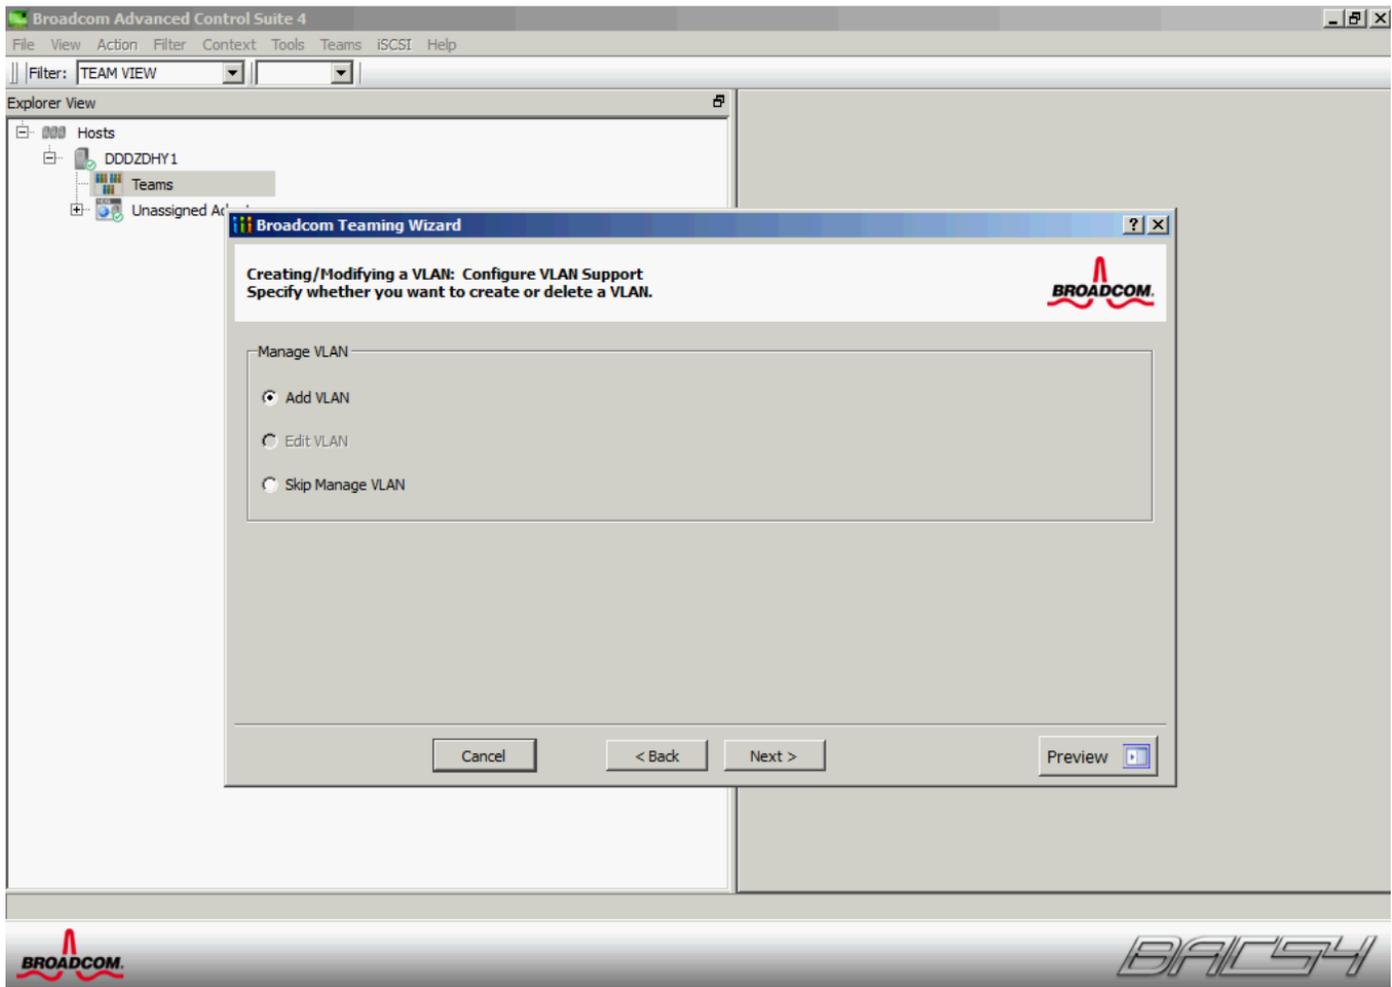
Figure 19   VLAN configuration on NIC team

14.   Enter **VLAN 5** for the VLAN Name (Figure 20) and click **Next**.
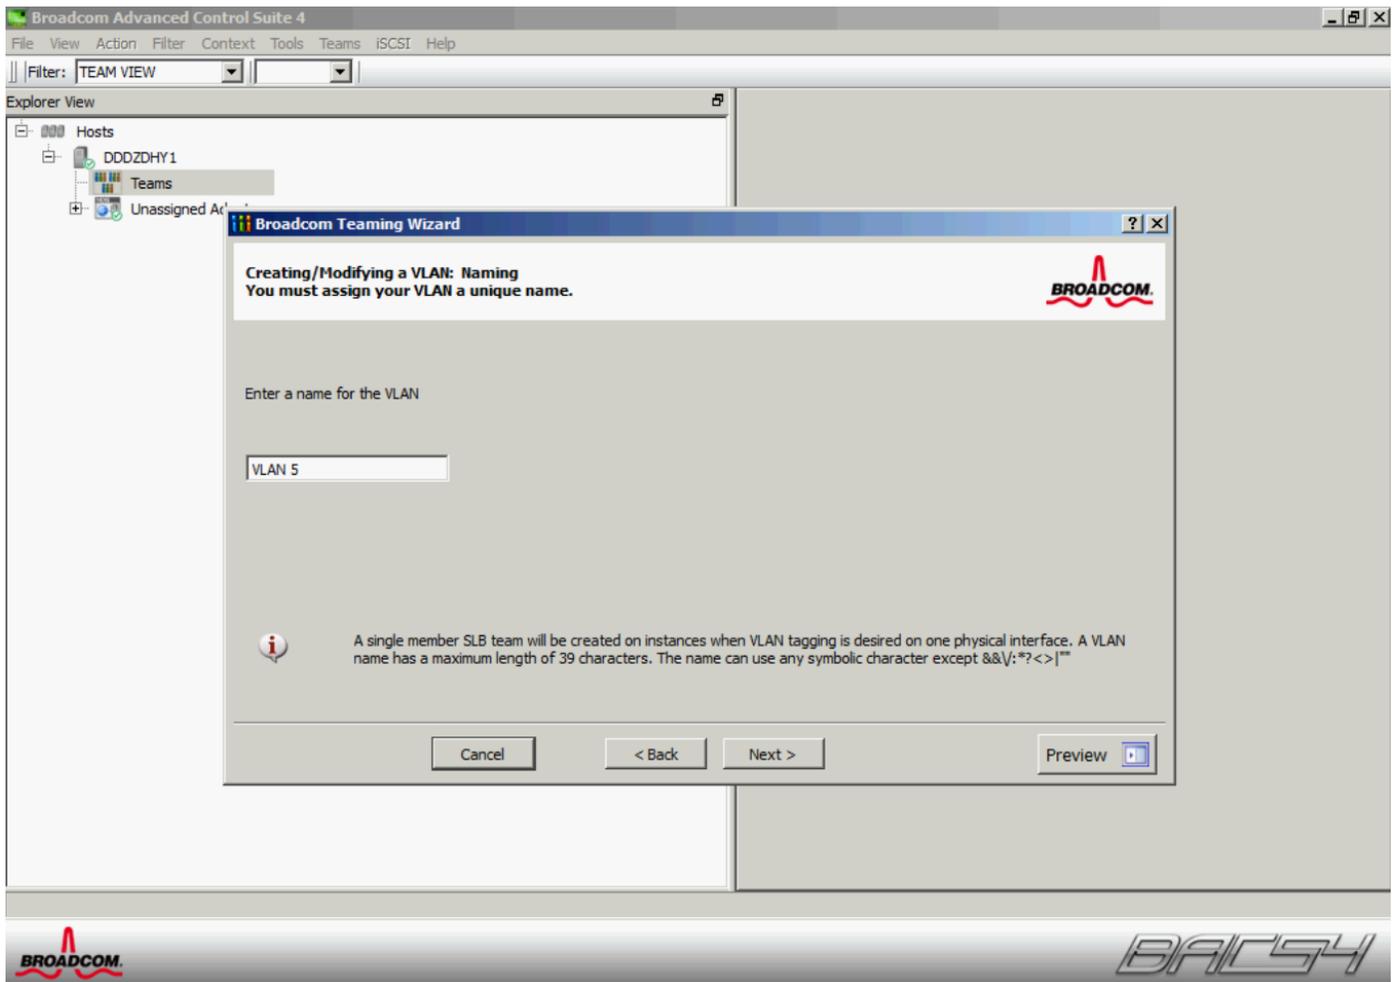
Figure 20   Naming the VLAN that will be added

15. Select the **Tagged** radio button (Figure 21) and click **Next**.
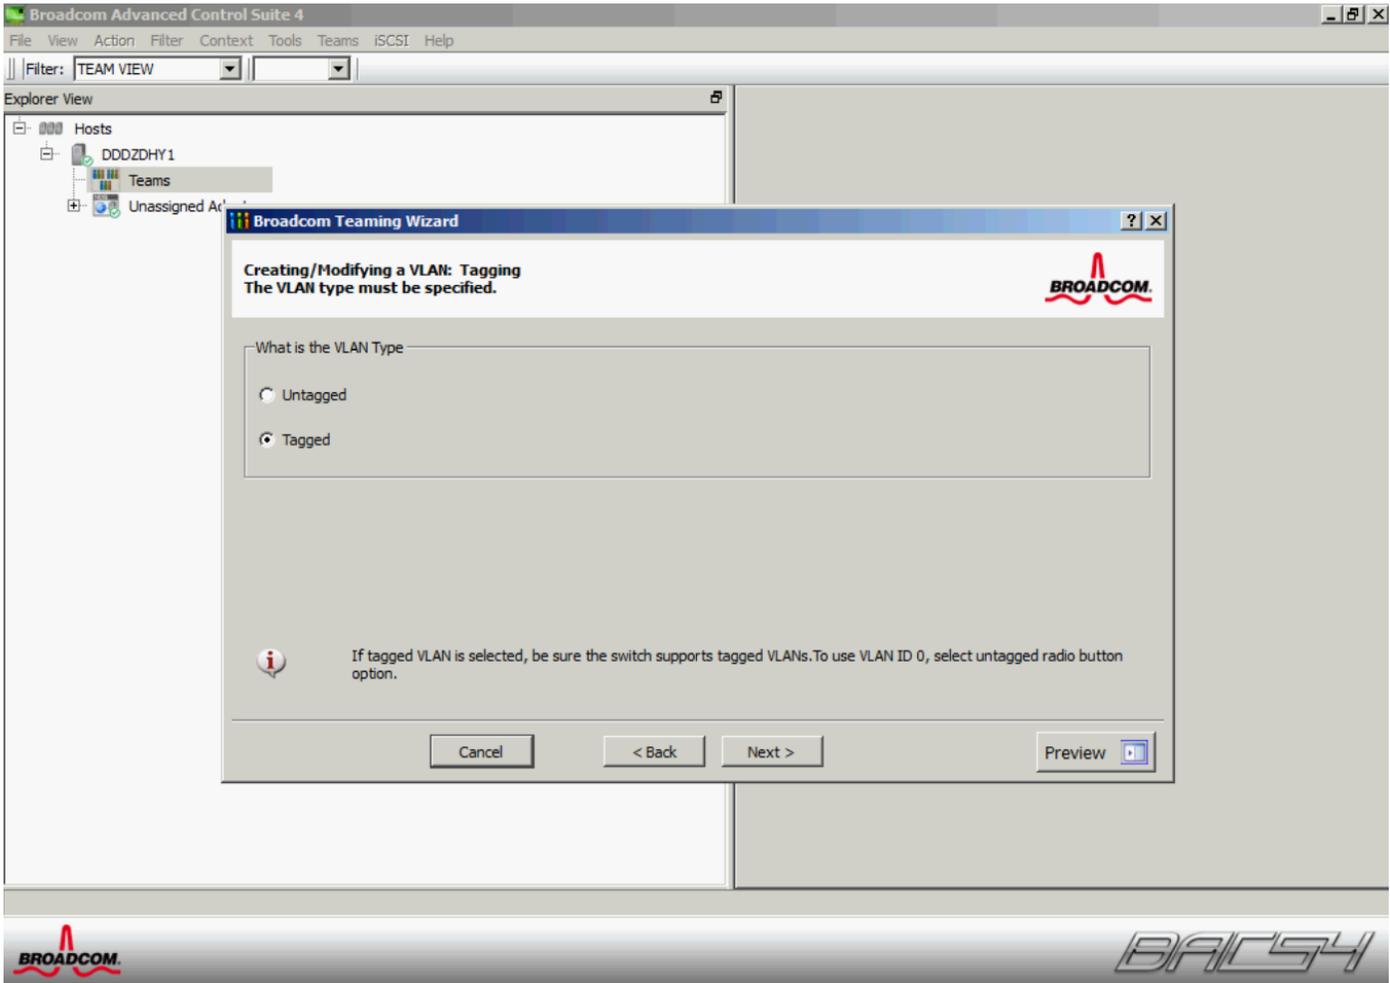
Figure 21   Selecting to have the VLAN tagged on the NIC team

Deployment/Configuration Guide

16. Enter **5** for the VLAN tag value (Figure 22) and click **Next**.
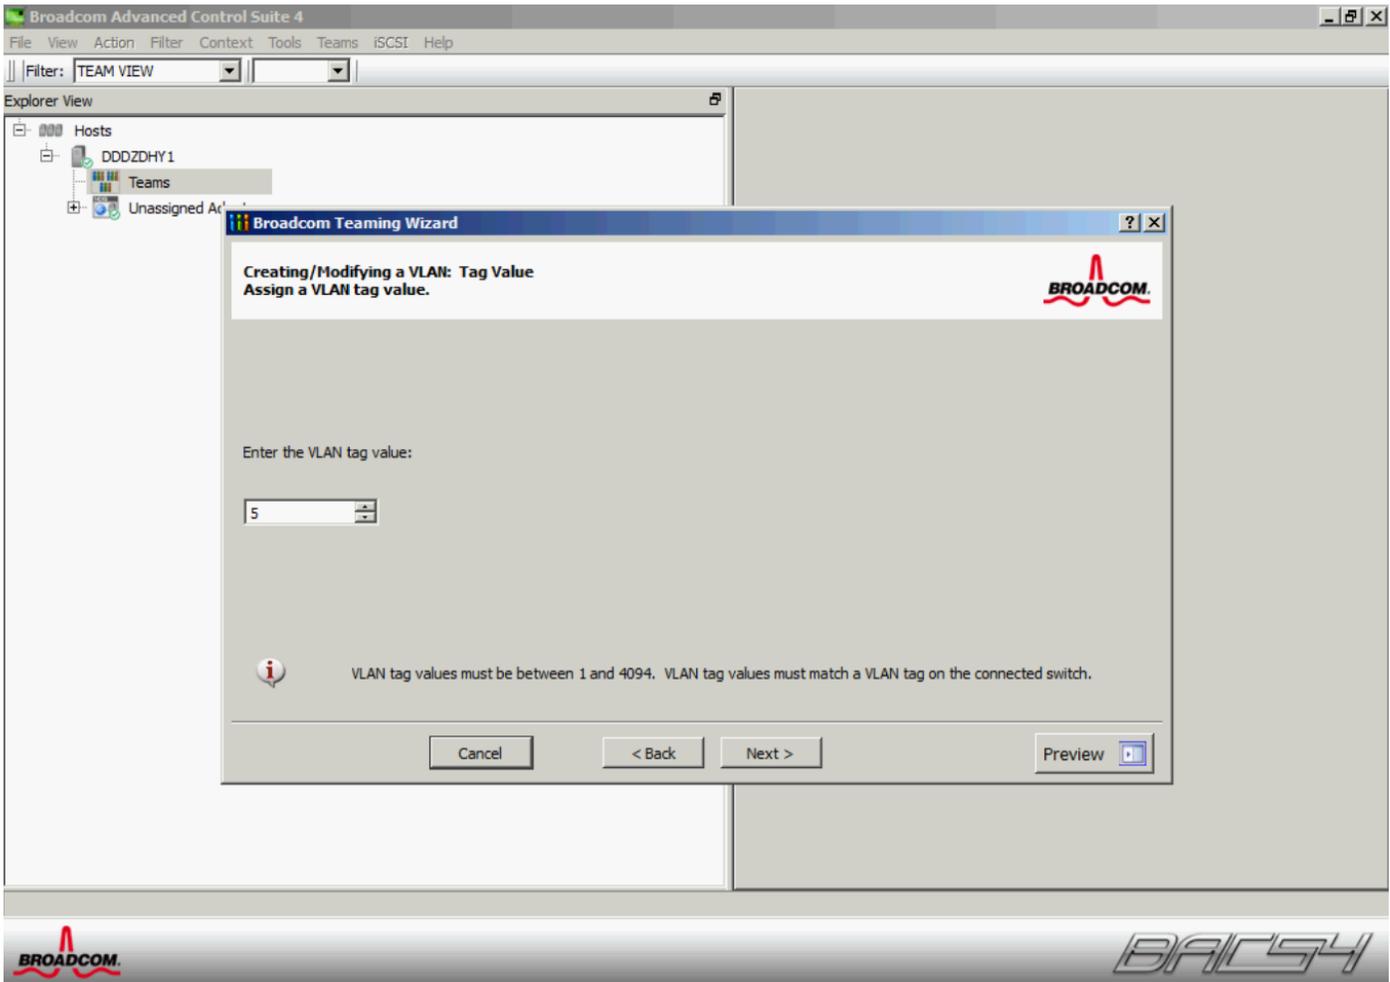
Figure 22   Selecting the correct VLAN ID

Dell Networking FC Flex IOM: Deployment of FCoE with Dell FC Flex IOM, Brocade FC switches,
and Dell Compellent Storage Array
Deployment/Configuration Guide

17. Since there are no other VLANs to manage in this example, select **No** (Figure 23) and click **Next**.
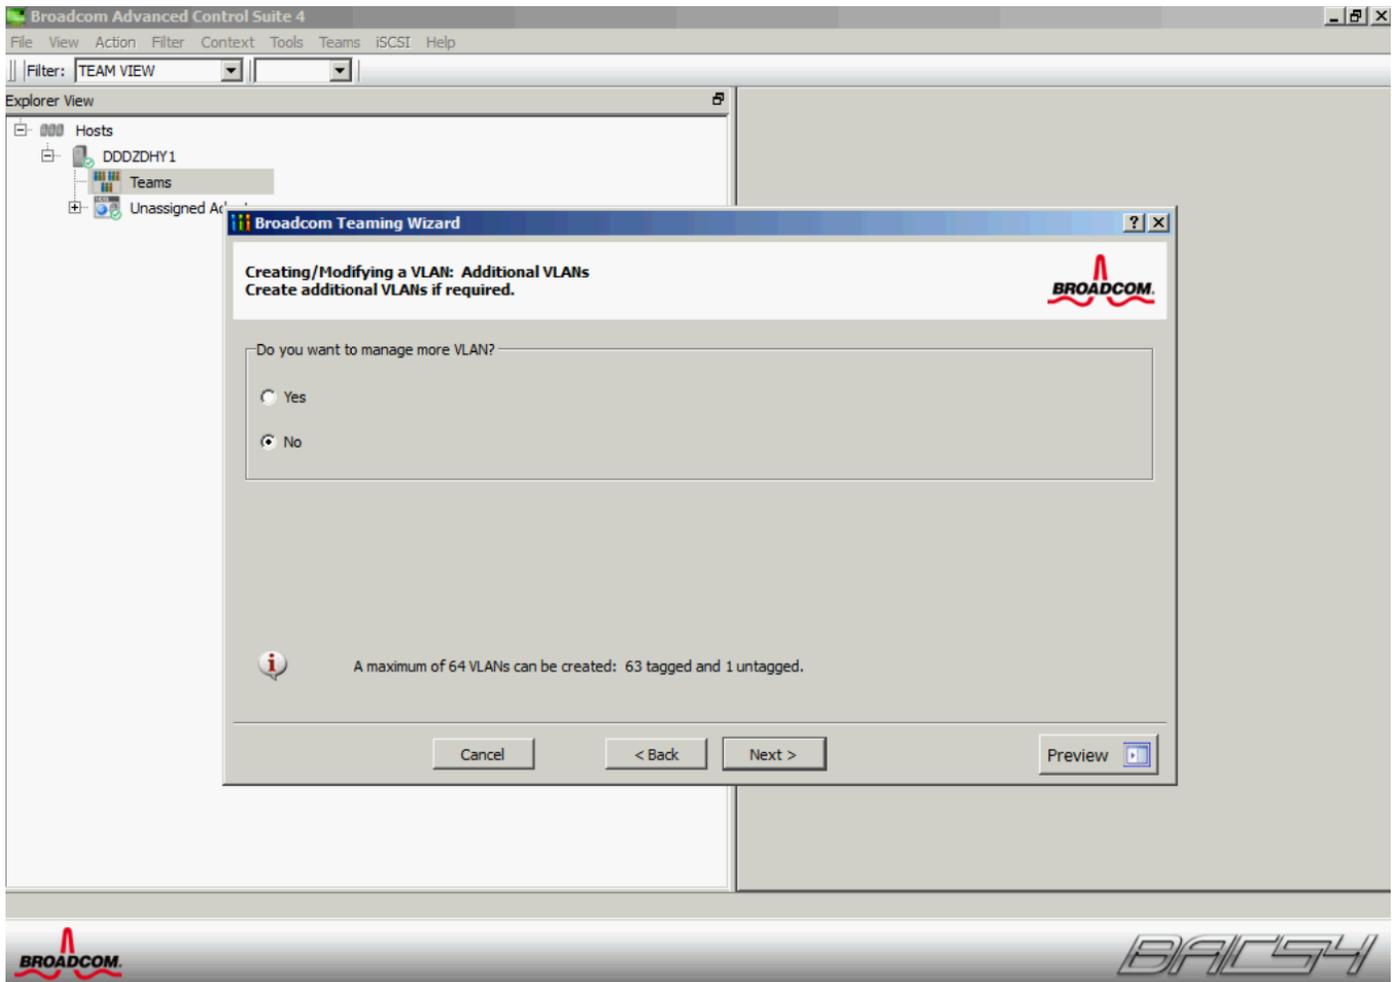
Figure 23   Confirming no other VLANs need to be added

18. Click **Finish** and then click **Yes** when prompted for confirmation (Figure 24).
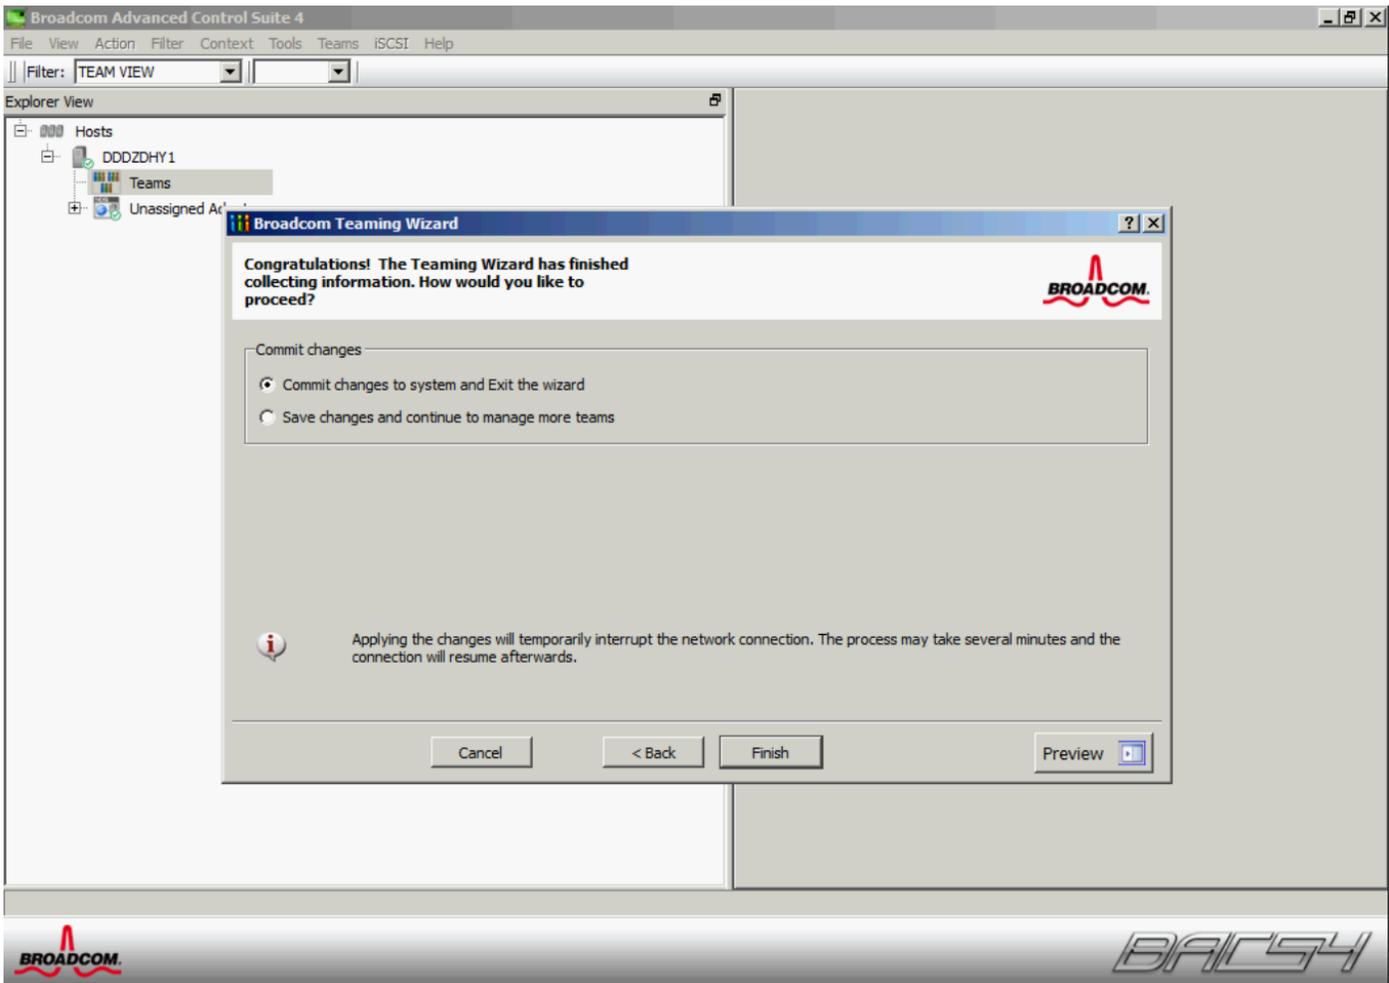
Figure 24   Committing the configuration for the NIC team

Dell Networking FC Flex IOM: Deployment of FCoE with Dell FC Flex IOM, Brocade FC switches,
and Dell Compellent Storage Array
Deployment/Configuration Guide

19. The CNA configuration for both LAN and SAN traffic is now complete. The created NIC team can be displayed by expanding the adapters under TEAM VIEW (Figure 25).
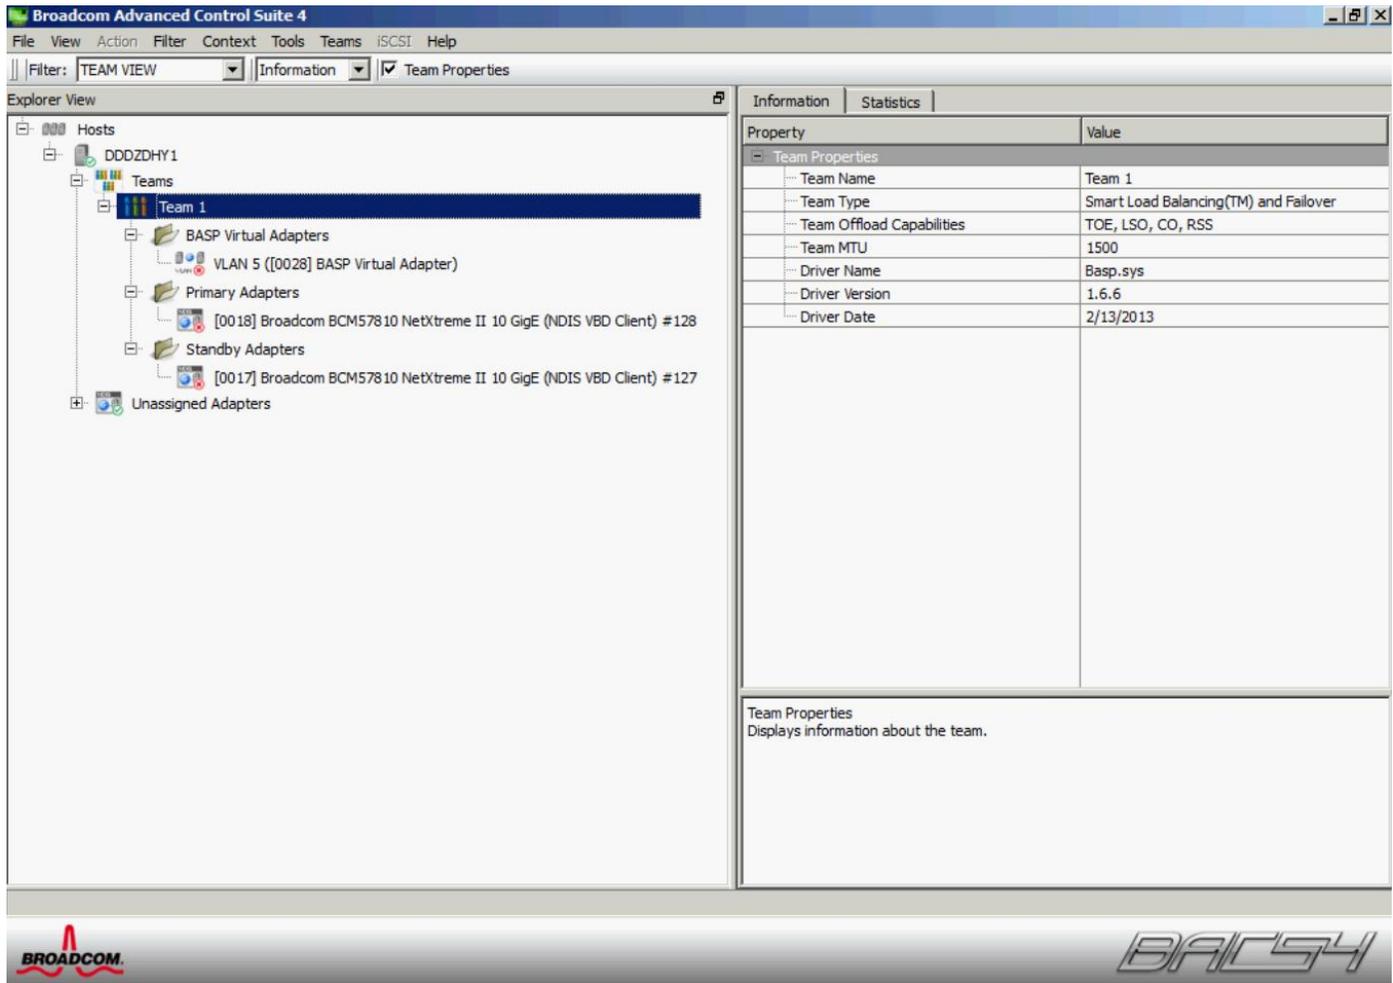
Figure 25   'TEAM VIEW' displaying the created NIC team

20. Go to **Control Panel** > **Network and Internet** > **Network** and **Sharing Center** > **Change adapter settings**.
The Network Connections are displayed. The two Broadcom 57810 CNA ports being utilized in this setup and the virtual adapter from the NIC teaming are highlighted in Figure 26.
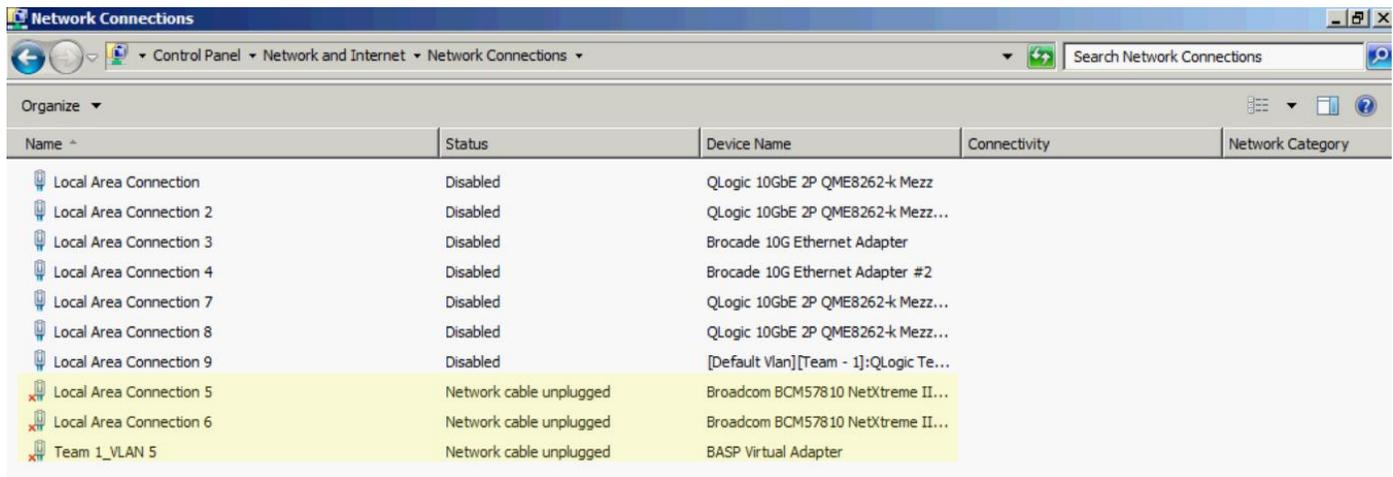
Figure 26   Viewing the NIC team and network connect in Windows Server 2008 R2 Enterprise

The correct IP address information now needs to be assigned to the teamed adapter.

21. Right click **Team1_VLAN5**, click **Properties**. Select **Internet Protocol Version 4 (TCP/IPv4)** and click the **Properties** button. Enter the respective IP information (Figure 27), and click the **OK** button.
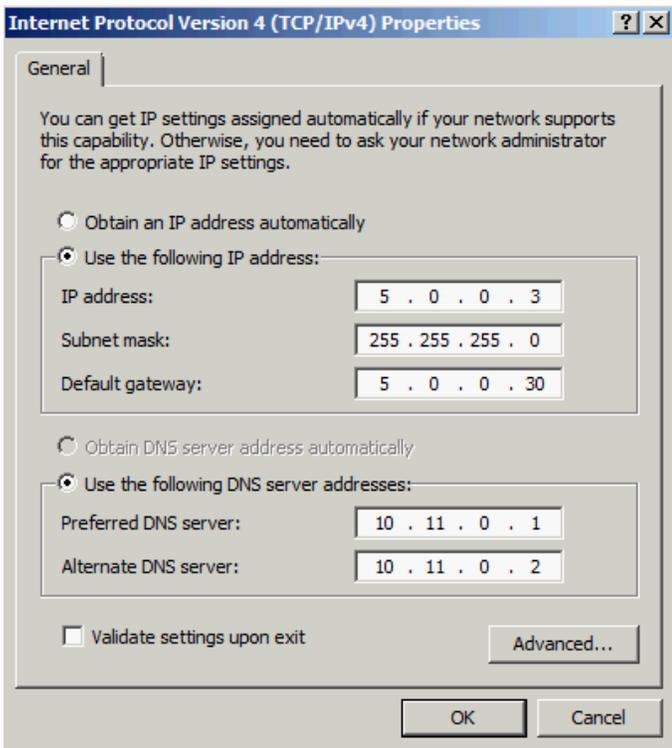
Figure 27   Configuring the NIC team interface with the correct IP information

The next step is to configure the Dell MXL switch w/ FC Flex IOM.

4.2   Dell MXL w/ FC Flex IOM Configuration

Below is the full configuration for the Dell MXL w/ FC Flex IOM for both fabric A and fabric B. Note, VLT is employed between the Dell S6000s down to the MXLs; this configuration provides active-active links of a port-channel that has members going to two separate switches for redundancy purposes.

Configuration steps:

1. Create the VLT Domain and LACP LAG up to the VLT (Enable RSTP as a best practice for configuring VLT).

2. Configure port to the CNA as a hybrid port. Create a LAN VLAN and tag it to both the **tengigabitethernet 0/4** interface going to the respective CNA and port channel going up to VLT.

3. Enable FC capability

4. Create DCB Map and configure the priority-based flow control (PFC) and enhanced transmission selection (ETS) settings for LAN and SAN traffic. Priorities or Class of Service (CoS) are mapped to priority groups using the **priority-pgid** command. In this example, priorities 0, 1, 2, 4, 5, 6, and 7 are mapped to priority group 0. Priority 3 is mapped to priority-group 1.

5. Create FCoE VLAN

6. Next, create a FCoE MAP so FCoE traffic is mapped to the respective VLAN. The FCoE MAP is applied to both the **tengigabitethernet 0/4** interface going to the respective CNA port and to the FC interface connecting to the FC switch. Note, on MXL/IOA w/ FC Flex IOM, FCoE is always mapped to priority 3.

7. Apply DCB map to downstream interface going to server.

8. Configure uplink failure detection.

The same procedure is repeated for the Dell MXL blade switch w/ FC Flex IOM connecting to fabric B. Note a different **fc-map** and FCoE VLAN should be used for each fabric.

Especially important to note is the fact that the same Ethernet port on the MXL where the FCoE MAP is applied is also untagged on the default VLAN 1. This is needed because the FIP protocol communicates over the default VLAN to discover the FCoE VLAN. The LAN traffic is tagged on VLAN 5.

Dell MXL Configuration (Fabric A)

/* Enable RSTP (Enabled due to VLT config on S6000s) */
> enable
> config terminal
> protocol spanning-tree rstp
> no disable
> exit

/* Configure Manangement Interface */
> interface management 0/0
> ip address 10.11.129.166/20
> no shut

/* Create LACP LAG for connecting to VLT LAG down from S6000s */
> interface range fortyGigE 0/33 - 37
> port-channel-protocol LACP
> port-channel 10 mode active
> no shutdown
> exit

/* Create LAN VLAN and tag interfaces */
> interface port-channel 10
> switchport
> no shutdown
> exit

> interface tengigabitethernet 0/4
> portmode hybrid
> switchport
> no shutdown
> exit

> interface vlan 5
> tagged tengigabitethernet 0/4
> tagged port-channel 10
> exit

/* Enable FC capability */
> feature fc

Dell MXL Configuration (Fabric A) cont.

```
/* Create DCB MAP */
> dcb-map SAN_DCB_MAP
> priority-group 0 bandwidth 60  pfc off
> priority-group 1 bandwidth 40 pfc on
> priority-pgid 0 0 0 1 0 0 0 0
> exit

/* Create FCoE VLAN */
> interface vlan 1002
> exit

/* Create FCoE MAP */
> fcoe-map SAN_FABRIC_A
> fabric-id 1002 vlan 1002
> fc-map 0efc02
> exit

/* Apply FCoE MAP to interface */
> interface range fibreChannel 0/41 - 44
> fabric SAN_FABRIC_A
> no shutdown
> exit

/* Apply FCoE MAP and DCB MAP to interface */
> interface tengigabitethernet 0/4
> dcb-map SAN_DCB_MAP
> fcoe-map SAN_FABRIC_A

/* Make server port an edge-port */
> spanning-tree rstp edge-port
> exit

/* Configure Uplink Failiure Detection */
> uplink-state-group 1
> downstream TenGigabitEthernet 0/4
> upstream Port-channel 10
> downstream auto-recover
> end

> write
```

Dell MXL Configuration (Fabric B)

/* Enable RSTP (Enabled due to VLT config on S6000s) */
```
> enable
> config terminal
> protocol spanning-tree rstp
> no disable
> exit
```

/* Configure Manangement Interface */
```
> interface management 0/0
> ip address 10.11.129.167/20
> no shut
```

/* Create LACP LAG for Connecting to VLT LAG down from S6000s */
```
> interface range fortyGigE 0/33 - 37
> port-channel-protocol LACP
> port-channel 11 mode active
> no shutdown
> exit
```

/* Create LAN VLAN and tag interfaces */
```
> interface port-channel 11
> switchport
> no shutdown
> exit

> interface tengigabitethernet 0/4
> portmode hybrid
> switchport
> no shutdown
> exit

> interface vlan 5
> tagged tengigabitethernet 0/4
> tagged port-channel 11
> exit
```

/* Enable FC capability */
```
> feature fc
```

Dell MXL Configuration (Fabric B) cont.

```
/* Create DCB MAP */
> dcb-map SAN_DCB_MAP
> priority-group 0 bandwidth 60  pfc off
> priority-group 1 bandwidth 40 pfc on
> priority-pgid 0 0 0 1 0 0 0 0
> exit

/* Create FCoE VLAN */
> interface vlan 1003
> exit

/* Create FCoE MAP */
> fcoe-map SAN_FABRIC_B
> fabric-id 1003 vlan 1003
> fc-map 0efc03
> exit

/* Apply FCoE MAP to interface */
> interface range fibreChannel 0/41 - 44
> fabric SAN_FABRIC_B
> no shutdown
> exit

/* Apply FCoE MAP and DCB MAP to interface */
> interface tengigabitethernet 0/4
> dcb-map SAN_DCB_MAP
> fcoe-map SAN_FABRIC_B

/* Make server port an edge-port */
> spanning-tree rstp edge-port
> exit

/* Configure Uplink Failiure Detection */
> uplink-state-group 1
> downstream TenGigabitEthernet 0/4
> upstream Port-channel 11
> downstream auto-recover
> end

> write
```

4.3 Dell Compellent Storage Array Configuration

The Dell Compellent Storage Center controllers are used to support various I/O adapters including FC, iSCSI, FCoE, and SAS. A Dell Compellent Storage Center consists of one or two controllers, FC switches, and one or more enclosures. In the current example, two Compellent SC8000 controllers, one Compellent SC220 enclosure, two FC switches, and one 4-port FC HBA card on each Compellent controller is used for the SAN network. The FC switches provide robust connectivity to servers, allowing for the use of multiple controllers and redundant transport paths.

SAS enclosures hold disks for data storage and connect to the controllers through back-end ports via SAS cables; you can see how the SC220 enclosure and controllers are cabled together in Figure 5 prior in this document.

To keep the diagram uncluttered yet detailed, the only connections not shown are the **eth0** ports on each controller connecting to the management network and the **eth1** port on each controller connecting to the **eth1** port on the other controller. The **eth0** connection supports system login and access for the software. It's used to send emails, alerts, SNMP traps, and Phone Home data. The **eth1** connection is used for dedicated Inter-Process Communication (IPC) between controllers in a dual-controller Storage Center. There is no default gateway for **eth1** and it does not need to be set. See the "CT-SC040 and SC8000 Connectivity Guide" and "Compellent Storage Center System Setup Guide" to get started on cabling and configuring your Compellent storage array.

In this example setup, two SC8000 controllers and one SC220 disk enclosure have been cabled together. There are two paths available from the server to the FC switches and four paths available from each FC switch to the Compellent storage array.

1. During initial configuration of the Compellent Storage Center, a disk pool labeled **Pool_1** was created consisting of seven 300 GB drives. The total disk space is 1.64 TB; this can be seen in the screen shot of the **Storage Center System Manager** GUI as shown below in Figure 28.
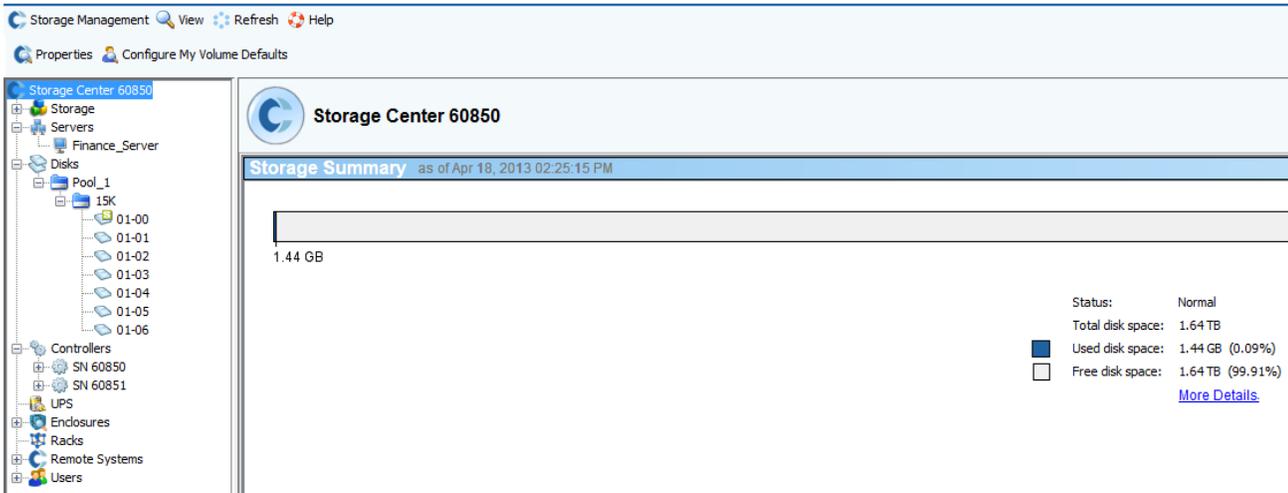
Figure 28   Storage Center System Manager GUI displays disk pool "Pool_1" with 1.64 TB free space

2. Since there are two fabrics, fabric A and fabric B, two fault domains are created. Domain 1 is already created by default and all the FC ports are currently in domain 1. To create another domain, click **Storage Management** on the top left of the webpage and then select **System > Setup > Configure Local Ports**.

3. Next, click the **Edit Fault Domains** button at the bottom right of the dialog box.

4. On the next dialog box click the **Create Fault Domain** button on the lower right of the dialog box. In the **Name** field type a name for the new domain. In this case, the name **Domain 2** was used. Make sure **FC** is selected in the **Type** field and click **Continue**. Figure 29 below shows a sceenshot of the created second domain.
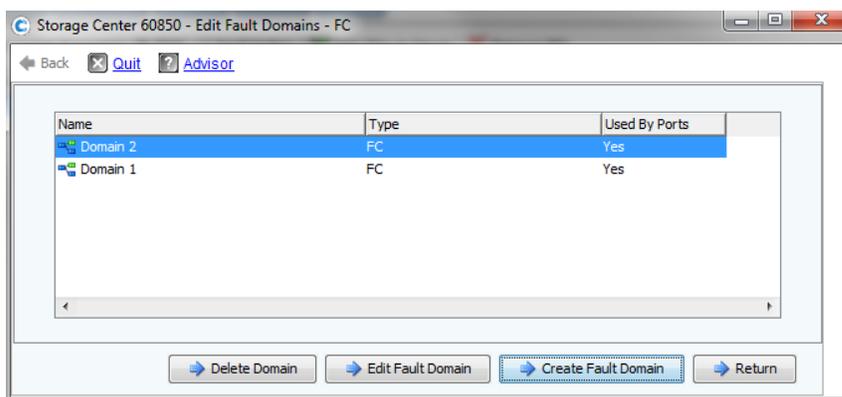
Figure 29   Creating an additional Fault Domain on Dell Compellent Storage Array

5. Now, navigate back to the **Configure Local Ports** dialog and select the appropriate domain to put each port in. Each fabric should be in its own Domain; below all ports going to fabric A are put in Domain 1 and all ports going to fabric B are put in Domain 2.
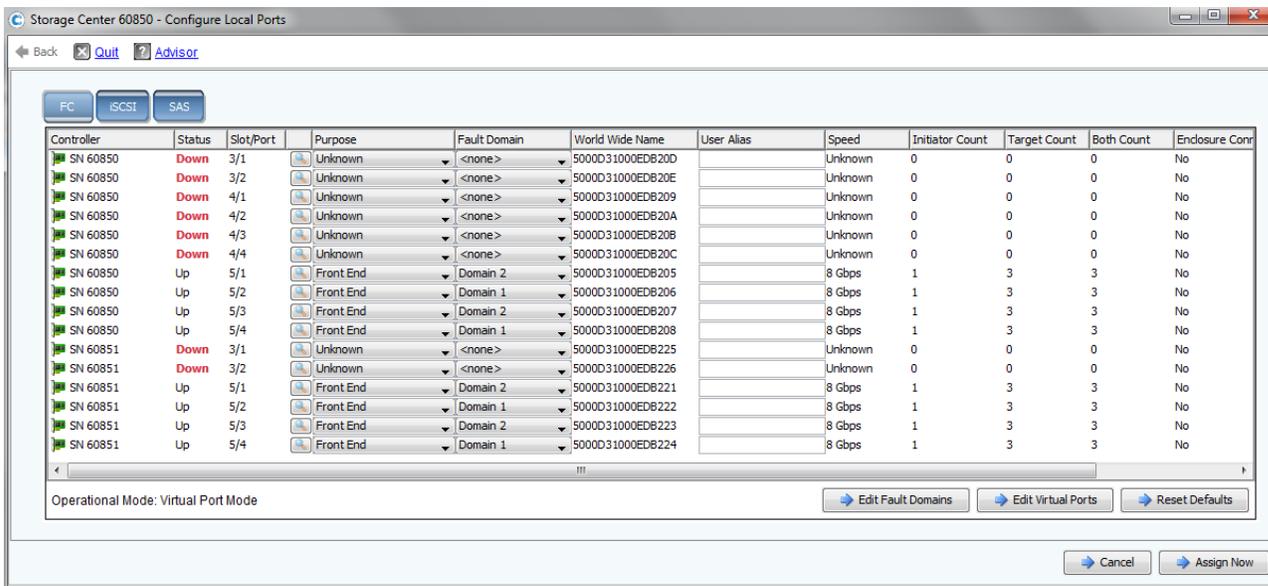
Figure 30   Assigning ports on Compellent Storage to respective Fault Domains

Note: If you get a warning that paths are not balanced, navigate to the left-hand pane, right click **Controllers** and select **Rebalance Local Ports**.

Next, a server object needs to be created and the respective FC ports have to be selected to be used by the server object.

6. Right click **Servers** on the left pane and select **Create Server**.

In Figure 31 below, you can see a server object named "Finance_Server" was created that includes both of the virtual FCoE ports on the CNA card; the storage array sees the ports as FC ports.
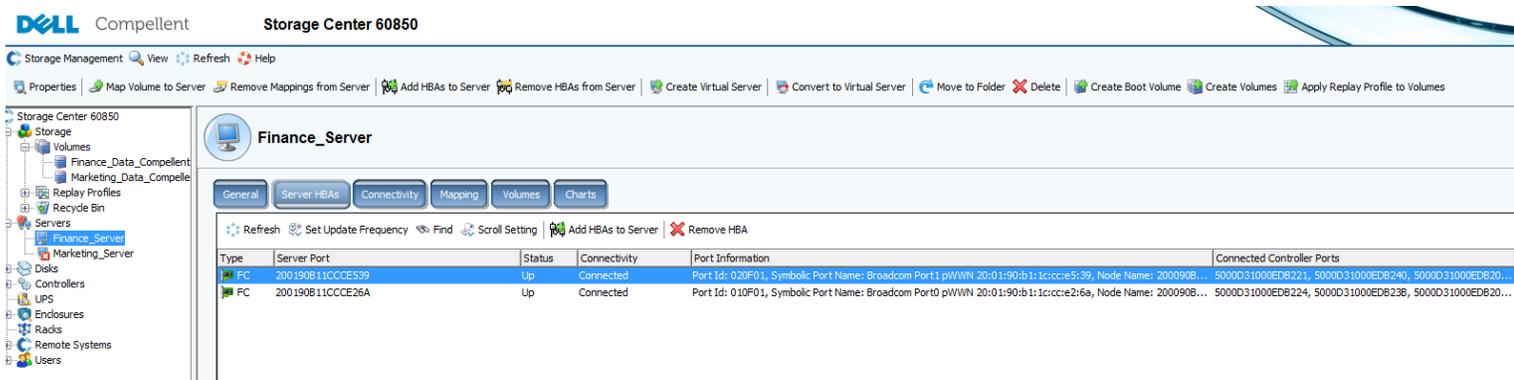
Figure 31   Added Dell PowerEdge Server HBAs to 'Server Object' on Dell Compellent Storage Array

The next step is to enable mulipathing on Windows Server 2008 R2 Enterprise.

7.   Navigate to **Start > Administrative Tools > Server Manager > Features > Add Features** and select **Multipath I/O**.

You can see in Figure 32 below that the **Multipath I/O** feature has been installed.
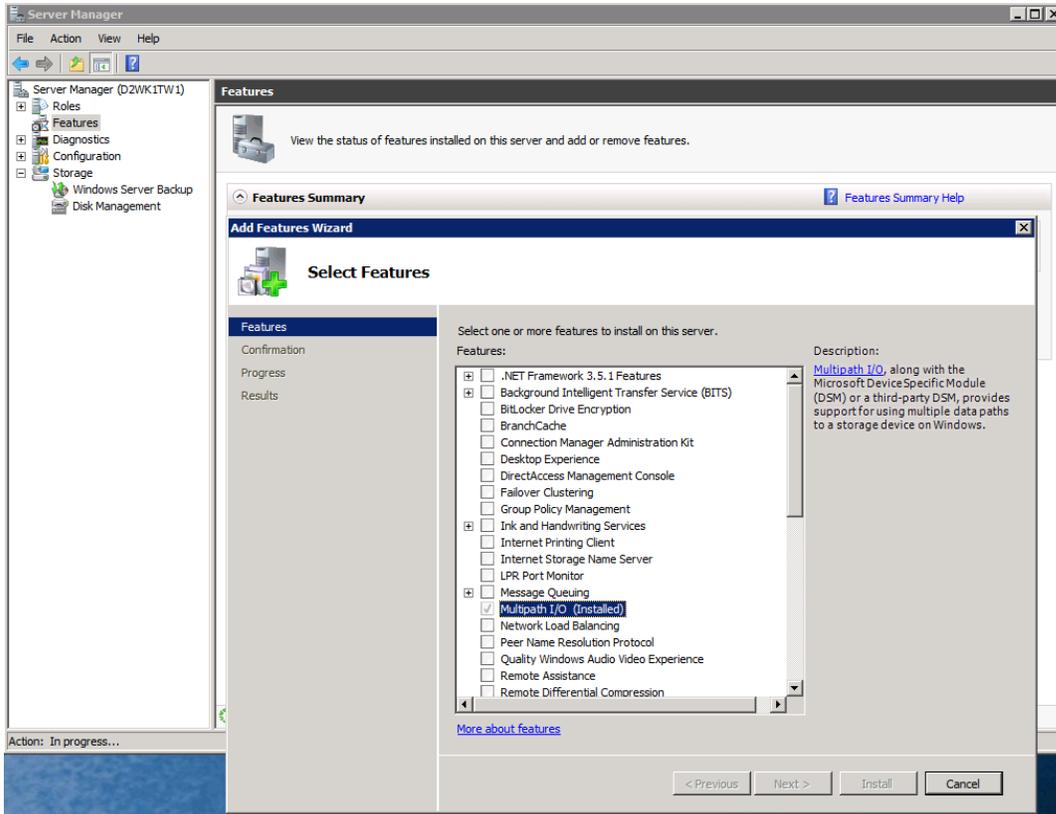
Figure 32   Installing Windows Server 2008 R2 Enterprise Multipath I/O feature

8.   Navigate to **Start > Control Panel > MPIO** and click the **Add** button. When prompted for a **Device Hardware ID**, input **COMPELNTCompellent Vol** and click the **OK** button. The system will need to be restarted for the changes to take effect. Figure 33 displays the **COMPELNTCompellent Vol** text that you should see on the **MPIO Devices** tab in **MPIO Properties** once the system is brought back up.
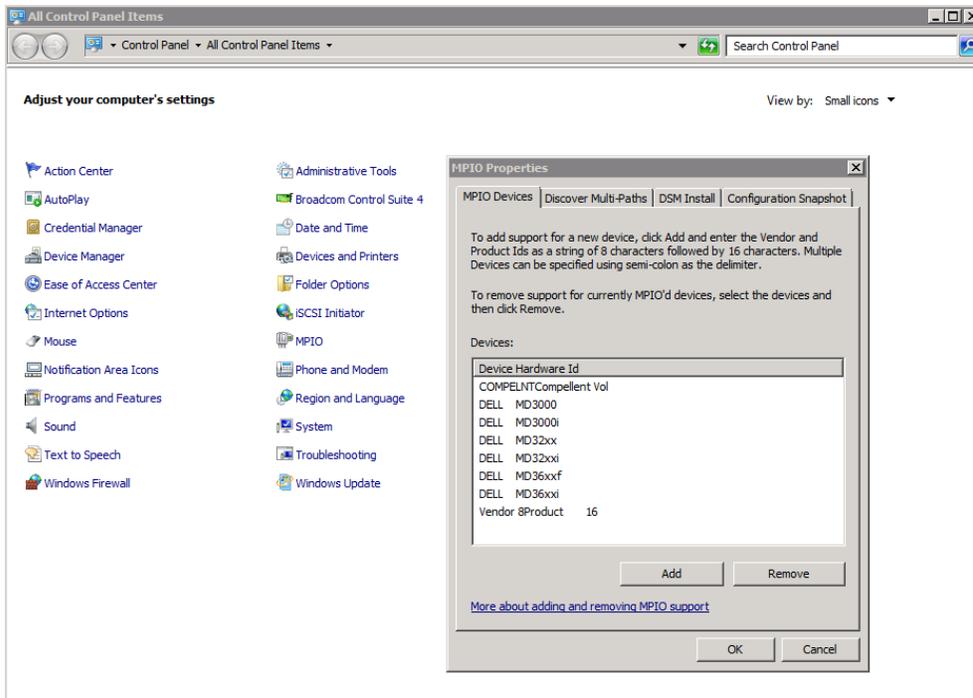
Figure 33   Installing Windows Server 2008 R2 Enterprise Multipath I/O for Compellent array

Next, create a volume and map it to a server object so the respective server can write to the FC storage array.

9.   Simply right click **Volumes** on the left-hand pane and select **Çreate Volume** to get started. During the process, you will be asked to select a **Replay Profile**; this is simply asking you how often **snapshots/recovery points** of the storage volume should be taken. A **snapshot/recovery point** allows you to revert a volume back to a certain point in time (for example if files are accidentally deleted). In Figure 34 below, you can see that a 20 GB volume named **Finance_Data_Compellent** has already been created. Figure 35 displays the dialog box where you can select a **Replay Profile**.

Figure 34   Created 20 GB "Finance_Data_Compellent" volume on Compellent array
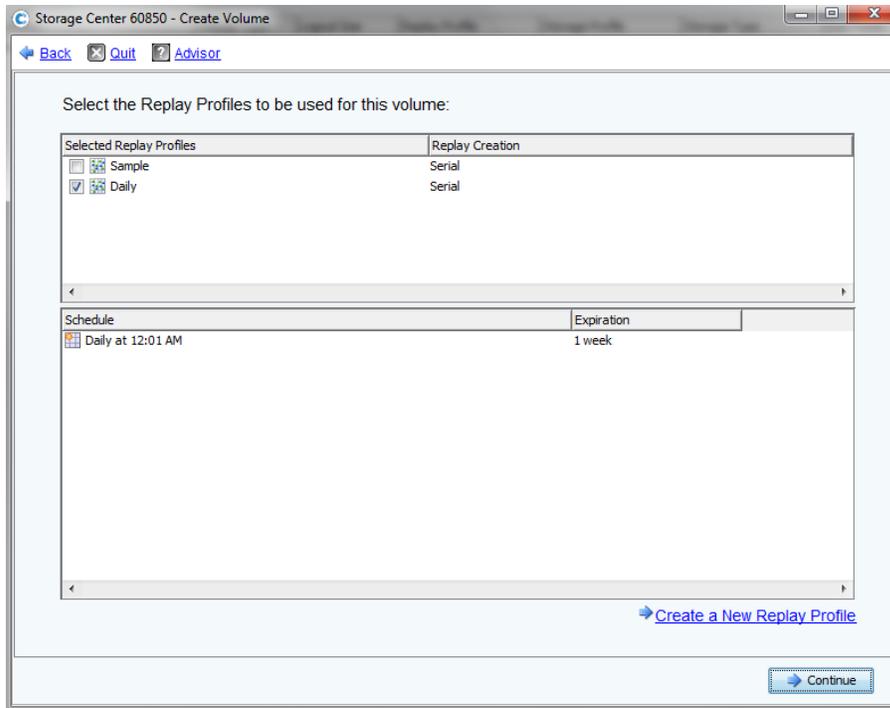
Figure 35   Confirming to keep the default value for 'Replay Profiles'

10. The last step in configuring the Dell Compellent Storage Center array is mapping the newly created volume to the server. Once you create the volume, you will be asked if you want to map it to a server object. You can do it at this time or later. If mapping the volume to a server object later, on the left-hand pane under **Storage > Volumes**, simply right click on the volume you just created and select **Map Volume to Server**. You can then select the respective server object that you created prior.

11. As soon as the FCoE adapter on the Windows server detects storage available for it, it will be detected in the Windows disk management administration tool after performing a disk scan. To perform a disk scan, right click **Disk Management** on the left-hand pane and select **Rescan Disks**. You must right click the detected virtual disk and initialize it. Below in Figure 36, you can see the disk (Disk 1) has already been initialized and formatted as **NTFS**.
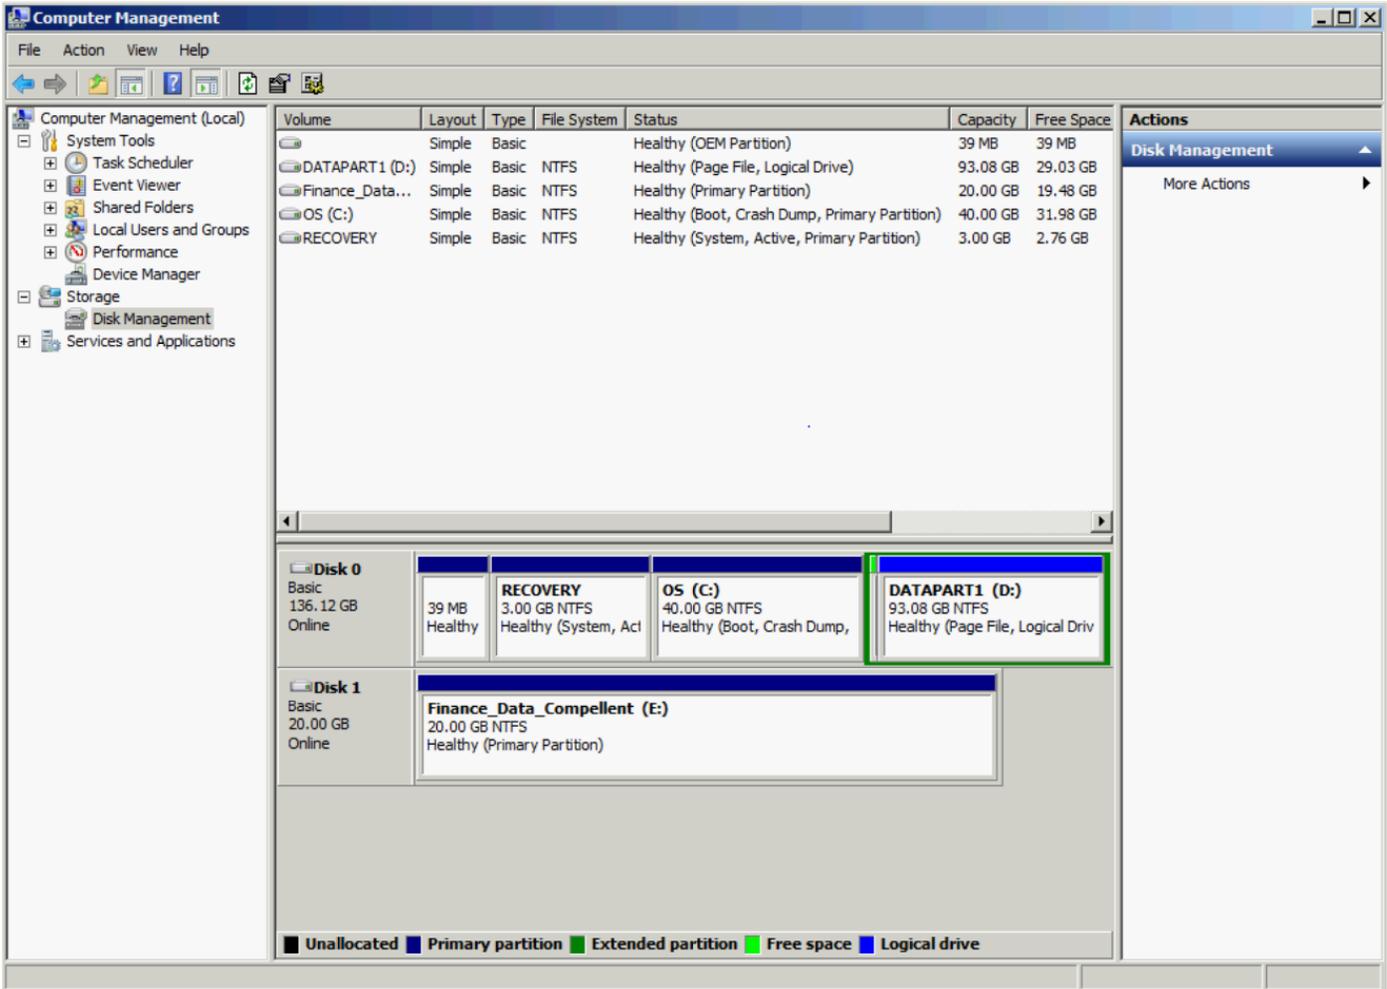
Figure 36   Initialized and formatted virtual disk within Windows Server 2008 R2 Enterprise

Note: The volume on the Compellent storage array displays in Windows just like a typical hard drive. Other than enabling FCoE on the CNA, no special configuration was needed.
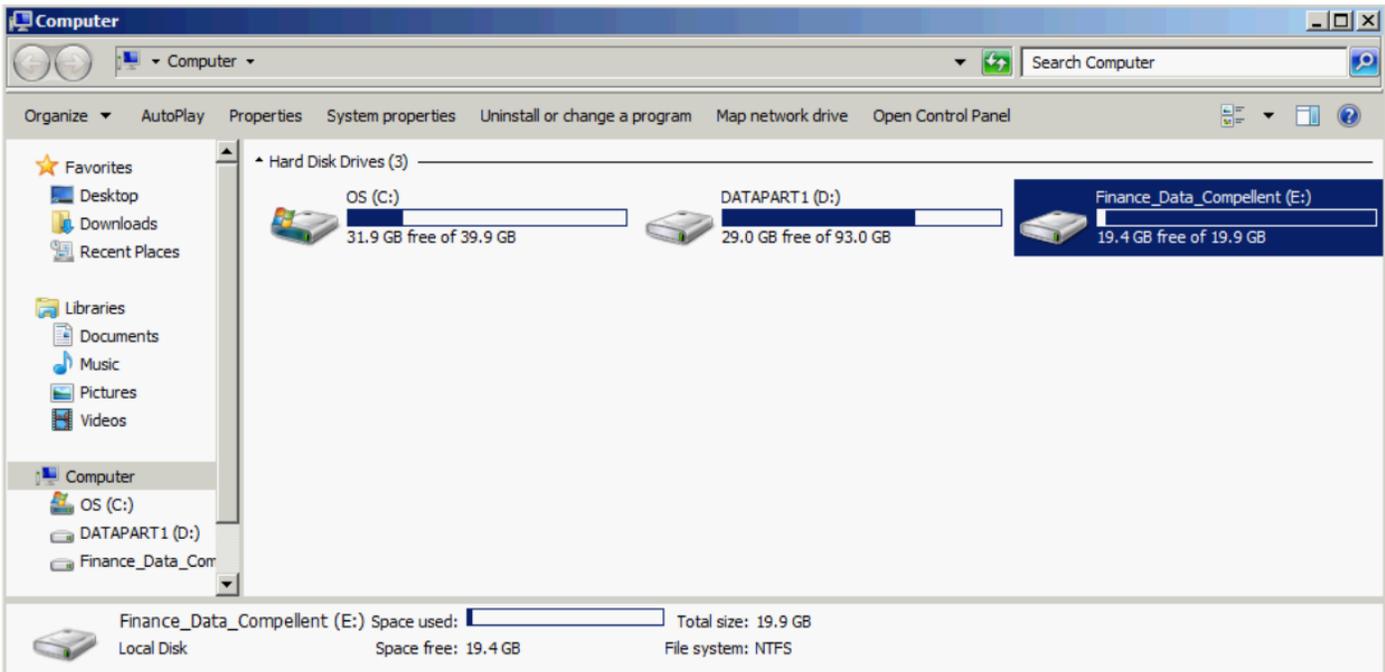
Figure 37   Remote storage on Compellent as seen in Windows as drive 'E:'

Compellent SC8000 Load Balancing Policy Options:

The Compellent SC8000 controller uses Microsoft Multipath I/O (MPIO) for load balancing over ports. Microsoft MPIO is a framework that allows administrators to configure load balancing and failover processes for FC and iSCSI connected storage devices. You can configure load balancing to use up to 32 independent paths from the connected storage devices. The MPIO framework uses Device Specific Modules (DSM) to allow path configuration. For Windows Server 2008 and above, Microsoft provides a built-in generic Microsoft DSM (MSDSM) and it should be used. For Windows Server 2003 only, Dell Compellent provides a DSM.

A load balance policy is used to determine which path is used to process I/O. Once the Compellent volume has been created and mapped accordingly, as will be demonstrated shortly, to see the selected MPIO policy in Windows Server 2008 R2 Enterprise navigate to **Start > Administrative Tools > Computer Management**. On the left-hand pane navigate to **Computer Management > Storage > Disk Management** and right click the disk created on the Compellent storage array and select **Properties**. Next, select the **Hardware** tab, click the **Properties** button at the bottom right, and select the **MPIO** tab. Figure 38 below displays what you should see. Note, for Windows Server 2008, the default MPIO Policy is **Round Robin**.

The Dell Compellent SC8000 will only support **Round Robin** or **Fail Over Only** MPIO policies. The other

policies require Asymmetric Logical Unit Access (ALUA) support on the array and additional support through a DSM.
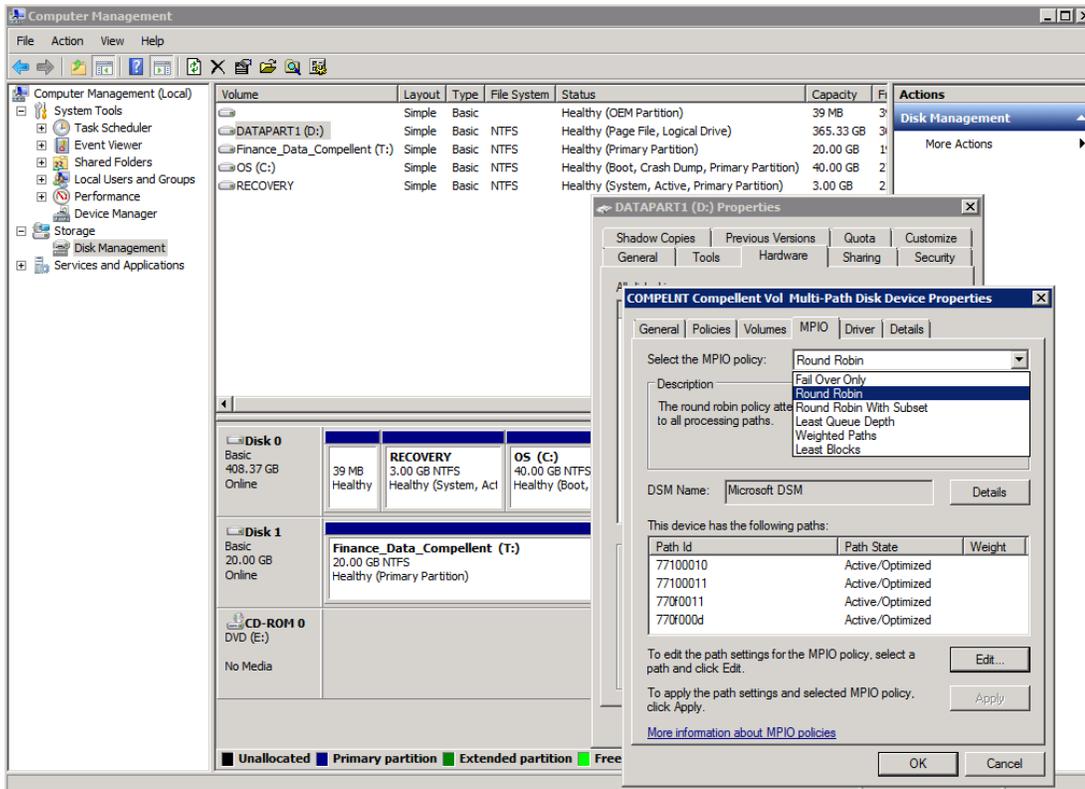
Figure 38   Checking MPIO settings in Windows Server 2008 R2 Enterprise

Additionally, there are two IO connection options available with the Dell Compellent Storage Center that allow multiple paths to be presented to the servers: **Legacy Ports** and **Virtual Ports**. You will be asked which one you would like to use when initially setting-up the Compellent Storage Center and configuring the FC IO cards. See the "Storage Center 6.2 System Setup Guide" for more information on initial setup of the Dell Compellent Storage Center.

In legacy mode, front-end IO ports (in this case FC ports) are broken into primary and reserve ports based on a fault domain. The reserve port is in a standby mode until a primary port fails over to the reserve port. In terms of MPIO, this requires twice the IO ports to enable multiple paths. For redundancy, a primary port connects to one controller, and the reserved port in that fault domain connects to the other controller. While this is a highly robust failover solution, it requires a large number of ports.

Dell Compellent introduced virtual ports in Storage Center 5.0. Virtual ports allow all front-end IO ports to be virtualized. All FC ports can be used at the same time for load balancing as well as failover to

another port. Although a virtual disk can still only be written to from the controller that owns the disk, virtual ports allow for better performance in terms of failover as the virtual connection can simply be moved to another physical port in the same fault domain. To use virtual ports, all FC switches and HBAs must support N_Port ID Virtualization (NPIV). See the "Dell Compellent Storage Center Microsoft Multipath IO (MPIO) Best Practices Guide" for more information on multipathing with Microsoft Windows Server 2008 R2 Enterprise.

4.4 Brocade 6505 FC Switch Configuration

Note, the Brocade 6505 FC switch has not been configured yet and connectivity from the server to the Compellent Storage Array is good. The reason for this is because, by default, the Brocade FC switch allows all devices to see and communicate with each other. However, this is not best practice.

Access control on a switched SAN fabric can be provided by the fabric itself via zoning. By creating zones, the fabric can be partitioned so that only specific devices can see and talk to each other. This is somewhat analogous to VLANs on a LAN. However, unlike VLANs, zones typically consist of very few devices such as one initiator and a few targets.

Best practice for zoning is to create a separate zone for each path that the host can see. This minimizes external chatter and when something changes all nodes won't be notified of the change. If there are few nodes connected and to minimize configuration and maintenance, some may prefer to group multiple WWPNs going to the same destination within one zone. Below, one zone is created for the FCoE port going to all four ports of the storage. Note the World Wide Port Number (WWPN) of the end node ports are used in the zoning configuration. We also could have created four zones – one per FCoE port and storage port combination.

The two FC switches being used are Brocade 6505s and the zoning configurations are below. The WWPNs starting with **20** are the virtual FCoE port WWPNs and the other WWPNs are for the Compellent storage array. In this example, the WWPN of the CNA virtual FCoE port is being utilized. When working with network adaptors that provide a MAC address as well as a World Wide Port and Node name, it's important to understand where these addresses originate. The M1000e chassis has a feature called Flexaddressing which allows for a virtual MAC address/WWPN to be linked to a server, therefore, if the CNA/adapter is changed later, zoning or configuration on switches does not need to be modified. For more information on this feature and how to implement it see Section 6 "M1000e FlexAddress enablement" of this document.

1. Fabric A Brocade FC Switch Zoning Configuration

```
> zonecreate financeServer1_p1_test, "20:01:90:b1:1c:cc:e2:6a; 50:00:d3:10:00:ed:b2:3b;
50:00:d3:10:00:ed:b2:3d; 50:00:d3:10:00:ed:b2:41; 50:00:d3:10:00:ed:b2:43"

> cfgcreate zoneCfg_test,"financeServer1_p1_test"

> cfgenable zoneCfg_test

> cfgsave
```

On the fabric A FC switch you can see the WWPN of the server FCoE port is **20:01:90:b1:1c:cc:e2:6a** and the WWPNs of the storage ports are **50:00:d3:10:00:ed:b2:3b**, **50:00:d3:10:00:ed:b2:3d**, **50:00:d3:10:00:ed:b2:41**, and **50:00:d3:10:00:ed:b2:43**. This zoning configuration is allowing all four storage ports to communicate only to each other and the server virtual FCoE port.

The zoning configuration can be confirmed with the **zoneshow** command.

```
Brocade_6505_1:admin> zoneshow
Defined configuration:
 cfg:    zoneCfg_test
                financeServer1_p1_test
 zone:   financeServer1_p1_test
                20:01:90:b1:1c:cc:e2:6a; 50:00:d3:10:00:ed:b2:3d;
                50:00:d3:10:00:ed:b2:43; 50:00:d3:10:00:ed:b2:3b;
                50:00:d3:10:00:ed:b2:41

Effective configuration:
 cfg:    zoneCfg_test
 zone:   financeServer1_p1_test
                20:01:90:b1:1c:cc:e2:6a
                50:00:d3:10:00:ed:b2:3d
                50:00:d3:10:00:ed:b2:43
                50:00:d3:10:00:ed:b2:3b
                50:00:d3:10:00:ed:b2:41

Brocade_6505_1:admin>
```

Figure 39   Displaying zoning configuration on fabric A Brocade FC switch

2. Fabric B Brocade FC Switch Zoning Configuration

```
> zonecreate financeServer1_p2_test, "20:01:90:b1:1c:cc:e5:39; 50:00:d3:10:00:ed:b2:3a;
50:00:d3:10:00:ed:b2:3c; 50:00:d3:10:00:ed:b2:40; 50:00:d3:10:00:ed:b2:42"

> cfgcreate zoneCfg_test,"financeServer1_p2_test"

> cfgenable zoneCfg_test

> cfgsave
```

On the fabric B FC switch you can see the WWPN of the server FCoE port is **20:01:90:B1:1C:CC:E5:39** and the WWPNs of the storage ports are **50:00:d3:10:00:ed:b2:3a**, **50:00:d3:10:00:ed:b2:3c**, **50:00:d3:10:00:ed:b2:40**, and **50:00:d3:10:00:ed:b2:42**.

The zoning configuration can be confirmed with the **zoneshow** command.

```
Brocade_6505_2:admin> zoneshow
Defined configuration:
 cfg:    zoneCfg_test
                financeServer1_p2_test
 zone:   financeServer1_p2_test
                20:01:90:b1:1c:cc:e5:39; 50:00:d3:10:00:ed:b2:3a;
                50:00:d3:10:00:ed:b2:3c; 50:00:d3:10:00:ed:b2:40;
                50:00:d3:10:00:ed:b2:42

Effective configuration:
 cfg:    zoneCfg_test
 zone:   financeServer1_p2_test
                20:01:90:b1:1c:cc:e5:39
                50:00:d3:10:00:ed:b2:3a
                50:00:d3:10:00:ed:b2:3c
                50:00:d3:10:00:ed:b2:40
                50:00:d3:10:00:ed:b2:42

Brocade_6505_2:admin>
```

Figure 40   Displaying zoning configuration on fabric B Brocade FC switch

To observe that the storage ports and FCoE ports are logged into the fabric, you can use the **nsshow** command on the Brocade FC switch. A small capture of the output of this command is shown below.

```
Brocade_6505_1:admin> nsshow
{
 Type Pid    COS     PortName                     NodeName                   TTL(sec)
 N    010300;        3;50:00:d3:10:00:ed:b2:08;50:00:d3:10:00:ed:b2:00; na
     FC4s: FCP
     PortSymb: [90] "Compellent Port QLGC FC 8Gbps; Slot=05 Port=04 in Controller: SN 60850 of Storage Center: "
     NodeSymb: [47] "Compellent Storage Center: Storage Center 60850"
     Fabric Port Name: 20:03:00:05:33:d3:f2:4f
     Permanent Port Name: 50:00:d3:10:00:ed:b2:08
     Port Index: 3
     Share Area: No
     Device Shared in Other AD: No
     Redirect: No
     Partial: No
 N    010301;        3;50:00:d3:10:00:ed:b2:3d;50:00:d3:10:00:ed:b2:01; na
     FC4s: FCP
     PortSymb: [110] "Compellent Port QLGC FC 8Gbps; Slot=05 Port=04 in Controller: SN 60850 of Storage Center: Storage Center 60850"
     NodeSymb: [47] "Compellent Storage Center: Storage Center 60850"
     Fabric Port Name: 20:03:00:05:33:d3:f2:4f
     Permanent Port Name: 50:00:d3:10:00:ed:b2:08
     Port Index: 3
     Share Area: No
     Device Shared in Other AD: No
     Redirect: No
     Partial: No
```

Figure 41   'nsshow' command output on fabric A Brocade FC switch

You can also see the node WWPN by looking at what is logged in on the physical port as shown in Figure 42 below.

```
Brocade_6505_1:admin> portshow 15
portIndex:  15
portName: port15
portHealth: HEALTHY

Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x24b03       PRESENT ACTIVE F_PORT G_PORT U_PORT NPIV LOGICAL_ONLINE LOGIN NOELP LED ACCEPT FLOGI
LocalSwcFlags: 0x0
portType:  24.0
POD Port: Port is licensed
portState: 1    Online
Protocol: FC
portPhys: 6     In_Sync        portScn:   32    F_Port
port generation number:    41572
state transition count:    25


portId:    010f00
portIfId:    4302000f
portWwn:    20:0f:00:05:33:d3:f2:4f
portWwn of device(s) connected:
       20:01:90:b1:1c:cc:e2:6a
       20:2c:d0:67:e5:a7:b0:80
Distance:  normal
portSpeed:  N8Gbps

FEC: Inactive
LE domain: 0
FC Fastwrite: OFF
Interrupts:        0         Link_failure: 5         Frjt:        0
Unknown:          55         Loss_of_sync: 11        Fbsy:        0
Lli:             214         Loss_of_sig:  12
Proc_rqrd:      1925         Protocol_err: 0
Timed_out:         0         Invalid_word: 22317
Rx_flushed:        0         Invalid_crc:  1281
Tx_unavail:        0         Delim_err:    498
Free_buffer:       0         Address_err:  0
Overrun:           0         Lr_in:        13
Suspended:         0         Lr_out:       3
Parity_err:        0         Ols_in:       0
2_parity_err:      0         Ols_out:      13
CMI_bus_err:       0

Port part of other ADs: No
Brocade_6505_1:admin>
```

Figure 42   'portshow 15' command output on fabric A Brocade FC switch

Another useful FC switch command to check what ports are connected is **switchshow**.

```
Brocade_6505_1:admin> switchshow
switchName:       Brocade_6505_1
switchType:       118.1
switchState:      Online
switchMode:       Native
switchRole:       Principal
switchDomain:     1
switchId:         fffc01
switchWwn:        10:00:00:05:33:d3:f2:4f
zoning:           ON (zoneCfg_test)
switchBeacon:     OFF
FC Router:        OFF
FC Router BB Fabric ID: 1
Address Mode:     0

Index Port Address Media Speed State      Proto
==================================================
    0    0  010000   id   N16  No_Light    FC  Disabled
    1    1  010100   id   N16  No_Light    FC  Disabled
    2    2  010200   id   N16  No_Light    FC  Disabled
    3    3  010300   id    N8  Online      FC  F-Port  1 N Port + 1 NPIV public
    4    4  010400   id   N16  In_Sync     FC  Disabled (Remote domain is RCS incapable)
    5    5  010500   --   N16  No_Module   FC
    6    6  010600   --   N16  No_Module   FC
    7    7  010700   --   N16  No_Module   FC
    8    8  010800   id    N8  Online      FC  F-Port  1 N Port + 1 NPIV public
    9    9  010900   id    N8  Online      FC  F-Port  1 N Port + 1 NPIV public
   10   10  010a00   id    N8  Online      FC  F-Port  1 N Port + 1 NPIV public
   11   11  010b00   id    N8  No_Light    FC
   12   12  010c00   id    N8  Online      FC  F-Port  20:29:d0:67:e5:a7:b0:7d
   13   13  010d00   id    N8  Online      FC  F-Port  20:2a:d0:67:e5:a7:b0:7e
   14   14  010e00   id    N8  Online      FC  F-Port  1 N Port + 1 NPIV public
   15   15  010f00   id    N8  Online      FC  F-Port  1 N Port + 1 NPIV public
   16   16  011000   --   N16  No_Module   FC
   17   17  011100   --   N16  No_Module   FC
   18   18  011200   --   N16  No_Module   FC
   19   19  011300   id   N16  No_Light    FC
   20   20  011400   --   N16  No_Module   FC
   21   21  011500   --   N16  No_Module   FC
   22   22  011600   --   N16  No_Module   FC
   23   23  011700   --   N16  No_Module   FC
Brocade_6505_1:admin>
```

Figure 43   'switchshow' command output on fabric A Brocade FC switch

4.5 Verification on the Dell MXL w/ FC Flex IOM

1. To see information on NPIV devices logged into the fabric, use the **show npiv devices** command on the Dell MXL switch as shown below. Note the FCoE MAC is **0e:fc:02:01:0f:01** (the FCoE Map + FC_ID as expected).

```
Dell_MXL_A1#show npiv devices

ENode[0]:
ENode MAC     :    90:b1:1c:cc:e2:6a
ENode Intf    :    Te 0/4
FCF MAC       :    d0:67:e5:a7:b0:80
Fabric Intf   :    Fc 0/44
FCOE Vlan     :    1002
Fabric Map    :    SAN_FABRIC_A
ENode WWPN    :    20:01:90:b1:1c:cc:e2:6a
ENode WWNN    :    20:00:90:b1:1c:cc:e2:6a
FCoE MAC      :    0e:fc:02:01:0f:01
FC-ID         :    01:0f:01
LoginMethod   :    FLOGI
Secs          :    9438
Status        :    LOGGED_IN
Dell_MXL_A1#
```

Figure 44   'show npiv devices' command output on fabric A Dell MXL switch

2. To see currently active FCoE VN_Port sessions, use the **show fip-snooping sessions** command.

```
Dell_MXL_A1#show fip-snooping sessions
Enode MAC          Enode Intf    FCF MAC           FCF Intf    VLAN    FCoE MAC          FC-ID      Port WWPN                 Port WWNN
---------          ----------    --------          --------    ----    --------          -----      ---------                 ---------
90:b1:1c:cc:e2:6a  Te 0/4        d0:67:e5:a7:b0:80 Fc 0/44     1002    0e:fc:02:01:0f:01 01:0f:01   20:01:90:b1:1c:cc:e2:6a   20:00:90:b1:1c:cc:e2:6a
Dell_MXL_A1#
```

Figure 45   'show fip-snooping enode' command output on fabric A Dell MXL switch

3. To see all FCoE end-node (ENodes) , use the **show fip-snooping enode** command.

```
Dell_MXL_A1#show fip-snooping enode
Enode MAC            Enode Interface       FCF MAC             VLAN        FC-ID
---------            ---------------       -------             ----        -----
90:b1:1c:cc:e2:6a    Te 0/4                d0:67:e5:a7:b0:80   1002        01:0f:01
Dell_MXL_A1#
```

Figure 46   'show fip-snooping enode' command output on fabric A Dell MXL switch

4. To see a list of configured fcoe-maps, use the **show fcoe-map brief** command.

```
Dell_MXL_A1#show fcoe-map brief
    Fabric-Name   Fabric-Id            Vlan-Id      FC-MAP    FCF-Priority  Config-State      Oper-State
    SAN_FABRIC_A      1002                1002      0efc02          128        ACTIVE              UP
Dell_MXL_A1#
```

Figure 47   'show fcoe-map brief' command output on fabric A Dell MXL switch

5. To see more detailed information on a given fcoe-map, use the **show fcoe-map
   <FCoE_MAP_NAME>** command. Notice below, the priority mapped to FCoE by default is **3**.

```
Dell_MXL_A1#show fcoe-map SAN_FABRIC_A

Fabric Name          SAN_FABRIC_A
Fabric Id            1002
Vlan Id              1002
Vlan priority        3
FC-MAP               0efc02
FKA-ADV-Period       8
Fcf Priority         128
Config-State         ACTIVE
Oper-State           UP
Members
Fc 0/41 Fc 0/42 Fc 0/43 Fc 0/44
Te 0/9 Te 0/4
Dell_MXL_A1#
```

Figure 48   'show fcoe-map SAN_FABRIC_A' command output on fabric A MXL switch

5 Dell IOA w/ FC Flex IOM Configuration

This example uses a similar topology as shown prior except instead of using the Dell MXLs in the Fabric A blade slot, the Dell IOAs with FC Flex IOMs in the Fabric B slot of the M1000e chassis are used. Another difference from the prior example in Section 4 is that instead of the Broadcom 57810 NDC CNA in the Fabric A slot of the server, the QLogic QME8262-k mezzanine CNA adapter inserted in the Fabric B slot of the server is used.

5.1 Dell QLogic QLE8262 CNA Configuration

The below shows an example of how to configure the QLogic QME8262-k mezzanine CNA adapter from the **QConvergeConsole** CLI management application.

QLogic offers CNAs in three formats for Dell 12G servers: QLE8262 standard PCI Express, QME8262-k mezzanine for Dell blade servers, and QMD8262-k for the Dell Network Daughter Card. The Dell QLogic QME8262-k mezzanine CNA allows for up to four partitions per physical port and eight partitions total per 2-port adapter. A partition can be looked upon as a virtual port.

This example displays Dell QLogic QME8262-k mezzanine CNA configuration in Microsoft Windows Server 2008 R2 Enterprise. By default, only the NIC functionality is enabled. FCoE must be manually enabled on the CNA for the virtual FCoE ports to be identified in Windows. The configuration of the CNA for FCoE is shown in Figures 48 - 51.

1. Once the Dell QLogic QME8262-k mezzanine CNA drivers and **QConvergeConsole** CLI are installed, double click the **QConvergeConsole** CLI shortcut in Windows and configure the CNA as shown below in Figure 51.

   You can see that **function 6** on **port 1** and **function 7** on **port 2** have been configured to handle FCoE, **function 0** on **port 1** and **function 1** on **port 2** have been configured to handle LAN traffic, and the other partitions have been disabled.

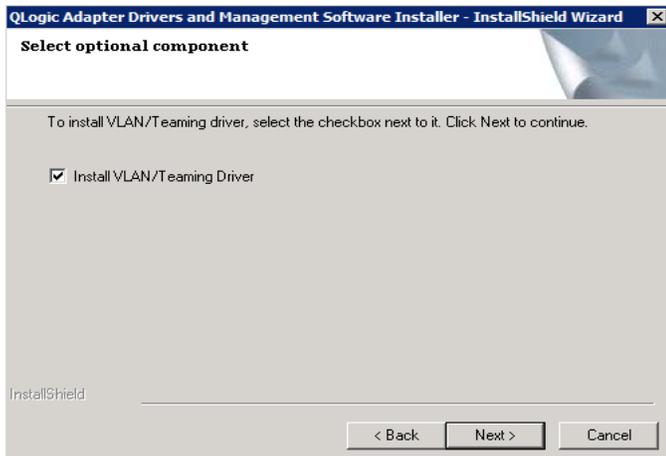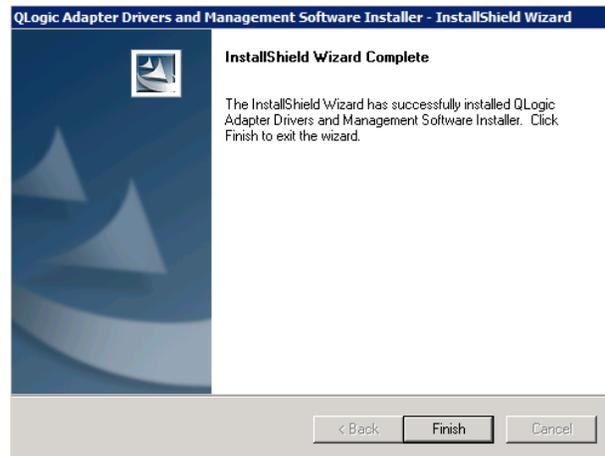Figure 49   QLogic QLE8262 CNA Configuration

```
QConvergeConsole

CLI - Version 1.1.0 (Build 9)

Main Menu

 1:  Adapter Information
 2:  Adapter Configuration
 3:  Adapter Updates
 4:  Adapter Diagnostics
 5:  Adapter Statistics
 6:  NIC Partitioning (NPAR) Information
 7:  NIC Partitioning (NPAR) Configuration
 8:  Refresh
 9:  Help
10:  Exit
```

Figure 50   Dell Qlogic QLE8262 CNA QConvergeConsole screen

```
      Please Enter Selection: 7

      QConvergeConsole

      CLI - Version 1.1.0 (Build 9)

  NIC Partitioning (NPAR) Configuration Selection

  1:   NPAR Configuration


       (p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit)
       Please Enter Selection: 1

       QConvergeConsole

       CLI - Version 1.1.0 (Build 9)

  NIC Partitioning (NPAR) Configuration

  1:   Change PCI Function Personality


       (p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit)
       Please Enter Selection: 1

       QConvergeConsole

       CLI - Version 1.1.0 (Build 9)

  NIC Partitioning (NPAR) Port Configuration Selection

  CNA Model QME8262 SN: RFE1314H60191
    1. Port    1
        Function    6 [Type: FCoE ] [WWPN: 20-01-90-B1-1C-CC-E5-76] Online
        Function    0 [Type: NIC  ] [MAC Address: 90:B1:1C:CC:E2:6C] Link Up
    2. Port    2
        Function    7 [Type: FCoE ] [WWPN: 20-01-90-B1-1C-CC-E5-77] Online
        Function    1 [Type: NIC  ] [MAC Address: 90:B1:1C:CC:E2:6D] Link Up


       (p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit)
       Please Enter Selection: _
```

Figure 51   Dell Qlogic QLE8262 CNA FcoE/NPAR Configuration

Creating a NIC Team

Since the NICs and HBAs are seen as virtual ports, you can treat them as separate entities and create a NIC team with the virtual CNA NIC ports. In Figures 51 - 54, you can see the two virtual NIC ports are NIC teamed using 'Failsafe Team'.

In this example, Windows Server 2008 R2 Enterprise is the operating system used.

2. To create a NIC team on the virtual NIC ports navigate to **Control Pane > Network and Internet > Network and Sharing Center > Change adapter settings** and right click one of the ports you wish to put in a NIC team. Click **Properties**. Click the **Configure** button. Next, click the **Team Management** tab as shown in Figure 52.
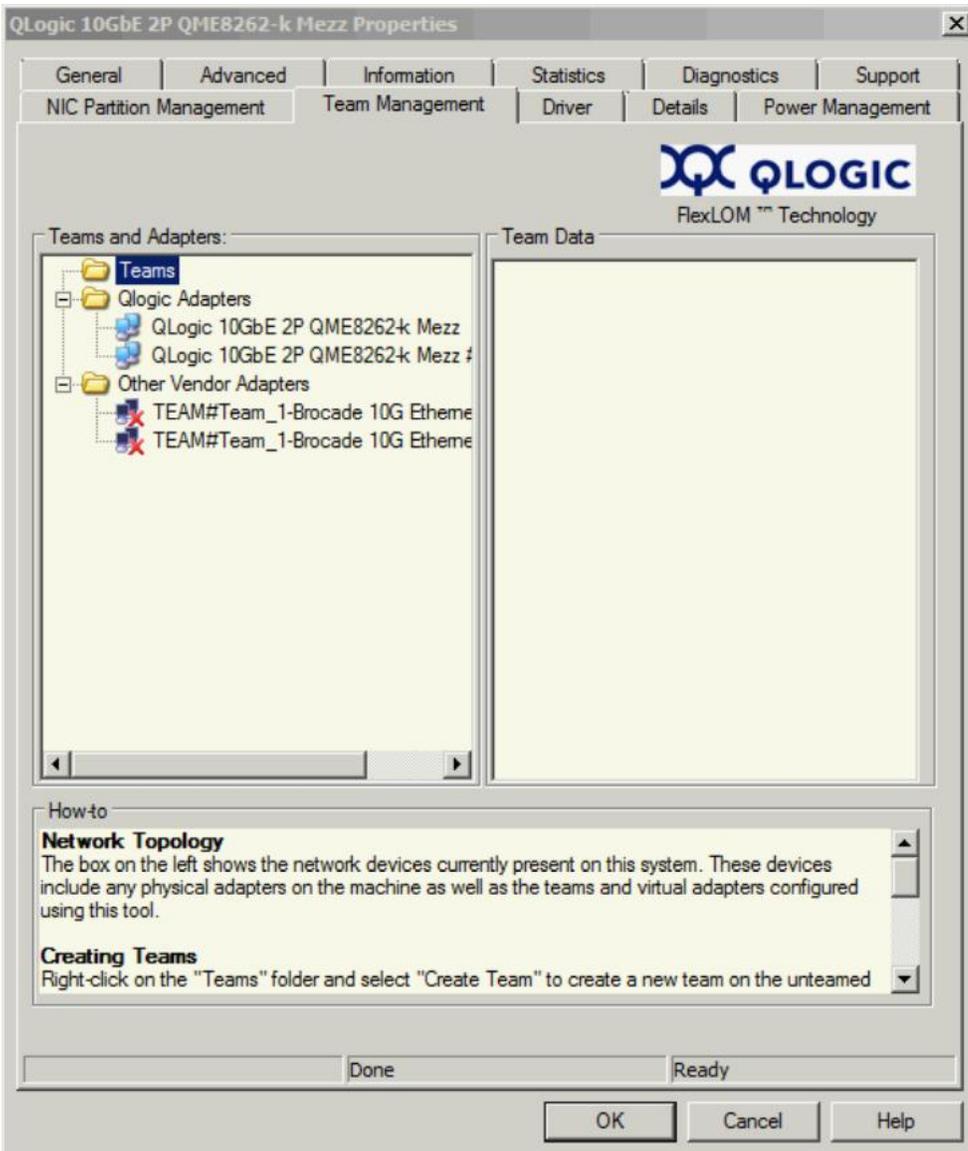
Figure 52   'Team Management' tab of port properties

3. Now right click on the **Teams** folder and click **Create Team**. Choose the type of NIC teaming you desire. In
    this example we will demonstrate with **Failsafe Team**. Next, select the ports to add to the NIC team and select
    the primary adapter. The rest of the settings we leave as default. Figure 53 displays the virtual port NIC team
    with two virtual NIC ports as members.
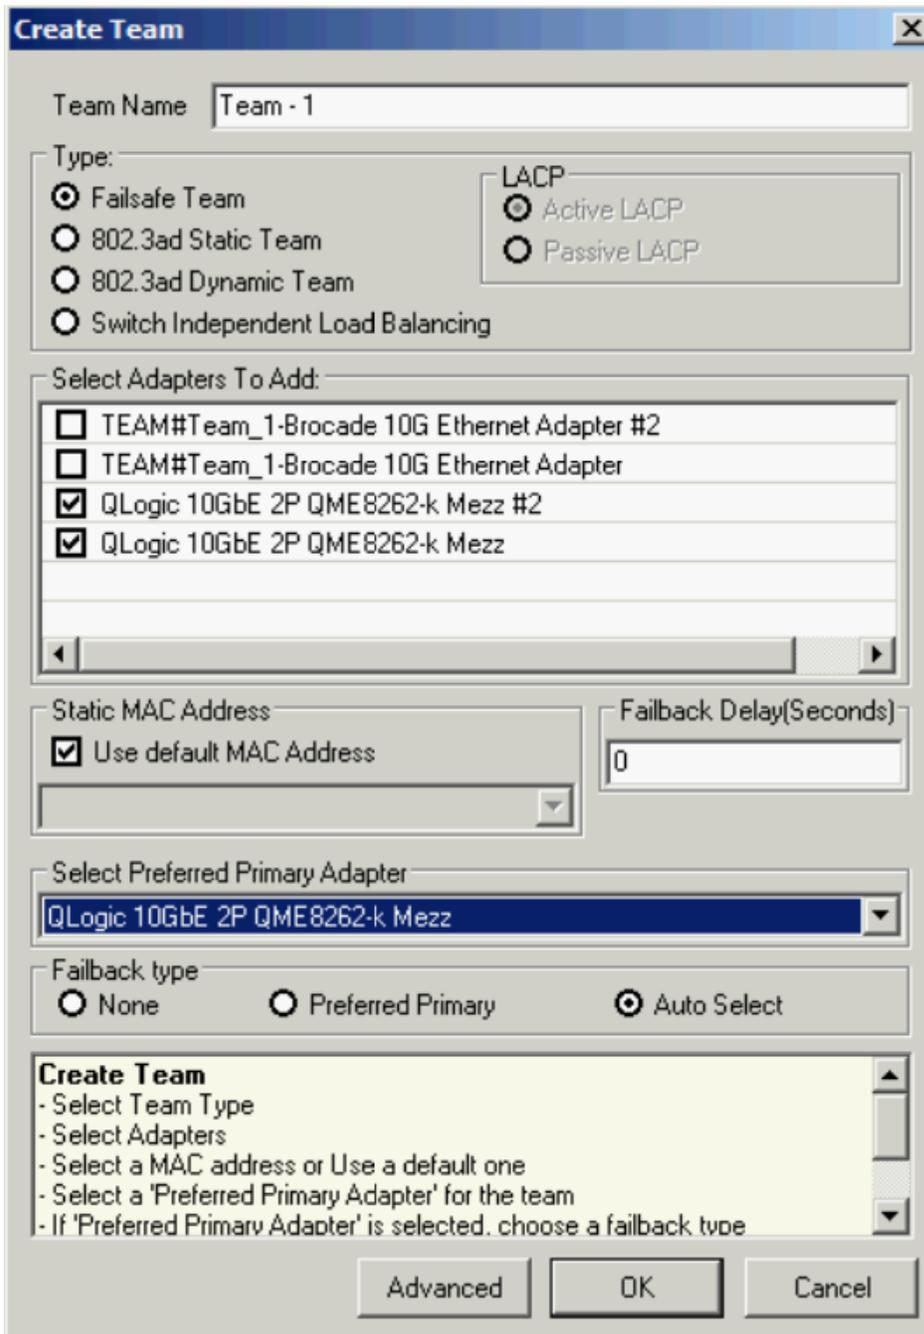
Figure 53   NIC teaming virtual NIC ports with Failsafe Team

Dell Networking FC Flex IOM: Deployment of FCoE with Dell FC Flex IOM, Brocade FC switches,
and Dell Compellent Storage Array
Deployment/Configuration Guide
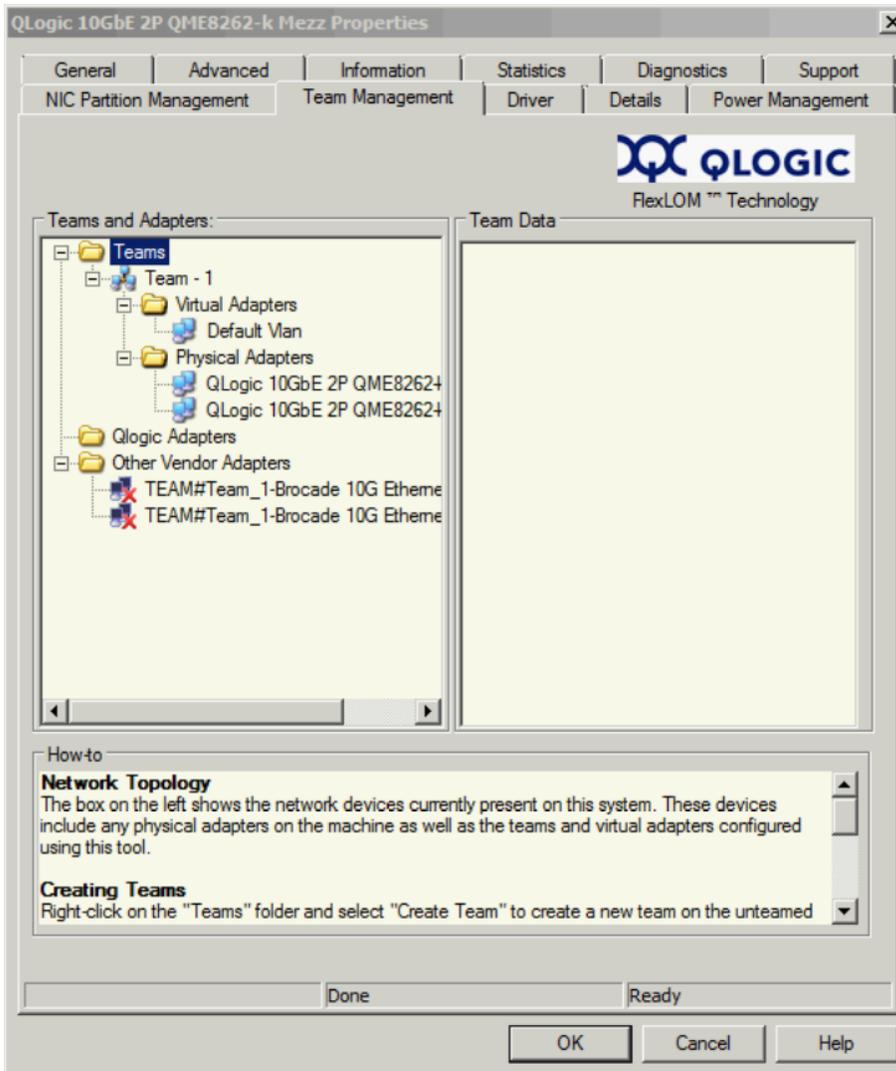
Figure 54   Dell QLogic QLE8262 adapter propertise displaying the created NIC team

As far as the network configuration for the LAN, since **Failsafe Team** is utilized, there is no special configuration that needs to be done on the IOA switches. You can simply have one link going to each IOA switch with one port in **active** mode and the other in **standby** mode.

4.   In this example, the LAN traffic is on VLAN 5. You can easily tag the NIC team with VLAN 5 by right clicking the VLAN name, selecting **Modify VLAN**, and entering the respective VLAN as shown below.
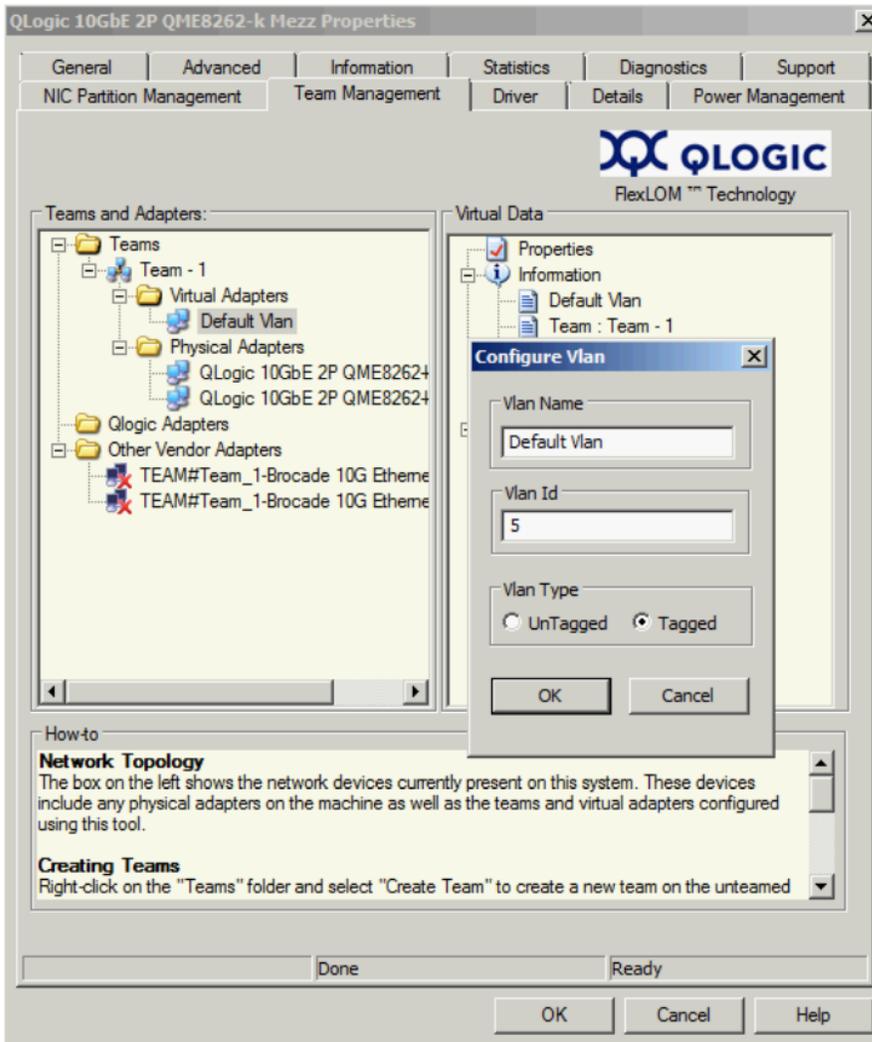
Figure 55   Tagging the NIC team with VLAN 5

The NIC team will now show in Windows as a new virtual adapter as shown in Figure 56 and Figure 57.

Figure 56   Virtual adapter network connection as seen in Windows

Network adapters
  [Default Vlan][Team - 1]:QLogic Teaming Virtual Adapter

Figure 57   NIC team virtual adapter as seen in 'Device Manager' in Windows

5.2   Dell IOA w/ FC Flex IOM Configuration

5.2.1   Default Configuration

If only layer 2 functionality is required, FC Flex IOM can be used with the Dell IOA blade switch to provide Ethernet-FC bridging capability instead of the Dell MXL blade switch which also provides advanced layer 3 functionality. In addition, in the default **standalone** mode, the Dell IOA blade switch requires zero-touch configuration. Simply insert the blade sw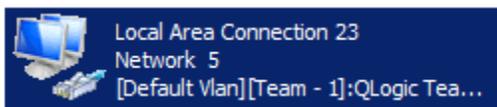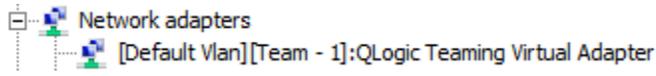itch into the Dell M1000e chassis and it functions as a NPIV Proxy Gateway switch with the default configuration provided. The below topology is slightly different than the topology shown prior with the MXL blade switches. The main difference is that instead of the Dell MXLs in Fabric A, we are now using the Dell IOAs in Fabric B of the M1000e chassis with FC Flex IOMs. Because we're using Fabric B of the M1000e chassis, we are also using the QLogic QME8262-k mezzanine CNA adapter inserted in the Fabric B slot of the server.

Note, by default, all uplink ports on a Dell IOA are part of the Port Channel 128 LAG. All 40 GbE ports are in quad mode (10 GbE mode) and also part of the Port Channel 128 LAG. In such a configuration, the upstream LAN switches should be configured correctly for the Port Channel 128 LAG consisting of 10 GbE ports. In such a configuration, a possible design may be such as the below. Here VLT can still be employed on the S6000s down to the IOA for the LAN traffic. This gives an extra level of switch-level redundancy for the ToR. In this configuration, the LAN connection from the server is active/standby while the SAN connection is active/active and utilizing a **round-robin** multipathing policy.
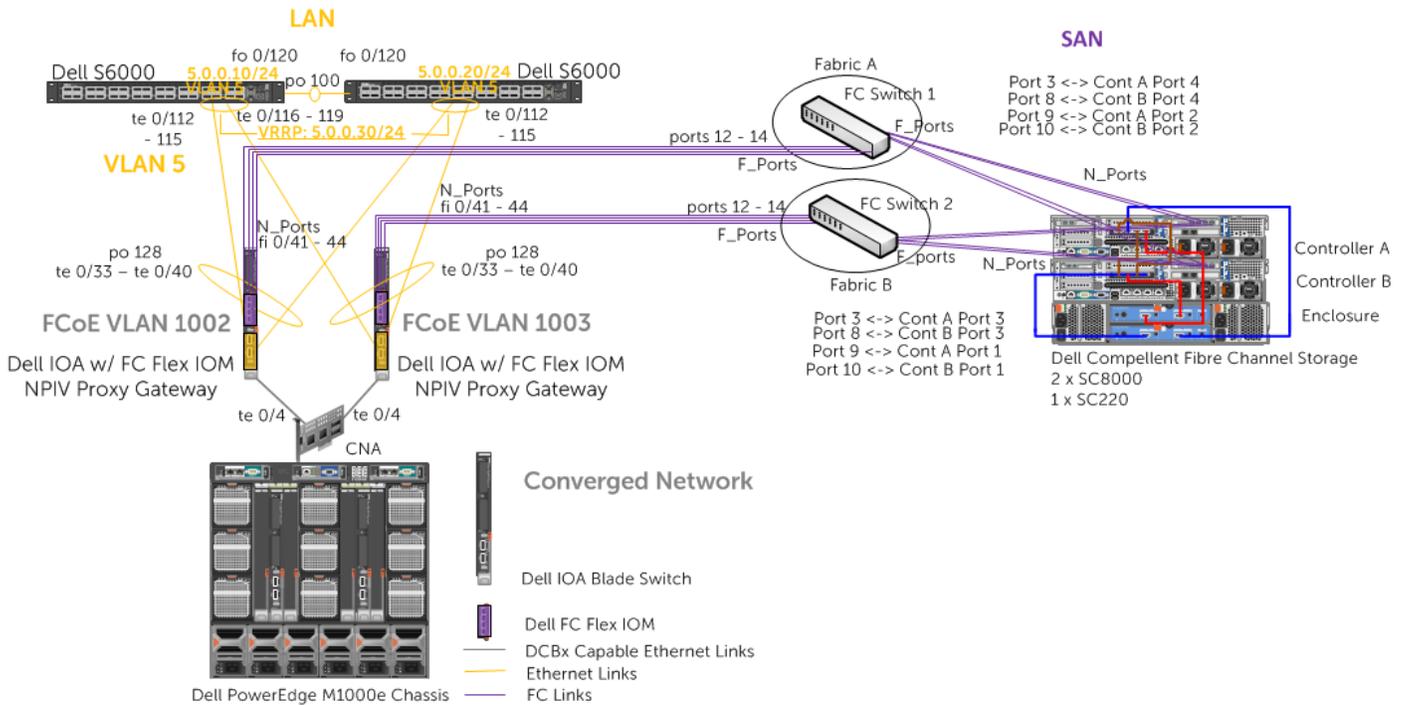
Figure 58   Dell IOA w/ FC Flex IOM acting as a NPIV Proxy Gateway and allowing for a converged
infrastructure

By simply inserting the Dell IOA blade switch w/ FC Flex IOM into the Dell M1000e chassis, you can see below that the end-node is able to log into the SAN A fabric without any configuration needed. In default **standalone** mode, the Dell PowerEdge M I/O Aggregator requires zero-touch configuration in terms of setting up the environment for FCoE. By default it is already preconfigured for NPIV Proxy Gateway when the FC Flex IO Module is detected. Also by default, all uplink ports are lagged together via LACP in **port-channel 128**. In figure 58 above, since we are using the two included 40 GbE ports on each IOA up to the S6000s. Ports **tengigabitethernet 0/33 – tengigabitethernet 0/40** are the ports that are being utilized to connect to the S6000 via 40 GbE cable. In this example 40 GbE active fiber DACs are used but 40 GbE copper DACs could also have been used. OM3 fiber and 8G FC optics are used to connect the FC ports from the FC Flex IOM to the FC ports on the Brocade FC switches.

The IO Aggregator is meant to be used for ease of deployment. All interfaces are tagged on all VLANs and untagged on the default VLAN. For this reason, since all DCB and FCoE configuration is also applied by default, there is no configuration needed on the Dell IOA. If it is desired to change the VLAN settings, the Chassis Management Controller (CMC) GUI can be employed.

It should be noted that the IOA has uplink failure detection by default, so if the upstream connectivity is down, the IOA automatically disables the port connecting to the server – **tengigabitethernet 0/4** in this case. This is done for proper failover upon upstream connectivity failure. Because of all the

preconfiguration already applied by default, once the server is configured properly, the IOA will automatically function as a FCoE NPIV Proxy Gateway. For the Dell, MXL setup in figure 5 prior in this document, we manually applied much of the same configuration such as uplink-failure detection.

By simply inserting the Dell IOA blade switch w/ FC Flex IOM into the Dell M1000e chassis, you can see below that the end-node is able to log into the SAN A fabric without any configuration needed.

```
Dell_IOA_B1#show fip-snooping enode
Enode MAC              Enode Interface       FCF MAC              VLAN       FC-ID
---------              ---------------       -------              ----       -----
90:b1:1c:cc:e5:76     Te 0/4                d0:67:e5:ac:ac:ae    1002       01:0f:01
Dell_IOA_B1#
```

Figure 59   'show fip-snooping enode' command output on fabric A IOA switch

By checking the **running-config** of the Dell IOA switch, the default configuration can be seen. The proper LACP, FCoE, and uplink failure detection configuration is already present by default so no further configuration needs to be done. Below, some relevant default configuration is shown in regards to DCB and FCoE. **Programmable-mux** mode can also be used if a manual configuration via CLI for specific requirements such as multiple upstream LAGs or multiple FC fabrics is needed. See section '5.2 Programmable-mux mode' for more information.

```
dcb-map SAN_DCB_MAP
 priority-group 0 bandwidth 30 pfc off
 priority-group 1 bandwidth 30 pfc off
 priority-group 2 bandwidth 40 pfc on
 priority-pgid 0 0 0 2 1 0 0 0
```

Figure 60   DCB map configuration on fabric A IOA switch

```
fcoe-map SAN_FABRIC
 description SAN_FABRIC
 fc-map 0efc00
 fabric-id 1002 vlan 1002
```

Figure 61   FCoE map configuration on fabric A IOA switch

```
interface TenGigabitEthernet 0/4
 mtu 12000
 portmode hybrid
 switchport
 auto vlan
 dcb-map SAN_DCB_MAP
 fcoe-map SAN_FABRIC
!
 protocol lldp
  advertise management-tlv management-address system-name
 no shutdown
```

Figure 62   Server connecting port configuration on fabric A IOA switch

```
interface TenGigabitEthernet 0/33
 mtu 12000
 dcb-map SAN_DCB_MAP
!
 port-channel-protocol LACP
  port-channel 128 mode active
!
 protocol lldp
  advertise management-tlv management-address system-name
  no advertise dcbx-tlv ets-reco
  dcbx port-role auto-upstream
 no shutdown
```

Figure 63   Port-channel member interface configuration on fabric A IOA switch

```
interface Port-channel 128
 mtu 12000
 portmode hybrid
 switchport
 fip-snooping port-mode fcf
 no shutdown
 link-bundle-monitor enable
```

Figure 64   Port-channel interface configuration on fabric A IOA switch

```
interface FibreChannel 0/41
 fabric SAN_FABRIC
 no shutdown
!
interface FibreChannel 0/42
 fabric SAN_FABRIC
 no shutdown
!
interface FibreChannel 0/43
 fabric SAN_FABRIC
 no shutdown
!
interface FibreChannel 0/44
 fabric SAN_FABRIC
 no shutdown
!
interface FibreChannel 0/45
 fabric SAN_FABRIC
 no shutdown
!
interface FibreChannel 0/46
 fabric SAN_FABRIC
 no shutdown
!
interface FibreChannel 0/47
 fabric SAN_FABRIC
 no shutdown
!
interface FibreChannel 0/48
 fabric SAN_FABRIC
 no shutdown
```

Figure 65   FC interface configuration

5.2.2   Programmable-mux mode

It's also possible to deploy Dell IOA w/ FC Flex IOM using **Programmable-mux** mode. This allows for cli configuration of the IOA similar to the MXL. This mode can be useful if you want to customize the configuration based on specific requirements such as multiple upstream LAGs or multiple FC fabrics. To set the IOA to **programmable-mux** mode use the **stack-unit <unit-number> iom-mode programmable-mux** command from config mode as shown further below in Figure 66. Once set to **programmable-mux** mode, a reload will be required and configiuration will no longer be automated and will be required via the CLI similar to that which was done in the MXL configuration in Section 4.2 "Dell MXL w/ FC Flex IOM Configuration."

```
Dell_IOA_B1(conf)#stack-unit 0 iom-mode ?
programmable-mux          Programmable Mux mode
stack                     Stack mode
standalone                Standalone mode
vlt                       Vlt mode
```

Figure 66   IOA mode options

5.3   Verification on the Dell IOA w/ FC Flex IOM

To see information on NPIV devices logged into the fabric, use the **show npiv devices** command on the Dell IOA switch as shown below. Note the FCoE MAC is **0e:fc:00:01:0f:01** (the FCoE Map + FC_ID as expected).

```
Dell_IOA_B1#show npiv devices

ENode[0]:
ENode MAC       :    90:b1:1c:cc:e5:76
ENode Intf      :    Te 0/4
FCF MAC         :    d0:67:e5:ac:ac:ae
Fabric Intf     :    Fc 0/44
FCOE Vlan       :    1002
Fabric Map      :    SAN_FABRIC
ENode WWPN      :    20:01:90:b1:1c:cc:e5:76
ENode WWNN      :    20:00:90:b1:1c:cc:e5:76
FCoE MAC        :    0e:fc:00:01:0f:01
FC-ID           :    01:0f:01
LoginMethod     :    FLOGI
Secs            :    242
Status          :    LOGGED_IN
Dell_IOA_B1#
```

Figure 67   'show npiv devices' command output on fabric A IOA switch

To see currently active FCoE VN_Port sessions, use the **show fip-snooping sessions** command.

```
Dell_IOA_B1#show fip-snooping sessions
Enode MAC          Enode Intf   FCF MAC          FCF Intf    VLAN    FCoE MAC          FC-ID     Port WWPN                Port WWNN
---------          ----------   -------          --------    ----    --------          -----     ---------                ---------
90:b1:1c:cc:e5:76  Te 0/4       d0:67:e5:ac:ac:ae Fc 0/44    1002    0e:fc:00:01:0f:01 01:0f:01  20:01:90:b1:1c:cc:e5:76  20:00:90:b1:1c:cc:e5:76
Dell_IOA_B1#
```

Figure 68   'show fip-snooping sessions' command output on fabric A IOA switch

To see all FCoE end-node (ENodes) , use the **show fip-snooping enode** command.

```
Dell_IOA_B1#show fip-snooping enode
Enode MAC             Enode Interface      FCF MAC             VLAN        FC-ID
---------             ---------------      -------             ----        -----
90:b1:1c:cc:e5:76     Te 0/4               d0:67:e5:ac:ac:ae   1002        01:0f:01
Dell_IOA_B1#
```

Figure 69   'show fip-snooping enode' command output on fabric A IOA switch

To see a list of configured fcoe-maps, use the **show fcoe-map brief** command.

```
Dell_IOA_B1#show fcoe-map brief
    Fabric-Name   Fabric-Id        Vlan-Id     FC-MAP     FCF-Priority  Config-State      Oper-State
    SAN_FABRIC    1002             1002        0efc00             128        ACTIVE              UP
Dell_IOA_B1#
```

Figure 70   'show fcoe-map brief' command output on fabric A IOA switch

To see more detailed information on a given fcoe-map, use the **show fcoe-map <FCoE_MAP_NAME>** command. Notice below, the priority mapped to FCoE by default is **3**.

```
Dell_IOA_B1#show fcoe-map SAN_FABRIC

Fabric Name          SAN_FABRIC
Fabric Id            1002
Vlan Id              1002
Vlan priority        3
FC-MAP               0efc00
FKA-ADV-Period       8
Fcf Priority         128
Config-State         ACTIVE
Oper-State           UP
Members
Fc 0/41 Fc 0/42 Fc 0/43 Fc 0/44 Fc 0/45 Fc 0/46 Fc 0/47 Fc 0/48
Te 0/2 Te 0/6 Te 0/4
Dell_IOA_B1#
```

Figure 71   'show fcoe-map SAN_FABRIC' command output on fabric A IOA switch

6 M1000e FlexAddress enablement

When working with network adaptors that provide a MAC address as well as a World Wide Port and Node name, it's important to understand where these addresses originate. The M1000e chassis has a feature called Flexaddressing which allows for a virtual MAC address/WWPN to be linked to a server, therefore, if the CNA/adapter is changed later, zoning or configuration on upstream switches does not need to be modified.

Flexaddressing is a pool of 208 MAC's and 64 WWN contained on an SD card that can either be added at the Dell factory to the M1000e's chassis management controller (CMC) or to an existing M1000e CMC at a datacenter via a customer kit. Referring to Figure 72 below, enabling Flexaddressing is facilitated by:
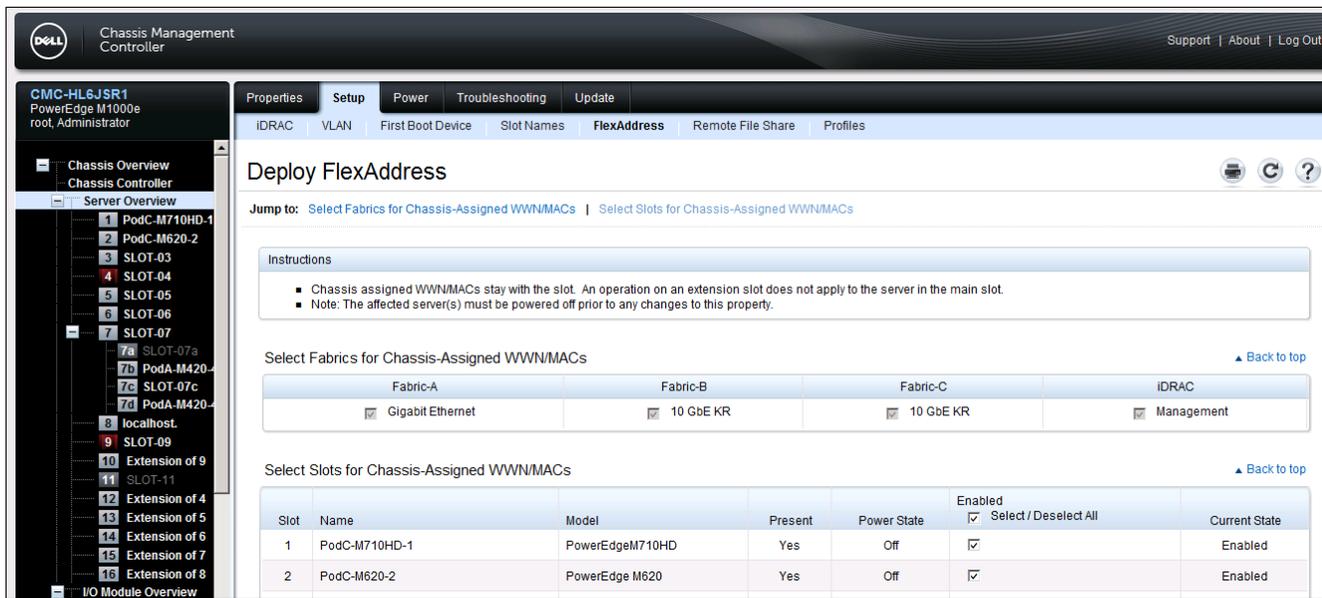

Figure 72   FlexAddress Enablement

1. Enter CMC's IP address

2. Click on Server Overview

3. Click on Setup

4. Click on FlexAddress

5. Choose to enable FlexAddressing to be either Fabric based or Slot based

Note: When the FlexAddress feature is deployed for the first time on a given server, it requires at least one power-down and power-up sequence of the server for the FlexAddress feature to take effect. Certain network configurations may require refreshing of network tables, such as ARP tables on IP switches and routers in order for new MAC/WWNs to be logged in.

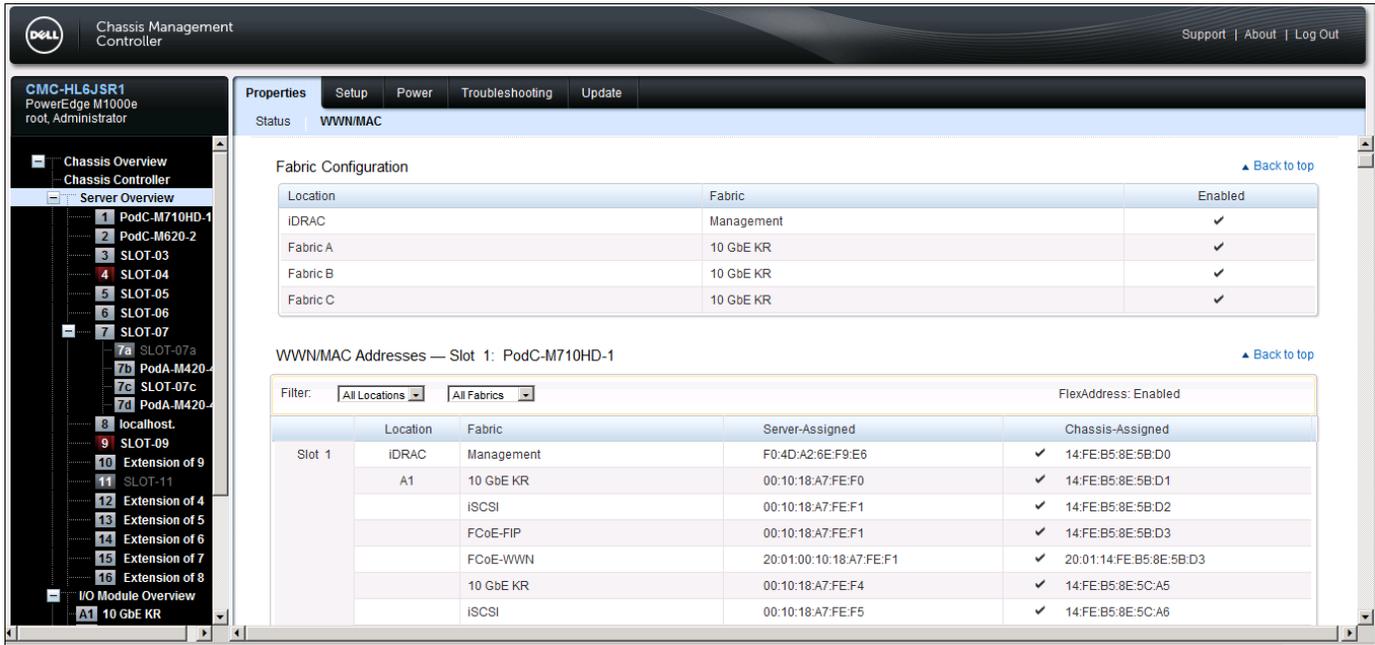To check that FlexAddressing is enabled:

Figure 73   FlexAddress Enablement Confirmation

1. Click on Server Overview

2. Click on Properties

3. Click on WWN/MAC

4. WWN/MAC Addresses section – FlexAddress: Enabled should be visible. Chassis-Assigned should have check marks next to each address.

Once FlexAddress has been confirmed as enabled, it's a simple matter of matching up the MAC or WWN of the network adaptor with the Server Slot and the Fabric that it is in.