Exploring Networking Configurations

for the PowerEdge 1655MC

The networking architecture integrated in the Dell™ PowerEdge™ 1655MC server blade system provides some significant advantages in terms of flexibility, resource consolidation, and cable management. However, because its design differs from that of a traditional server, deploying and managing the PowerEdge 1655MC requires that administrators have a detailed understanding of the architecture. This article describes major networking scenarios and provides deployment strategies.

BY YINGLIN YANG, MIKE J. ROBERTS, AND BRIAN DOTY

The Dell™ PowerEdge™ 1655MC server blade system integrates up to six server blades, two Ethernet switch modules, and an embedded remote management module into a highly dense, highly integrated 3U enclosure. Each integrated switch module is an independent Layer 2 switching device.[1]

In this highly integrated networking architecture, the lines between the management responsibilities of the server administrator and the network administrator begin to blur. Server administrators generally configure the operating system (OS) and network interface cards (NICs) and cable the server to the switch, while network administrators configure and manage the switches and routers. The highly integrated network infrastructure of the server blade[2] brings these distinct roles closer together, so server administrators might become responsible for configuring and managing the integrated switches. This article explains how to configure and deploy the PowerEdge 1655MC within an existing network infrastructure.

Deploying the PowerEdge 1655MC

The Dell PowerEdge 1655MC fits into many network environments immediately out of the box.[3] This server blade system can communicate with any Ethernet switches (such as Dell PowerConnect™ switches), hubs, routers, or NICs that support IEEE® 802.3, 802.3u, or 802.3ab standards. As a factory default, the four uplink (external) 10/100/1000BaseT ports of the integrated switches are set to auto-negotiation mode. In this mode, the four uplink ports support Auto-MDIX (medium dependent interface crossover), which automatically detects and adjusts the transmit and receive (Tx and Rx) signals to match those of a link partner. Auto-MDIX allows administrators to use either straight-through or crossover cables to connect an integrated switch to existing networking devices, switches, routers, or NICs.

For quick and simple deployment, administrators can plug a Category 5 (Cat 5) unshielded twisted pair (UTP) cable into the top switch (switch 1) of the PowerEdge 1655MC enclosure. Doing so provides connectivity to the

[1] For more information about integrated switches, see "Understanding the Gigabit Ethernet Switch Integrated in a PowerEdge Blade Server" by Yinglin Yang in Dell *Power Solutions*, August 2002.

[2] For more information about server blade architecture, see "Server Blades: An Emerging Server Architecture" by Mike J. Roberts in Dell *Power Solutions*, March 2002.

[3] For more information, see "Examining Network Performance of the PowerEdge 1655MC Blade Server" by Yinglin Yang and Mike J. Roberts in Dell *Power Solutions*, November 2002.

first embedded LAN on Motherboard (LOM), LOM 1, on each server blade. Plugging the UTP cable into switch 1 also allows network access to the PowerEdge 1655MC Embedded Remote Access/MC (ERA/MC) through port 1/11. Similarly, plugging a cable into the bottom switch (switch 2) enables access to LOM 2 on each server blade.

The integrated switches also have the IEEE 802.1D Spanning Tree Protocol (STP) enabled by default. STP automatically resolves bridge loops and provides some non-real-time network fault tolerance. For example, if two Cat 5 cables from the same external switch are plugged into the same integrated switch, STP automatically blocks one of the links to resolve the loop. If the active (forwarding) link fails, STP recalculates the spanning tree, removes the block on the other link, and sets the blocked link to a forwarding state, allowing the network to remain active until the problem can be repaired.

In the factory default configuration, the integrated switch is set as a Dynamic Host Configuration Protocol (DHCP) client. This feature allows quick deployment; DHCP automatically assigns an IP address to the switch, so that it can immediately be managed through the Web graphical user interface (GUI) or the Telnet command-line interface (CLI). Administrators also can assign a static IP address, network mask, and default gateway to the switch.

Configuring integrated switch uplink ports for forced mode

The ports on data center switches are often configured in forced mode: they are hard set to a specific speed, such as 100 Mbps full-duplex or 1000 Mbps full-duplex. In this case, administrators can use the Web GUI or CLI to configure the uplink ports of the PowerEdge 1655MC integrated switches in forced mode as well. From the CLI of the integrated switch, administrators can, for example, enter the following command to set port 7 to 100 Mbps full-duplex forced mode:

```
config
interface e 1/7
speed-duplex 100full
no negotiation
exit
```

The integrated switch has 11 ports; only ports 1/7, 1/8, 1/9, and 1/10 are configurable for speed and duplex. Ports 1/1 to 1/6, the switches' downlink (internal) ports, have solely 1000 Mbps full-duplex capability. Port 1/11 has 100 Mbps half-duplex capability, and it is reserved for management use.

During forced-mode operation, the Auto-MDIX feature of the integrated switch's uplink ports is no longer available. Thus, specific cables must be used to connect the integrated switch to external third-party switches, NICs, or routers. NIC ports or the Ethernet ports of routers are usually wired as MDI, whereas the normal ports of switches are wired as MDIX (MDI with internal crossover). To connect the uplink port of the integrated switch in the PowerEdge 1655MC to the MDI port of an external device, a crossover cable is required. To connect the uplink port of the integrated switch to the MDIX port on an external device, a straight-through cable is used.

Troubleshooting tips for cabling

To verify that the switches have been cabled correctly, administrators can try the following:

- Check the link (green) LED on the RJ45 connectors of the integrated switch. A green LED indicates that the network link has been established. If the green LED is off, check the cable and then ensure the corresponding port on the external switch is set to the right speed and duplex.
- Ping the IP addresses of the server blades or the integrated switches. Successfully pinging these IP addresses shows that network connectivity has been established. When pinging the switch, ensure that the management virtual LAN (VLAN) is on the same VLAN as the switch port through which the ping is performed. The switch management interface is statically set to VLAN 1, as are all switch ports by default.

Configuring the server blade system on two separate networks

Connecting servers to two mutually exclusive networks is one common network scenario for Web, application, and security servers. One network might be private, such as an intranet, management, or backup network, and another might be a public network, such as the Internet. The PowerEdge 1655MC is also well suited for a tiered network architecture.

Simple cable configurations for better performance

To configure the system for two separate networks, administrators can simply run a cable from the private network to the top switch and a second cable from the public network to the bottom switch (see the solid lines in Figure 1). This method provides simple

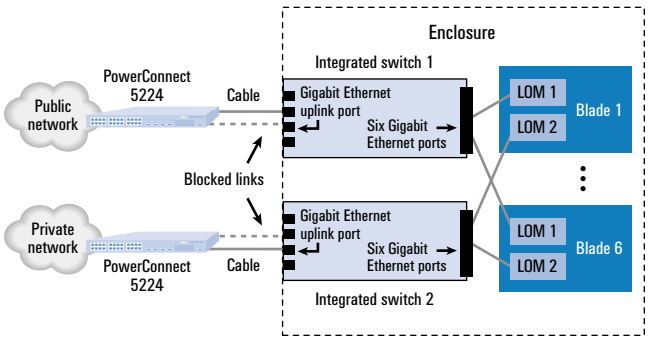

Figure 1. Simple solution for configuring the server blade system on two separate networks

and quick deployment, but it does not offer any fault tolerance for the uplink cables, switches, or LOMs, and it does not optimize bandwidth usage.

Alternatively, administrators can plug one or two more cables from the same network or subnet into the same integrated switch (see the dashed lines in Figure 1). This configuration offers a simpler solution and limited cable redundancy. STP will block redundant links. This configuration provides non-real-time failover of the network links and PowerEdge 1655MC switches. If an uplink fails, running network traffic will be interrupted until STP enables the standby link.

Channeling to maximize performance

Neither of the two plug-and-play methods maximizes performance and availability. For better results, administrators can configure a port channeling (also known as port trunking or link aggregation) protocol on the ports of the integrated switches and those of the external switch (see Figure 2). The channeling control protocol runs across these physical links and logically aggregates the links into a single virtual connection. The channel supports link failover for availability and provides improved bandwidth.

Port channeling protocols comprise two types of link aggregation: static and dynamic. The integrated switches in the PowerEdge

SETTING UP STATIC AND DYNAMIC CHANNELS

Figures A and B provide reference CLI commands for configuring static and dynamic channels on different switch products. These commands should be repeated on all ports that need to be configured.

Only the uplink ports 1/7, 1/8, 1/9, and 1/10 of the integrated switch in the PowerEdge 1655MC support port channeling protocols. Because two to four ports are required to form a channel, the commands for the integrated switch shown in Figure B should be repeated for the other ports to be included in the channel.

Troubleshooting tips for checking ports and switches on static and dynamic port channels

- To verify channel and member configuration on the PowerEdge 1655MC integrated switch, the PowerConnect 5224 switch, and the PowerConnect 3248 switch, enter:

```
show port-channel [1or2]
show interface status port-channel
show interface status ethernet [module/port]
```

- To verify channel and member configuration on the Cisco Catalyst® set-based switch or the PowerEdge 1655MC integrated switch, enter:

```
show port channel (for static channels)
show lacp-channel (for dynamic channels)
show port [module/port]
```

- To verify channel and member configuration on the Foundry FastIron® switch, enter:

```
show trunk
show interface ethernet [module/port]
```

Standards or switch products	Commands for configuring a static channel	Terms used in the various products
IEEE standards	Not applicable	Static link aggregation
Dell PowerEdge 1655MC integrated switch	config interface port-channel 1 exit interface ethernet 1/7 channel-group 1 exit	Port trunking
Dell PowerConnect 5224 and 3248	config interface port-channel 1 exit interface ethernet [module/port] channel-group 1 exit	Port trunking
Cisco Catalyst switch (set-based)	set port channel [port_list] on	EtherChannel®
Foundry FastIron	config t trunk server ethernet [module/port] to [module/port] exit	Trunk

Figure A. Reference commands to set up a static channel on various switch products

Standards or switch products	Commands for configuring a dynamic channel	Terms used in the various products
IEEE standards	Not applicable	Dynamic link aggregation or Link Aggregation Control Protocol (LACP)
PowerEdge 1655MC integrated switch	config interface ethernet 1/7 lacp exit	LACP
PowerConnect 5224 and 3248	config interface ethernet 1/x [where x is the port number used] lacp exit	LACP
Cisco Catalyst switch (set-based)	set channelprotocol lacp [module] set port lacp-channel [port_list] mode active	LACP
Foundry FastIron	config t interface ethernet [module/port] link-aggregation passive exit	LACP

Figure B. Reference commands to set up a dynamic channel on various switch products

1655MC support both. Administrators can use the CLI or Web GUI to set the static and dynamic channeling protocols on the integrated switch, as well as on other switch products existing in network infrastructures. For the appropriate commands, see the "Setting up static and dynamic channels" sidebar.

Achieving high availability through virtual teaming

Some environments, such as server farms or clusters, require high networking availability. Properly configuring the networking capabilities of the PowerEdge 1655MC can help achieve this goal. The Broadcom® Smart Load Balancing™ (SLB) technology enables bidirectional load balancing of IP traffic across multiple team members (NICs). Using channeling and NIC teaming can help to create a highly redundant networking configuration within the PowerEdge 1655MC system (see Figure 3).

To enable the two integrated switches to join a single subnet, an SLB virtual team is created on the two LOMs of each server blade. The SLB team binds the two physical LOMs into a single virtual network interface to the OS. The SLB team provides load balancing as well as device failover.

Because the SLB team does not use any probe packets to provide redundancy and failover, it does not consume any network bandwidth. Instead, to determine whether to trigger a failover, the virtual team checks the link status of the local network interfaces belonging to the team members. The local network interfaces in Figure 3 are the links between the LOMs and the internal ports of the integrated switches (hardwired through the midplane of the enclosure). The SLB team cannot probe the link status of the uplink ports, so administrators should aggregate, or channel, at least two of the uplink ports to achieve link redundancy. Commands and troubleshooting tips for configuring the 4-port channel on the uplink ports of the two integrated switches appear in the "Setting up static and dynamic channels" section of this article. When properly configured with channeling and SLB teams, the PowerEdge 1655MC can maintain network connectivity if any of the LOMs, internal switches, or uplink ports fails.

IDENTIFYING OS-SPECIFIC DEVICE NAMES FOR THE LOMs

To complete the network configurations–by assigning IP addresses to each LOM, for example–administrators must correctly identify the OS-specific network device names with their corresponding hardware device identifier. On a server blade that is running the Red Hat Linux OS, the device name associated with LOM 1 is eth0, and the name associated with LOM 2 is eth1. On the Microsoft Windows 2000 Server OS, the name Broadcom NetXtreme™ Gigabit Ethernet* is associated with LOM 1, and Broadcom NetXtreme Gigabit Ethernet 2 is the name associated with LOM 2.

*Gigabit Ethernet indicates compliance with IEEE 802.3ab and does not connote speeds of 1 Gbps.

Configuring an SLB team

Administrators can configure an SLB team by using the Broadcom Advanced Server Control Suite, a GUI for the Broadcom Advanced Server Program (BASP). This application software provides the configurations for load-balancing, fault-tolerance, and VLAN features. When running the Microsoft® Windows® 2000 Server or Advanced Server OS, administrators can use the following procedure to install and configure the SLB team:

1. Insert the Dell Server Assistant CD and allow it to auto-start.
2. Select the PowerEdge 1655MC and the Windows 2000 operating system.
3. Choose the Broadcom driver family download and save the zip file to the server.
4. Unzip the files and then run the setup program located in the MgmtApps directory.
5. Install the Control Suite application.
6. Choose "Yes" to install the BASP GUI application software and intermediate driver.

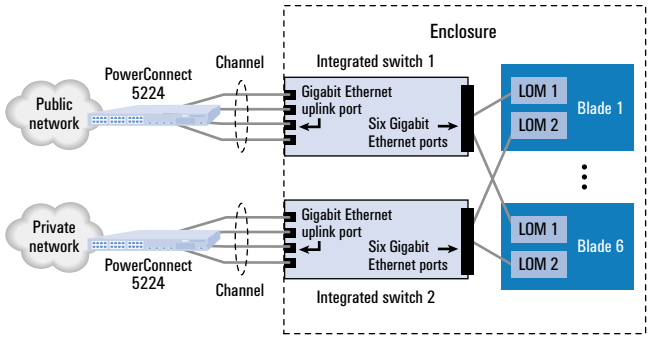

Figure 2. Port channeling solution for configuring the server blade system on two separate networks

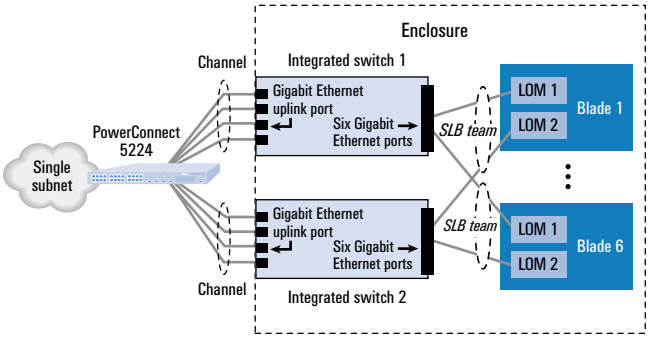

Figure 3. Configuration for redundancy and availability in a single subnet

7. Launch BASP by double-clicking the icon in the system tray or in the control suite.
8. Select the Load Balance/Virtual LAN tab.
9. Click on the Create Team button.
10. Name the team and choose the Smart Load Balance and Fail Over option.
11. Select the SLB team in the configuration box.
12. Select each of the adapters in the Available Adapters list and add them to the Load Balanced Members list.
13. Click "OK".

When running the Red Hat® Linux® OS, administrators can use the following procedure to configure the SLB team:

1. Install the RPM™ (Red Hat Program Manager) package of BASP:

```
rpm -i basplnx-version.i386.rpm
```

2. Copy a sample configuration script from the /etc/basp/samples directory to the /etc/basp directory. Note that the configuration script name must be prefixed with `team-`.

3. Edit the script so that it resembles the following:

```
# create the virtual adapter of the teaming.
The ID starts with 0.
TEAM_ID=0
TEAM_TYPE=0
TEAM_NAME=tSLB

# 1st physical interface in the team
TEAM_PA0_NAME=eth0
TEAM_PA0_ROLE=0

# 2nd physical interface in the team
TEAM_PA1_NAME=eth1
TEAM_PA1_ROLE=0
```

4. Start the teaming:

```
service basp start
```

Increasing redundancy with two external switches

The PowerEdge 1655MC can be configured to provide a highly redundant network environment when plugged into a single external switch. However, administrators may want to provide a further level of redundancy that accounts for failure of an external switch. Figure 4 shows the necessary network connections for the PowerEdge 1655MC system in this environment. Administrators implement this redundancy by creating two 2-port channels between each external switch and integrated switches 1 and 2. In this configuration, even if an external switch fails, the server blades maintain seamless network connectivity.

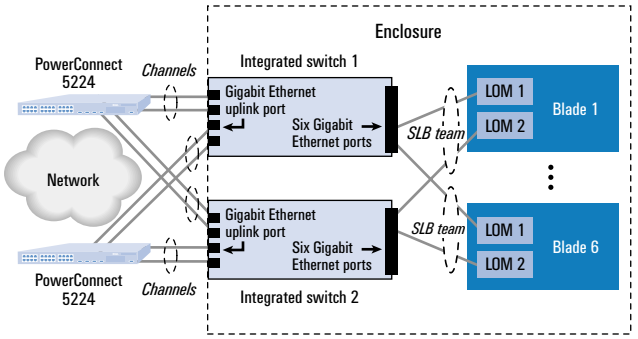

Figure 4. Configuration for redundancy with two external switches

Troubleshooting tips for networks configured to achieve redundancy

To troubleshoot network configurations with a single subnet or two external switches, administrators may:

- Verify that a compatible channel configuration is configured on both the external and the integrated switches. Configuring this incorrectly may cause a broadcast storm.
- Verify the configuration of the network interface of the virtual SLB team. At the command prompt of a server blade running the Microsoft Windows OS, enter:

```
ipconfig /all
```

  If the server blade is running Red Hat Linux, enter:

```
ifconfig
```

- Bring down the two team members before creating a virtual team on them. On a server blade running Red Hat Linux, enter:

```
ifconfig eth0 down
ifconfig eth1 down
```

  In Windows-based server blades, Windows automatically disables LOM 1 and LOM 2 when enabling SLB.

- Assign and ping the IP addresses of the virtual team. For Red Hat Linux, do not assign any IP addresses to eth0 and eth1.

Setting up server blades in VLANs

In some cases, the network topology requires the server blades to join more than one VLAN. Also, administrators may want to use VLANs rather than two physical networks (such as those shown in Figures 1 and 2) to create a private backup or management network and a public network.

SETTING UP VLANS

Figure C shows commands for configuring VLAN channels on different switch products.

Configuring VLANs on SLB teams

Administrators can use the Broadcom Advanced Server Control System (BASCS) GUI to easily create VLANs on SLB teams. For systems running Red Hat Linux, edit the Broadcom Advanced Server Program (BASP) configuration script to configure VLAN 10 and 20 on the SLB teams:

```
#create the virtual adapter of the teaming.
   The ID starts with 0.
TEAM_ID=0
TEAM_TYPE=0
TEAM_NAME=tSLB

# 1st physical interface in the team
TEAM_PA0_NAME=eth0
TEAM_PA0_ROLE=0
# 2nd physical interface in the team
TEAM_PA1_NAME=eth1
TEAM_PA1_ROLE=0

# 1st virtual interface of vlan 10 in the team
TEAM_VA0_NAME=sw0
TEAM_VA0_VLAN=10

# 2nd virtual interface of vlan 20 in the team
TEAM_VA1_NAME=sw1
TEAM_VA1_VLAN=20
```

To start the virtual adapters configured in the script, enter:

```
service basp start
```

Standards or switch products	Commands for configuring a VLAN channel	Terms used in the various products
IEEE standards	Not applicable	802.1Q VLAN
PowerEdge 1655MC integrated switch	config vlan database vlan 10 name vlan10 media ethernet state active vlan 20 name vlan20 media ethernet state active exit interface port-channel 1 switchport allowed vlan add 10 tagged switchport allowed vlan add 20 tagged exit interface e 1/x **[where x is the port number]** switchport allowed vlan add 10 tagged switchport allowed vlan add 20 tagged exit	VLAN, tagged, untagged
PowerConnect 5224 and 3248	config vlan 10 name vlan10 media ethernet state active vlan 20 name vlan20 media ethernet state active exit interface port-channel 1 switchport allowed vlan add 10 tagged switchport allowed vlan add 20 tagged end	VLAN, tagged, untagged
Cisco Catalyst switch (set-based)	Set channelprotocol pagp | lacp [module] Set port lacp-channel 5/1-2 mode active (or: set port channel 5/1-2 on) Set vlan 10 Set vlan 20 Set trunk 5/1 on dot1q	Dot1q vlan, Trunk for tagged, VLAN for untagged
Foundry FastIron	Config t interface ethernet [module/port] link-aggregation passive (or: trunk server ethernet [module/port] to [module/port]) exit vlan [vlan_id] tagged [module/port] to [module/port] exit	VLAN, tagged, untagged

Figure C. Reference commands to configure VLAN channels

Troubleshooting tips for VLAN configurations

- Verify the channel configuration
- To verify the VLAN configurations on the PowerEdge 1655MC integrated switch, the PowerConnect 5224 switch, and the PowerConnect 3248 switch, enter:

```
show vlan
show vlan [vlan_id]
show interface switchport ethernet
[module/port]
show interface switchport port-channel [1 or 2]
```

- To verify the VLAN configurations on the Cisco Catalyst set-based switch, enter:

```
show vlan
show vlan [vlan_id]
show trunk
```

- To verify the VLAN configurations on the Foundry FastIron switch, enter:

```
show vlan
show vlan [vlan_id]
```

- To verify the VLAN configuration of the network interface in Microsoft Windows, at the command prompt enter:

```
ipconfig /all
```

- Under Red Hat Linux, enter:

```
ifconfig
```

- Assign and ping the IP addresses of the VLANs.

When properly configured with channeling and SLB teams, the PowerEdge 1655MC can maintain network connectivity if any of the LOMs, internal switches, or uplink ports fails.

The integrated switch and LOM driver both support the IEEE 802.1Q VLAN standard. Properly configuring the integrated switches and LOMs of server blades can enable the servers to join and communicate on the desired VLANs. Figure 5 shows a configuration using two VLANs to map private and public networks connected to the PowerEdge 1655MC system. This configuration combines the use of VLANs with SLB teams to take advantage of the system's network bandwidth and fault-tolerance capabilities while servicing multiple subnets.

This scenario requires administrators to create VLAN 10 and VLAN 20 on the SLB team of each server blade and on each channel of the integrated switch. First, administrators should create a 4-port channel on each integrated switch and an SLB team on the two LOMs of each server blade. Then administrators should configure the VLAN on the port channel and the SLB teams, respectively. Commands to configure VLANs and troubleshooting tips are included in the "Setting up VLANs" section in this article. To configure more than two VLANs, administrators can simply repeat the VLAN creation commands.

Securely managing the integrated switch

Administrators can manage and configure integrated switches either through the serial port (out of band) or through the network (in band). In-band management protocols include Telnet, HTTP, and Simple Network Management Protocol (SNMP). These in-band management interfaces allow administrators to manage the switch remotely; however, remote access to the switches may present a security concern for administrators.

Disabling remote management capability

Administrators may want to disable the remote management interfaces on the integrated switch for security purposes. To disable the HTTP agent, enter the following command at the switch's CLI:

```
no ip http server
```

To turn off the SNMP management interface, enter:

```
no snmp-server community <community-string>
```

The Telnet agent cannot be disabled on the integrated switch. To make the switch inaccessible through Telnet, remove any IP addresses assigned to the switch by entering the following commands in the CLI:

```
config
interface vlan 1
no ip address
exit
```

These sample commands assume that the switch management IP interface belongs to VLAN 1, the default setting for all switch ports and the switch management IP interface. These commands will also disable the other two in-band management protocols (HTTP and SNMP) of the switch.

Assigning the switch management interface to a separate VLAN

In many cases, data center administrators require remote management capability and a high degree of security. A potential solution is to assign the switch management IP interface to a separate VLAN.

For example, to implement the scenario shown in Figure 6, first create port channel 1, which aggregates four uplink ports of integrated switch 1. Second, using the CLI of the integrated switch,

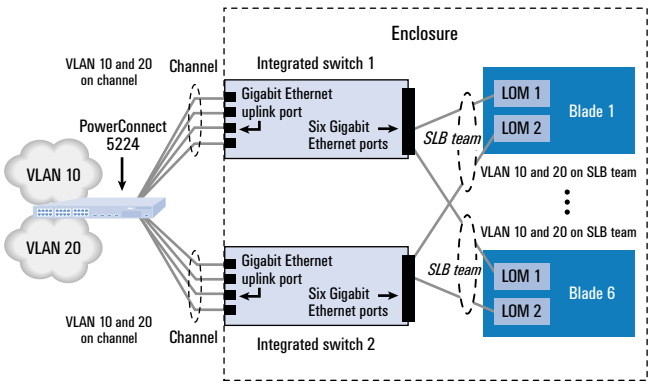

Figure 5. Combining VLANs and SLB teams

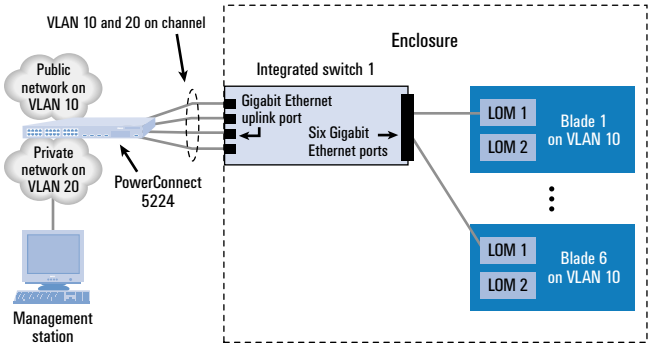

Figure 6. Assigning the switch management IP interface to a separate VLAN

configure VLAN 10 and VLAN 20 on the channel:

```
config
vlan database
vlan 10 name vlan10 media ethernet state active
vlan 20 name vlan20 media ethernet state active
exit
interface port-channel 1
switchport allowed vlan add 10 tagged
switchport allowed vlan add 20 tagged
end
```

Third, assign an IP address to the switch management IP interface in VLAN 20:

```
config
interface vlan 20
ip address [ip address, subnet mask]
exit
ip default-gateway [default gateway address]
exit
```

These instructions will create a switch management IP interface to VLAN 20.

Fourth, assign all internal ports 1/1 through 1/6 to VLAN 10, enabling all server blades to be reached by the public network on VLAN 10:

```
interface 1/1
switchport allowed vlan add 10
switchport native vlan 10
exit
```

Repeat these commands on the interfaces 1/2 through 1/6.

Now all the server blades are configured to communicate on VLAN 10, the public network, while the management interface of the integrated switch is on VLAN 20, the private network. The server blades and the switch are on virtually separate LANs.

Protecting access to integrated switches

In addition to remote access security, the integrated switches support two tiers of passwords—user and enable—that can be encrypted by commands. The switch also supports remote authentication dial-in user service (RADIUS), which provides a means of authenticating to an external RADIUS server. This capability protects access to the integrated switches in the same manner that many external network resources are protected.

Maximizing PowerEdge 1655MC network capabilities

Properly configuring the PowerEdge 1655MC can help administrators to easily deploy the server blade system into different networking environments. Configuration considerations include cabling, server network conductivity, channeling protocols, NIC teaming, VLANs, and security options. By using wise deployment strategies, administrators can achieve superior networking performance and availability from the PowerEdge 1655MC.

Yinglin Yang (yinglin_yang@dell.com) is a development engineer in the Server Networking and Communication Group at Dell. He has an M.S. in Electrical Engineering from Texas Tech University and a B.S. in Electrical Engineering from Chongqing University in China.

Mike J. Roberts (mike_j_roberts@dell.com) is a product manager at Dell. He is responsible for the Dell server blade product line. Mike has a B.A. from Marquette University and an M.B.A from The University of Texas at Austin.

Brian Doty (brian_doty@dell.com) is a test engineer consultant in the Dell Product Group, where he leads test efforts on Dell servers. Brian has a B.S. in Computer Engineering from the University of Kansas.