

Deploying the ProLiant BL p-Class GbE2 Interconnect Switch into a Cisco-based Network

HOWTO



Abstract.....	3
Introduction.....	3
Terminology.....	3
Typographical conventions	4
Critical features for successful deployment	4
Virtual local area network	4
VLAN tagging.....	4
IP management interface	5
Spanning tree protocol	5
Spanning tree groups.....	6
Bridging protocol data unit.....	6
Root bridge	6
Bridge priority.....	6
Port cost	6
Port priority.....	6
Multiple spanning tree groups.....	7
VLAN and STG configuration guidelines.....	7
Multi-link trunking.....	8
Load balancing.....	8
Multi-link trunking and spanning tree	8
Multi-link trunking configuration guidelines.....	9
Common topological examples	10
Topology 1: Fully meshed	11
VLAN configuration.....	12
Spanning tree configuration.....	13
Port configuration	14
Multi-link trunking and EtherChannel.....	14
Topology 2: Partial mesh	15

VLAN configuration	15
Spanning tree configuration.....	16
Port configuration	17
Multi-link trunking and EtherChannel	17
Topology 3: Straight-through	18
VLAN configuration	18
Spanning tree configuration.....	19
Port configuration	19
Multi-link trunking and EtherChannel	20
Topology summary	20
Securing the GbE2 Interconnect Switch	21
Management interfaces	21
Command line interface	21
Browser based interface.....	21
Setting source IP address range.....	21
SNMP management.....	21
RADIUS.....	22
Passwords	22
Additional best practices	22
Appendix A: GbE2 Interconnect Switch architecture	24
Appendix B: GbE2 Interconnect Switch default settings	25
For more information.....	30

Abstract

This HOWTO provides best practice guidelines and configuration examples for installation of the ProLiant BL p-class GbE2 Interconnect Switch into a Cisco-based network. This guide is meant to be a tool to help direct decisions in planning, optimization and securing the GbE2 Interconnect Switch environment. While the best practices and configurations examples in this document could be used in real world environments, they are to be used only as guidelines. This HOWTO does not serve as a replacement for the GbE2 Interconnect Switch user guides, rather it is meant to serve as a supplement to this documentation.

The intended audience for this paper includes engineers and system administrators familiar with the ProLiant BL p-Class GbE2 Interconnect Switch. For readers not familiar with GbE2 Interconnect Switch, please see the [ProLiant BL p-Class GbE2 Interconnect Switch Overview](#) white paper¹ as well as the user documentation that shipped with the GbE2 Interconnect Switch.

Introduction

This HOWTO identifies best practice guidelines and configuration examples for installation of the ProLiant BL p-class GbE2 Interconnect Switch into a Cisco-based network consisting of redundant Catalyst 6509 switches with the Catalyst switch operating system (CatOS). However, the examples in this document can be used as general guidelines appropriate for network infrastructures consisting of other Cisco switches, with the CatOS or Internetwork Operating System Software (IOS), and network devices from other vendors including Nortel, Extreme, Foundry, 3Com, etc.

The GbE2 Interconnect Switch is intended for applications that require up to 1000 megabits per second (Mb/s) Gigabit Ethernet network adapter (NIC) consolidation, advanced network feature support (including future planned options for layer 3 and 4-7 switching), server blade Fibre Channel pass-through, and future upgradeability for 10 Gigabit Ethernet bandwidth connectivity to the network. For additional information on the GbE2 Interconnect Switch, please see the [ProLiant BL p-Class GbE2 Interconnect Switch Overview](#) white paper¹.

For best practice guidelines for the entire p-Class system, see the [HP ProLiant BL System Best Practices Guide](#)² and the [HP ProLiant BL System Common Procedures Guide](#)³

Terminology

The terminology that differs between the Cisco Catalyst 6509 switch and the GbE2 Interconnect Switch documentation is identified in Table 1.

Table 1. Network terminology cross reference

HP ProLiant GbE2 Interconnect Switch	Cisco Catalyst 6509 switch
VLAN tagging, 802.1Q tagging	trunking, VLAN or 802.1Q encapsulation
port VLAN identification (PVID)	VLAN identification (VLANID)
link aggregation, multi-link trunking (MLT)	EtherChannel, channeling
spanning tree protocol group (STG)	spanning tree instance
IEEE 802.1s, multiple spanning tree	per VLAN spanning tree (PVST), PVST+

¹ Available at http://wws1.pro.compaq.com/support/reference_library/viewdocument.asp?source=5982-2175EN.xml&dt=21.

² Available at http://wws1.pro.compaq.com/support/reference_library/viewdocument.asp?source=351359-001.xml&dt=264.

³ Available at http://wws1.pro.compaq.com/support/reference_library/viewdocument.asp?source=351360-001.xml&dt=264.

Typographical conventions

The following table describes the switch command typographic styles used in this guide:

Table 2. Switch command typographical conventions

HP Typeface	Meaning	Example
AaBbCc123	This type displays in command examples and shows text that must be typed in exactly as shown.	/cfg/vlan
<i><AaBbCc123></i>	This italicized type displays in command examples as a parameter placeholder. Replace the indicated text with the appropriate real name or value when using the command. Do not type the brackets.	/cfg/vlan <vlan number>

To distinguish between ProLiant BL p-Class GbE2 Interconnect Switch and Catalyst 6509 commands, each command will be preceded by a GbE2>> and 6509#, respectively.

Critical features for successful deployment

Understanding VLANs and VLAN tagging (VLAN trunking), spanning tree protocol, and multi-link trunking (channeling) is critical to the successful deployment of the GbE2 Interconnect switch. Each of these topics is covered providing a high-level primer inclusive of GbE2 Interconnect Switch command introduction and general configuration guidelines. Specific commands and configuration steps follow in the section titled "Common topological examples". For additional information, refer to the [HP ProLiant BL p-Class GbE2 Interconnect Switch Application Guide](#)⁴ chapters 3 and 4.

Virtual local area network

A virtual local area network (VLAN) is a network topology configured according to a logical scheme rather than the physical layout. VLANs are used to logically segment traffic into different broadcast domains allowing packets to be forwarded only between ports within the VLAN. This enhances performance by conserving bandwidth and improves security by limiting traffic to specific domains.

The standard practice of configuring VLANs on an Ethernet switch is by assigning each port to a specific VLAN. In this port-based VLAN implementation, the switch identifies the specific VLAN membership of a packet per the port on which it was received. Individual VLANs are defined via a configurable VLAN number. The VLAN number is known as port VLAN identification (PVID) on GbE2 Interconnect Switches and VLAN identification (VLANID) on Cisco Catalyst switches. The GbE2 Interconnect Switch allows any PVID value from 2 to 4095 with PVID 1 reserved as the default VLAN. The default GbE2 Interconnect switch configuration has all ports assigned to PVID 1.

The IEEE industry standard for VLANs is 802.1Q. Each GbE2 Interconnect Switch supports 255 port-based IEEE 802.1Q VLANs. The GbE2 Interconnect Switch VLAN menu can be found under:

```
GbE2>> /cfg/vlan <vlan number>
```

VLAN tagging

VLAN tagging (often called VLAN trunking or encapsulation by Cisco) is the process of inserting into a data frame a tag identifying its VLAN membership. VLAN tagging allows each switch port to belong to multiple VLANs and provides the information switches need to create VLANs across the network.

⁴ Available at http://wws1.pro.compaq.com/support/reference_library/viewdocument.asp?source=331403-001.xml&dt=264.

Switch ports may be configured as tagged or untagged. A tagged port may receive tagged or untagged frames and is capable of forwarding the frames appropriately. When a VLAN tagged frame arrives at a tagged port, the switch looks at the PVID in the tag to determine its VLAN membership before switching the packet to the correct port. If an untagged frame arrives on a tagged port, the switch will tag the frame with the PVID of that port. If a frame exits the switch via a tagged port, any tag will remain on the frame unchanged as it exits.

An untagged port is only capable of switching untagged frames. Therefore, an untagged port will only see and accept incoming untagged frames. Frames received by the untagged port will be forwarded without any changes to the frame. For frames exiting the switch via an untagged port, any tag will be stripped from the frame before its forwarded.

GbE2 Interconnect Switch ports may be individually configured as tagged or untagged using the following command:

```
GbE2>> /cfg/port <port number>/tag ena
```

When implementing VLAN tagging on the GbE2 Interconnect Switch, the PVID values must be established correctly between devices communicating in the VLAN. This option is found under:

```
GbE2>> /cfg/port <port number>/pvid <PVID number>
```

IP management interface

The IP management interface provides management access to the GbE2 Interconnect Switch over an IP network. By default, the IP management interface is configured to request its IP address from a bootstrap protocol (BOOTP) server, but the IP address may also be assigned manually resulting in BOOTP being disabled.

Carefully consider how VLANs are configured within the GbE2 Interconnect Switch to ensure remote communication to the switch remains possible. In order to access the GbE2 Interconnect Switch for remote configuration, SNMP trap messages, and other remote management functions, confirm at least one IP management interface on the switch has a VLAN defined.

It is possible to inadvertently disable access to management functions if the port associated with the IP management interface is excluded from VLAN membership. Likewise, if all IP interfaces remain within the default VLAN (VLAN 1) and all ports are configured for a different VLAN, such as VLAN 2, then GbE2 Interconnect Switch management features are effectively disabled. To avoid these situations, it is suggested that all ports used for remote GbE2 Interconnect Switch management remain on the default VLAN and that an IP management interface be assigned to the default VLAN.

On the GbE2 Interconnect Switch, assign the IP management interface to a VLAN using the commands:

```
GbE2>> /cfg/ip/if <number>/mask <mask>/addr <address>  
GbE2>> /cfg/ip/if <number>/vlan <vlan#>/ena/apply
```

Spanning tree protocol

Spanning tree protocol (STP) is used to ensure that redundant paths within a layer 2 network do not result in broadcast loops. For a layer 2 Ethernet network to function correctly, only one active path may forward frames between any two switches at a given time.

Redundant connections between network switches can create loops or multiple forwarding paths. In layer 2 networks, these loops cause duplicate packets to be forwarded to the same destination over and over again until the network is completely saturated, which in turn prevents valid traffic from traversing the network. STP configures the network by allowing a switch to use the most efficient path while forcing the remaining redundant paths into a standby (blocked) state. If the forwarding path fails, STP automatically activates a standby path to sustain network operations.

Spanning tree groups

STP examines the network topology and defines a tree structure spanning all switches in a given layer 2 network domain. These layer 2 network domains are called spanning tree groups (STG). STGs are created by assigning a group of layer 2 switches to be part of a separate layer 2 network domain. When STP examines the network topology it only considers eliminating loops within a single STG. Within a layer 2 domain, there may be multiple STGs each operating its own individual STP algorithm to eliminate layer 2 loops.

The IEEE industry standard for STP is defined in 802.1D. The GbE2 Interconnect Switch meets the IEEE 802.1D standard and further provides interoperability with Cisco's Per VLAN Spanning Tree Plus (PVST+) via the use of STGs; refer to the "Multiple spanning tree groups" section for more information on PVST+.

Bridging protocol data unit

All network devices that are members of a spanning tree send out packets called bridging protocol data units (BPDU). A BPDU is a 64-byte packet sent by all switches participating in the spanning tree protocol providing information about each other. The BPDU includes information known as switch or bridge priority, port cost, and port priority used to establish a spanning tree root switch and which paths to designate as forwarding and blocking.

Root bridge

The STP root switch (or root bridge) is the base of the spanning tree topology much like the roots of a tree. All redundant paths to the root bridge within the spanning tree network are placed in the blocked mode. The root bridge is chosen by all the switches based on the results of the BPDU exchange process.

Bridge priority

The bridge priority is used to determine what switch is the root bridge. Bridge priority is a numerical value that may be configured on a switch. The lower a bridge's priority value, the greater the chance it has of becoming the root bridge. If all switches are configured with the same default bridge priority setting, the switch with the lowest MAC address in the STP network becomes the root switch. Bridge priority is automatically assigned by the STP process, or may be manually configured on the GbE2 Interconnect Switch using the following command:

```
GbE2>> /cfg/stp <stg number>/brg/prior <new bridge priority>
```

Port cost

The port cost is a value assigned to each switch port. The port cost information is exchanged within the BPDU to help determine the lowest cost path to the root switch. The port with the lowest cost path is used as the forwarding port between two segments in the STG. All remaining paths within each segment are placed in a blocked state.

The objective is to use the fastest links ensuring the route with the lowest cost is chosen. The spanning tree protocol assigns lower values to high-bandwidth ports, such as Gigabit Ethernet, to encourage their use. The cost of a port also depends on whether the port operates at full-duplex (lower cost) or half-duplex (higher cost). For example, a 100-Mb/s (Fast Ethernet) link has a STP assigned "cost" of 10 in half-duplex mode, and a cost of 5 in full-duplex mode. Port cost is automatically assigned by the STP process, or manually set on the GbE2 Interconnect Switch using the following command:

```
GbE2>> /cfg/stp <stg number>/port<number>/cost <1-65535>
```

Port priority

The port priority is yet another STP value assigned to each switch port. In case of identical port costs, the port priority is used as a tie breaker to determine the lowest path cost to the root switch and the resulting forwarding port for each segment. Therefore, in a network topology segment that has multiple paths with the same port cost, the port with the lowest port priority becomes the designated

port for the segment. It is also possible for the ports to have identical port priorities. If this is the case, the port number becomes the final decision criteria. Port priority is automatically assigned by the STP process, or manually set on the GbE2 Interconnect Switch using the following command:

```
GbE2>> /cfg/stp <stg number>/port <port number>/prior <1-255>
```

Multiple spanning tree groups

The IEEE 802.1D standard considers the network topology of all the switches participating in the spanning tree network as one broadcast domain or one spanning tree group (STG). It does not consider the logical VLAN implementation. Ports within different VLANs are logically separated broadcast domains. With the 802.1D implementation, paths that form physical loops within the network may be placed in a blocking state even though the VLAN topology would have not caused a layer 2 broadcast storm.

To prevent this, the IEEE standard 802.1s was adopted as an extension to the original 802.1D standard. It allows multiple STGs within a network switch taking into consideration the VLAN logical topology. Forwarding and blocking decisions are now made according to the BPDU information within its own broadcast domain. IEEE 802.1s utilizes the 802.1Q VLAN tagging method in its implementation. Prior to the adoption of 802.1s, Cisco developed a similar protocol known as Per VLAN Spanning Tree (PVST). PVST uses the Cisco proprietary Intra Switch Link (ISL) method of VLAN tagging. A more recent update to the protocol known as PVST+ provides the same functionality as PVST, but utilizes the 802.1Q VLAN tagging method.

The GbE2 Interconnect Switch integrates into a PVST+ environment through the use of STGs. In the GbE2 implementation, an administrator creates an STG and then assigns a VLAN to it. This differs from the Cisco implementation where an administrator creates a VLAN and then a spanning tree instance (i.e. STG) is automatically assigned to it. The PVST+ interoperability feature on the GbE2 Interconnect Switch includes the following:

- Tagged ports may belong to more than one STG, but untagged ports can belong to only one STG.
- When a tagged port belongs to more than one STG, egress BPDUs are tagged to identify their STG membership.
- An untagged port cannot span multiple STGs.
- Sixteen STGs are supported per GbE2 Interconnect Switch.
- The default STG 1 can hold multiple VLANs, all other STGs (groups 2–16) can hold one VLAN.

On each GbE2 Interconnect Switch, the six external ports (ports 19-24) and the crosslink ports (ports 17-18) are by default in STG 1. The STG can be changed for each port using the following command:

```
GbE2>> /cfg/stp <stg number>/port <port number>
```

VLAN and STG configuration guidelines

When creating a VLAN on the GbE2 Interconnect Switch, that VLAN automatically belongs to the default STG 1. To add the VLAN in another STG, it must be assigned to another STG. Keep the following rules in mind when creating VLANs and assigning STGs:

- The default VLAN (VLAN 1) cannot be removed from the default STG 1.
- VLANs must be contained within a single STG; a VLAN cannot span multiple STGs.
- When a VLAN spans multiple switches, the VLAN must be within the same STG (have the same STG ID) across all the switches.

- If ports are tagged, all trunked ports can belong to multiple STGs.
- A port that is not a member of any VLAN cannot be added to a STG. The port must be added to a VLAN, and that VLAN added to the desired STG.
- Tagged ports can belong to more than one STG, but untagged ports can belong to only one STG.
- When a tagged port belongs to more than one STG, the egress BPDUs are tagged to distinguish the BPDUs of one STG from those of another STG.
- An untagged port cannot span multiple STGs.
- When a port is removed from a VLAN that belongs to an STG, that port will also be removed from the STG. However, if that port belongs to another VLAN in the same STG, the port remains in the STG.
- An STG cannot be deleted, only disabled. If you disable the STG while it contains VLAN members, STP will be off on all ports belonging to that VLAN.
- If STP any port in the trunk is set to forwarding, the remaining ports in the trunk will also be set to forwarding

Multi-link trunking

Multi-link trunking (MLT), also known as link aggregation and port trunking (and EtherChannel by Cisco), combines multiple physical switch ports into a single logical port called a trunk. The bandwidth of the trunk is the multiple of the bandwidth of the individual links. An algorithm automatically applies load balancing to the ports in the trunk. A port failure within the group causes the network traffic to be directed to the remaining ports. Load balancing is maintained whenever a link in a trunk is lost or returned to service.

The industry standard for multi-link trunking is IEEE 802.3ad. Cisco has developed a similar multi-link trunking method known as EtherChannel. The GbE2 Interconnect Switch supports twelve IEEE 802.3ad (without LACP⁵) trunks per switch interoperable with EtherChannel. Each trunk may contain two to six ports providing a 12-Gbps aggregate throughput full duplex.

Load balancing

Within the trunk, the load distribution is determined by information embedded within the data frame. For traffic that does not contain IP information, the GbE2 Interconnect Switch elects the port with the lowest port number in the trunk to be the designated port for forwarding traffic. For traffic that contains IP addresses, the GbE2 Interconnect Switch will calculate the designated trunk port for forwarding traffic by using the statistical load balancing algorithm that considers the packet's source and destination IP addresses.

Multi-link trunking and spanning tree

A typical network is designed with multiple links between switches to provide increased bandwidth and redundant connections. In layer 2 networks, redundant links between switches create loops or multiple forwarding paths resulting in broadcast storms. The spanning tree protocol will identify these loops and place ports in a blocked state to eliminate the possibility of multiple forwarding paths. However, this defeats the purpose of using multiple connects between switches for increased bandwidth. MLT can be used to provide redundant links while ensuring that STP does not block this redundancy. Within a multi-link trunk, all the individual ports are seen as one logical by the spanning tree protocol.

⁵ Link aggregation control protocol (LACP) is an enhancement over EtherChannel and other static multi-link trunking methods. LACP dynamically learns about the link status and takes decisions on which links to use for and load balancing and failback in case of link failure. As a result, IEEE 802.3ad with LACP is often called dynamic trunking.

Multi-link trunking configuration guidelines

When creating trunks, consider the following configuration rules that determine how a trunk reacts in the network topology.

- Confirm the GbE2 Interconnect Switch ports to be trunked are set to enabled.
- All trunks must originate from one device, and lead to one destination device. For example, it is not possible to combine a port from two different switches into one trunk.
- Any physical switch port can belong to only one trunk.
- Trunking from non-HP devices must comply with Cisco EtherChannel technology.
- All ports within a trunk (trunk members) must be assigned to the same VLAN configuration before the trunk can be enabled.
- All ports within the trunk must be configured to full duplex.
- If the VLAN settings of any one trunk member are modified, the change cannot be applied until the VLAN settings of all trunk members are modified.
- When an active GbE2 Interconnect Switch port is configured in a trunk, the port becomes a trunk member using the following trunk command:

```
GbE2>> /cfg/trunk/trunk <trunk group>/add <port number>/ena
```

The spanning tree parameters for the port will change to reflect the new trunk settings.

- All trunk members must be in the same STG. If all ports are tagged, then all the ports within trunk can belong to multiple STGs; otherwise, only one STG membership is allowed.
- When a trunk is enabled, the spanning tree participation setting of the trunk takes precedence over that of any individual trunk member.
- If the spanning tree protocol participation of any trunk member is changed to enabled or disabled, the spanning tree participation of all members of that trunk changes similarly.
- A trunk member cannot be a monitoring port in a port mirroring configuration.
- Trunks act as a single logical port, but cannot be monitored by a monitor port; however, individual trunk members can.
- The port speeds of each trunk member must be the same.

Common topological examples

Three common topological configurations are provided that highlight the main requirements for a successful configuration of the GbE2 Interconnect Switch into a Cisco Catalyst 6509 network environment. Actual deployment scenarios may differ from the provided examples. Configuration guidelines and configuration settings for VLAN, spanning tree, multi-link trunking, and the port configuration are provided for each topology.

The provided configuration steps are not the complete set necessary for a fully functioning system. Use these examples as guidelines when deploying the GbE2 Interconnect Switch in any specific environment. This information is intended to supplement the user documentation included with the GbE2 Interconnect Switch and the Catalyst 6509 switches.

The configuration steps assume the GbE2 Interconnect Switch configuration is set to factory default. For a list of GbE2 Interconnect Switch default settings, see Appendix B. To set the GbE2 Interconnect Switch configuration settings to the factory default, use the following procedure:

1. Select the configuration block per the command:

```
GbE2>> /boot/conf
```

2. The system indicates which configuration block is currently set to be loaded at the next reset and prompts you to enter a new choice. Enter "factory" as the configuration block.
3. To make the new configuration block changes active, the GbE2 Interconnect Switch must be reset per the command:

```
GbE2>> /boot/reset
```

You are then prompted to confirm your request.

CAUTION: Prior to changing the configuration block, it is recommended any desired changes to the current configuration block be saved, see chapter 6 of the [HP ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide⁶](#).

NOTE: Resetting the GbE2 Interconnect Switch causes the spanning tree protocol to restart. This process can be lengthy depending on the topology of the network.

All topologies assume a p-Class server blade enclosure with eight ProLiant BL20p series servers and two GbE2 Interconnect Switches operating at layer 2. The two GbE2 Interconnect Switches are connected to two Cisco Catalyst 6509 switches creating a redundant architecture. The Catalyst 6509 switches are in turn connected to each other operating at layer 3. Each layer 2 configuration utilizes two VLANs: the default VLAN (VLAN 1) for data and a second VLAN (VLAN 2) for Integrated Lights-Out (iLO) Advanced management. The separate iLO management VLAN isolates management traffic from the rest of the network minimizing the risk of a remote network device intercepting sensitive data. To avoid layer 2 loops and to interoperate with Cisco's PVST+, a STG for each VLAN is configured for the topologies requiring the use of spanning tree. For redundancy, EtherChannel compatible multi-link trunking is configured on each topology where multiple links exist between switches.

The configuration of each topology provides tradeoffs in the areas of performance, availability, design complexity, and the need for spanning tree (Table 3). Each topology utilizes all four exterior Gigabit Ethernet ports (uplinks ports 19-22), but in a different design configuration. The front panel Gigabit Ethernet ports (ports 23 and 24) are not connected to the Cisco-based network and are available for local switch management, port analysis, and other administration tasks, and as additional uplinks. Each topology could be modified by connecting these front panel ports to the network increased bandwidth, added availability, creation of a remote port diagnostic network, etc.

⁶ Available at http://wws1.pro.compaq.com/support/reference_library/viewdocument.asp?source=331404-001%20.xml&dt=264.

Table 3. Topology overview

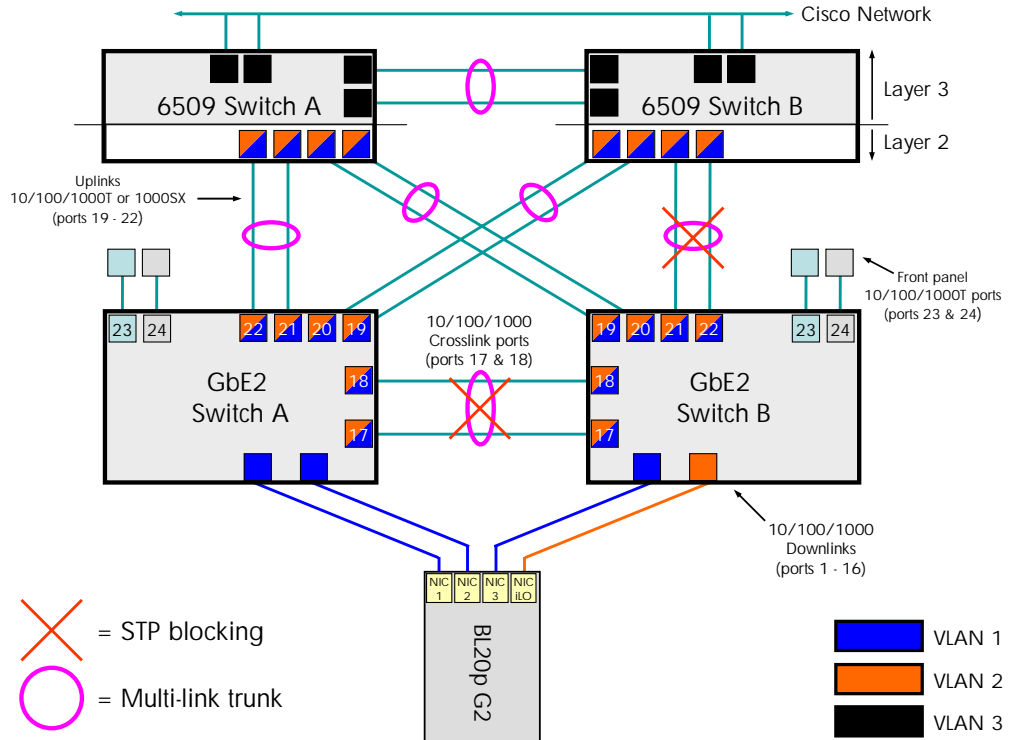
Topology	Benefits	Drawbacks
1	<ul style="list-style-type: none"> • Fully meshed design optimizing availability: <ul style="list-style-type: none"> – Maximum resiliency even with a dual link or dual non-like switch failure. – Loss of up to eight physical links without an interruption in service. • Separate VLAN isolating management traffic from the rest of the network for increased security. 	<ul style="list-style-type: none"> • Requires spanning tree protocol and its resulting convergence time delays. • Design complexity creates potential for added support. • Non-optimal throughput from GbE2 Interconnect Switches to the Catalyst 6509 switches.
2	<ul style="list-style-type: none"> • Near fully meshed providing a high level of availability <ul style="list-style-type: none"> – Service maintained after a dual link failure. – Single switch failure maintains connectivity to two other switches. • Less complex configuration as compared to topology 1 decreasing support requirements. • Separate VLAN isolating management traffic from the rest of the network for increased security. 	<ul style="list-style-type: none"> • Requires spanning tree protocol and its resulting convergence time delays. • Non-optimal availability. • Configuration is still reasonably complex. • Non-optimal throughput from GbE2 Interconnect Switches to the Catalyst 6509 switches.
3	<ul style="list-style-type: none"> • Good level of availability, service maintained after a dual link failure. • Optimal throughput between the GbE2 Interconnect Switches and the Cisco network. • Spanning tree protocol is not required. • Topology simplification decreasing support requirements. • Separate VLAN isolating management traffic from the rest of the network for increased security. • From any single GbE2 Interconnect Switch uplink, communicate to all ProLiant BL server NICs and manage both switches. 	<ul style="list-style-type: none"> • Reduced availability: <ul style="list-style-type: none"> – Greater possibility of certain failures causing partial to total loss of service. – Loss of a switch results in all traffic converging through surviving switch, potentially causing performance issues.

Topology 1: Fully meshed

Topology 1 is a fully meshed configuration designed for maximum fault resiliency (Figure 1). It optimizes system availability by utilizing multiple links to ensure no single switch failure or even multiple dissimilar switch failures will result in an outage. However, this configuration requires the use of the spanning tree protocol. A pair of links between the GbE2 Interconnect Switch and Catalyst 6509 switches must be placed in an STP blocking state reducing overall throughput. Additionally, the design complexity of this configuration creates the potential for added support.

Topology 1 is ideal for network administrators who desire maximum availability at the expense of complexity, some throughput, and the management and potential convergence delays of STP.

Figure 1. Topology 1 fully meshed architecture



VLAN configuration

Per Figure 1, this configuration utilizes two VLANs for GbE2 Interconnect Switch ports and ports connecting the GbE2 switches to the 6509 switches: the default VLAN (VLAN 1) for data and a second VLAN that will be created (VLAN 2) for iLO Advanced management. Any remaining Catalyst 6509 ports must be separate from VLANs 1 and 2 and are collectively represented in Figure 1 as VLAN 3.

NOTE: VLAN tagging (trunking) must be enabled on all ports within VLANs 1 and 2.

To configure the VLANs on the switches, perform the following:

1. On both Cisco 6509 switches, set the preferred VLAN trunk protocol mode:

```
6509# set vtp mode <mode>
```

2. On both Cisco 6509 switches, configure both VLANs (VLAN 1 and 2) per the command:

```
6509# set vlan 1-2
```

3. On both Cisco 6509 switches, enable 802.1Q tagging (VLAN trunking) on all ports connected to both GbE2 Interconnect Switches per the command:

```
6509# set trunk <module number>/<port number> nonegotiate dot1q 1-2
```

4. On both GbE2 Interconnect Switches (GbE2 Switch A and GbE2 Switch B), enable tagging (VLAN encapsulation) on the four uplink ports (19-22) and the two crosslink ports (17-18) per the command:

```
GbE2>> /cfg/port <number>/tag ena
```

5. On both GbE2 Interconnect Switches, create the management VLAN 2 per the command:

```
GbE2>> /cfg/vlan 2/ena
```

6. On the GbE2 Switch B, add iLO ports to VLAN 2. This includes the GbE2 Switch B ports 1, 3, 5, 7, 9, 11, 13, and 15. Use the following command to add the ports to VLAN 2:

```
GbE2>> /cfg/vlan 2/add <port number>
```

7. On both GbE2 Interconnect Switches, add ports 17-18 and 19-22 to VLAN 2 using the command identified in step 5.

Spanning tree configuration

To eliminate the physical loops in this topology, STP is required on the ports connecting the GbE2 Interconnect Switches and the Cisco Catalyst 6509 switches and on the GbE2 Interconnect Switch crosslink ports. This configuration also utilizes two STGs to separate the logical loops. To configure the STP on the switches, perform the following:

1. On the 6509 Switch A, set the bridge priority to 0 for both VLANs 1 and 2, per the commands:

```
6509# set spantree priority 0 1  
6509# set spantree priority 0 2
```

2. On 6509 Switch B, set the bridge priority to a slightly higher value than 6509 Switch A to ensure that if the primary root bridge fails, this second Catalyst 6509 switch becomes the root bridge. This allows the Catalyst 6509 switches to control the network and centralizes the administration. Use the following commands:

```
6509# set spantree priority 4096 1  
6509# set spantree priority 4096 2
```

NOTE: Do not alter the default bridge priority for either GbE2 Interconnect Switch. This will ensure that one of the Catalyst 6509 switches always becomes the root bridge.

3. On both GbE2 Interconnect Switches create a second STG (STG2) and add VLAN 2 to it. To perform these tasks, use the following command:

```
GbE2>> /cfg/stp 2/on/add 2
```

NOTE: When adding a port to a VLAN that belongs to an STG, the port is also added to the STG. However, if the port being added to the VLAN is untagged and already a member of another STG, that port will not be added to an additional STG because an untagged port cannot belong to more than one STG.

4. Modify the port cost to 100 on GbE2 Switch B, ports 17, 18, 21 and 22 using the command:

```
GbE2>> /cfg/stp 2/port <port number>/cost 100
```

This ensures that STP will behave in a predictable manner by blocking the GbE2 Switch B crosslink ports 17 and 18 and the uplink ports 21 and 22, and by placing all uplinks on GbE2 Switch A in a forwarding state.

Port configuration

Configure the ports on the Cisco and HP blade switches by performing the steps listed below.

1. On both Cisco Catalyst 6509 switches, configure the port speed and negotiation settings on all ports connected to the GbE2 Interconnect Switches per the commands:

```
6509# set port speed <module_number>/<port_number> auto
6509# set port negotiation <module_number>/<port_number> enable
```

2. Configure the port speed and negotiation for GbE2 Interconnect Switch ports 19-22 (on both GbE2 Interconnect Switches), per the commands:

```
GbE2>> /cfg/port <port number>/gig/auto on
GbE2>> / cfg/port <port number>/gig/mode full
```

NOTE: This step is only necessary if the GbE2 Interconnect Switch default configuration has been modified.

3. Utilizing Figure 1 as a reference, connect:
 - GbE2 Switch A ports 21 and 22 to 6509 Switch A
 - GbE2 Switch A ports 19 and 20 to 6509 Switch B.
 - GbE2 Switch B ports 21 and 22 to 6509 Switch B
 - GbE2 Switch B ports 19 and 20 to 6509 Switch A.

NOTE: Ports 17 and 18 on each GbE2 Interconnect Switch are already connected across the server blade enclosure backplane and do not require any further physical connectivity.

Multi-link trunking and EtherChannel

This topology requires the creation of two trunks on each GbE2 Interconnect Switch and three EtherChannel groups on each the Catalyst 6509 switch. Additionally, a default preconfigured trunk (trunk 1) exists for the crosslink ports between the GbE2 Interconnect Switches. Other trunks may also be present, as shown for VLAN 3 in Figure 1.

NOTE: For this topology, GbE2 Interconnect Switch ports 19 and 20 represent trunk 2 and ports 21 and 22 represent trunk 3.

1. On both Cisco Catalyst 6509 switches, configure EtherChannel on the ports connected to both GbE2 Interconnect Switches per the command:

```
6509# set port channel <module number>/<port number> mode on
```

2. Configure trunk 2 on each GbE2 Interconnect Switch per the commands:

```
GbE2>> /cfg/trunk 2
GbE2>> ena
GbE2>> add 19
GbE2>> add 20
```

3. On each GbE2 Interconnect Switch repeat the above steps, but for trunk 3 using ports 21 and 22.

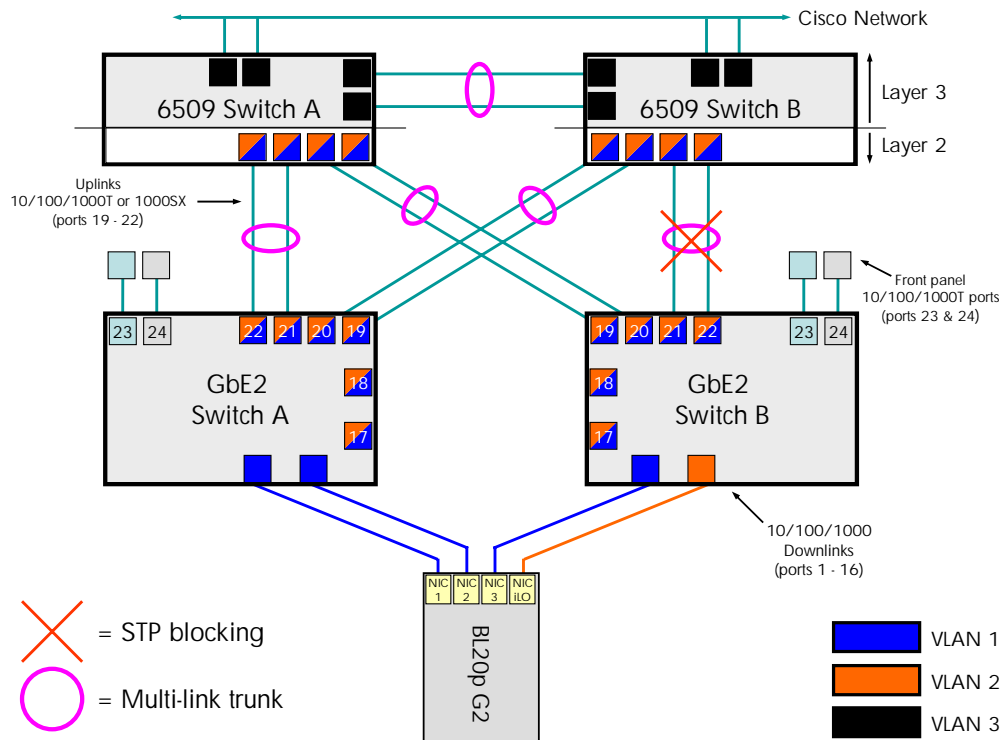
NOTE: It should not be necessary to configure trunk 1 as it is part of the GbE2 Interconnect Switch default configuration. However, if the default configuration has been modified, configure trunk 1 using ports 17 and 18.

Topology 2: Partial mesh

Topology 2 is very similar in design to topology 1 except the GbE2 Interconnect Switch crosslink ports are manually disabled (Figure 2). This configuration maintains a high level of availability to ensure the loss of up multiple physical links without an interruption in service. The removal of the crosslink ports or switch-to-switch links between the GbE2 Interconnect Switches decreases availability to some degree. However, it has the added positive effect of a less complex configuration as compared to topology 1 thereby decreasing support requirements.

Topology 2 is ideal for network administrators who need to maintain a high level of availability while minimizing some design complexity.

Figure 2. Topology 2 partial mesh architecture



VLAN configuration

Like topology 1, this configuration utilizes the default VLAN and iLO VLAN 2 for GbE2 Interconnect Switch ports and ports connecting the GbE2 switches to the 6509 switches. Any remaining Catalyst 6509 ports must be separate from VLANs 1 and 2 and are collectively represented in Figure 2 as VLAN 3.

NOTE: VLAN tagging (trunking) must be enabled on all ports within VLANs 1 and 2.

To configure the VLANs on the switches, perform the following:

1. On both Cisco 6509 switches, set the preferred VLAN trunk protocol mode:

```
6509# set vtp mode <mode>
```

2. On both Cisco switches, configure both VLANs (VLANs 1 and 2) per the command:

```
6509# set vlan 1-2
```

3. On both Cisco switches, enable 802.1Q tagging (VLAN trunking) on all ports connected to both of the GbE2 Interconnect Switches per the command:

```
6509# set trunk <module number>/<port number> nonegotiate dot1q 1-2
```

4. On both GbE2 Interconnect Switches (GbE2 Switch A and GbE2 Switch B), enable tagging (VLAN encapsulation) on the four uplink ports (19-22) per the command:

```
GbE2>> /cfg/port <port number>/tag ena
```

5. On both GbE2 Interconnect Switches, create the management VLAN 2 per the command:

```
GbE2>> /cfg/vlan 2/ena
```

6. On the GbE2 Switch B, add iLO ports to VLAN 2. This includes the GbE2 Switch B ports 1, 3, 5, 7, 9, 11, 13, and 15. Use the following command to add the ports to VLAN 2:

```
GbE2>> /cfg/vlan 2/add <port number>
```

7. On both GbE2 Interconnect Switches, add ports 19-22 to VLAN 2 using the command identified in step 5.

Spanning tree configuration

To eliminate the physical loops in this topology, STP is required on the ports connecting the GbE2 Interconnect Switches and the Cisco Catalyst 6509 switches. This configuration also utilizes two STGs to separate the logical loops. To configure the STP on the switches, perform the following:

1. On the 6509 switch A, set the bridge priority to 0 for VLANs 1 and 2, per the commands:

```
6509# set spantree priority 0 1  
6509# set spantree priority 0 2
```

2. On 6509 Switch B, set the bridge priority to a slightly higher value than 6509 Switch A to ensure that if the primary root bridge fails, this second Catalyst 6509 switch becomes the root bridge. This allows the Catalyst 6509 switches to control the network and centralizes the administration. Use the following commands:

```
6509# set spantree priority 4096 1  
6509# set spantree priority 4096 2
```

NOTE: Do not alter the default bridge priority for either GbE2 Interconnect Switch. This will ensure that one of the Catalyst 6509 switches always becomes the root bridge.

3. On both GbE2 Interconnect Switches create a second STG (STG2) and add VLAN 2 to it. To perform these tasks, use the following command:

```
GbE2>> /cfg/stp 2/on/add 2
```

NOTE: When adding a port to a VLAN that belongs to an STG, the port is also added to the STG. However, if the port being added to the VLAN is untagged and already a member of another STG, that port will not be added to an additional STG because an untagged port cannot belong to more than one STG.

4. Modify the port cost to 100 on GbE2 Switch B, ports 21 and 22 using the command:

```
GbE2>> /cfg/stp 2/port <port number>/cost 100
```

This ensures that STP will behave in a predictable manner by blocking the GbE2 Switch B uplink ports 21 and 22 and by placing all uplinks on GbE2 Switch A in a forwarding state.

Port configuration

Configure the ports on the Cisco and HP blade switches by performing the steps listed below.

1. On both Cisco Catalyst 6509 switches, configure the port speed and negotiation settings on all ports connected to the GbE2 Interconnect Switches per the commands:

```
6509# set port speed <module_number>/<port_number> auto
6509# set port negotiation <module_number>/<port_number> enable
```

2. Configure the port speed and negotiation for GbE2 Interconnect Switch ports 19-22 (on both GbE2 Interconnect Switches), per the commands:

```
GbE2>> /cfg/port <port number>/gig/auto on
GbE2>> / cfg/port <port number>/gig/mode full
```

NOTE: This step is only necessary if the GbE2 Interconnect Switch default configuration has been modified.

3. On both GbE2 Interconnect Switches, disable ports 17 and 18 per the command:

```
GbE2>> /cfg/port <port_number>/disable
GbE2>> apply
GbE2>> save
```

4. Utilizing Figure 2 as a reference, connect:

- GbE2 Switch A ports 21 and 22 to 6509 Switch A
- GbE2 Switch A ports 19 and 20 to 6509 Switch B.
- GbE2 Switch B ports 21 and 22 to 6509 Switch B
- GbE2 Switch B ports 19 and 20 to 6509 Switch B.

Multi-link trunking and EtherChannel

This topology requires the creation of two trunks on each GbE2 Interconnect Switch and three EtherChannel groups on each the Catalyst 6509 switch. Other trunks may also be present, as shown for VLAN 3 in Figure 2.

NOTE: For this topology, GbE2 Interconnect Switch ports 19 and 20 represent trunk 2 and ports 21 and 22 represent trunk 3.

1. On both Cisco Catalyst 6509 switches, configure EtherChannel on the ports connected to both GbE2 Interconnect Switches per the command:

```
6509# set port channel <module number>/<port number> mode on
```

2. Configure trunk 2 on each GbE2 Interconnect Switch per the commands:

```
GbE2>> /cfg/trunk 2
GbE2>> ena
GbE2>> add 19
GbE2>> add 20
```

3. On each GbE2 Interconnect Switch repeat the above steps, but for trunk 3 using ports 21 and 22.

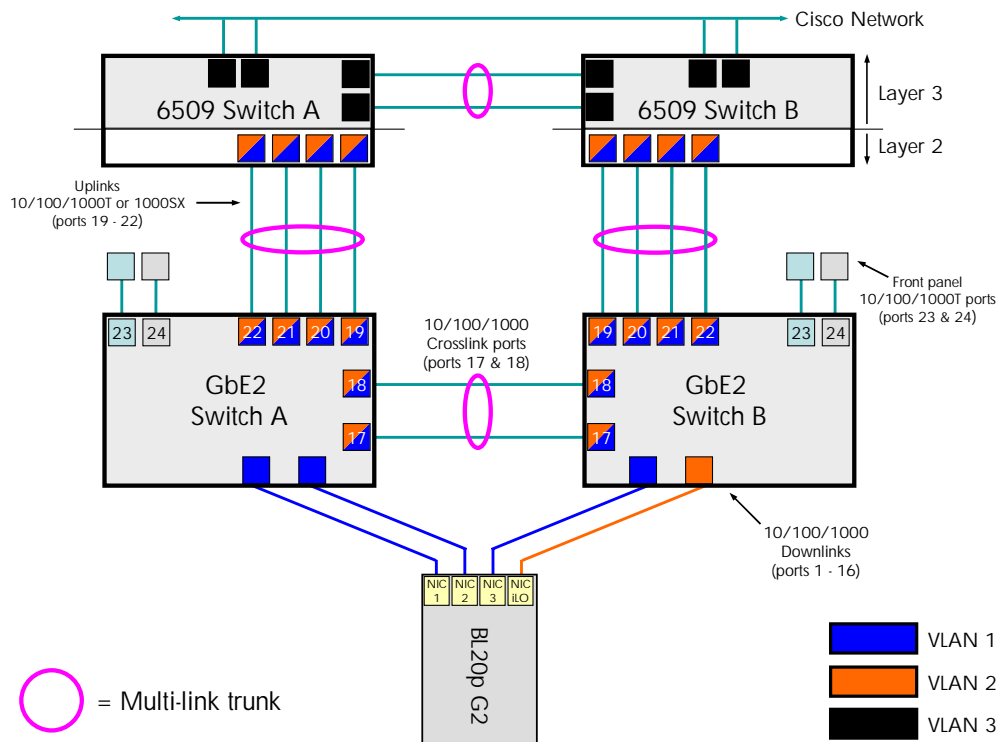
Topology 3: Straight-through

Topology 3 is a “straight-through” design providing a simplified architecture with maximum throughput (Figure 3). Spanning tree protocol is not required further simplifying this configuration. However, this topology has reduced availability; certain failures can cause partial to total loss of service with the greater possibility of a performance bottleneck.

The GbE2 Interconnect switch crosslinks are enabled in this configuration. These Gigabit Ethernet links permit management of both switches and access to all ProLiant BL server NICs from any single GbE2 Interconnect Switch uplink. The crosslinks may be disabled in this configuration, but it is not advised. In this case, a failure of any one switch (whether GbE2 Interconnect Switch or Catalyst 6509) would cause the loss of service to one half the NICs on each ProLiant server.

Topology 3 is ideal for network administrators who desire a simplified architecture that provides high levels of performance at the expense of some availability.

Figure 3. Topology 3 straight-through architecture



VLAN configuration

Consistent with the other two topologies, this configuration utilizes the default VLAN and iLO VLAN 2 for GbE2 Interconnect Switch ports and ports connecting the GbE2 switches to the 6509 switches. Any remaining Catalyst 6509 ports must be separate from VLANs 1 and 2 and are collectively represented in Figure 3 as VLAN 3.

NOTE: VLAN tagging (trunking) must be enabled on all ports within VLANs 1 and 2.

To configure the VLANs on the switches, perform the following:

1. On both Cisco 6509 switches, set the preferred VLAN trunk protocol mode:

```
6509# set vtp mode <mode>
```

2. On both Cisco switches, configure both VLANs (VLAN 1 and 2) per the command:

```
6509# set vlan 1-2
```

3. On both Cisco switches, enable 802.1Q tagging (VLAN trunking) on all ports connected to both of the GbE2 Interconnect Switches per the command:

```
6509# set trunk <module number>/<port number> nonegotiate dot1q 1-2
```

4. On both GbE2 Interconnect Switches (GbE2 Switch A and GbE2 Switch B), enable tagging (VLAN encapsulation) on the four uplink ports (19-22) and the two crosslink ports (17-18) per the command:

```
GbE2>> /cfg/port <port number>/tag ena
```

5. On both GbE2 Interconnect Switches, create the management VLAN 2 per the command:

```
GbE2>> /cfg/vlan 2/ena
```

6. On the GbE2 Switch B, add iLO ports to VLAN 2. This includes the GbE2 Switch B ports 1, 3, 5, 7, 9, 11, 13, and 15. Use the following command to add the ports to VLAN 2:

```
GbE2>> /cfg/vlan 2/add <port number>
```

7. On both GbE2 Interconnect Switches, add ports 19-22 to VLAN 2 using the command identified in step 5.

Spanning tree configuration

This topology does not require spanning tree to be enabled as by design no loops are present.

1. On both GbE2 Interconnect Switches, disable spanning tree, per the commands:

```
GbE2>> /cfg/stp 1/off  
GbE2>> save
```

CAUTION: If the GbE2 Interconnect Switches are already connected to the Cisco network, perform step 3 in the "Port configuration" section before disabling spanning tree.

Port configuration

Configure the ports on the Cisco and HP blade switches by performing the steps listed below.

1. On both Cisco Catalyst 6509 switches, configure the port speed and negotiation settings on all ports connected to the GbE2 Interconnect Switches per the commands:

```
6509# set port speed <module_number>/<port_number> auto  
6509# set port negotiation <module_number>/<port_number> enable
```

2. Configure the port speed and negotiation for GbE2 Interconnect Switch ports 19-22 (on both GbE2 Interconnect Switches), per the commands:

```
GbE2>> /cfg/port <port number>/gig/auto on  
GbE2>> / cfg/port <port number>/gig/mode full
```

NOTE: This step is only necessary if the GbE2 Interconnect Switch default configuration has been modified.

3. On both GbE2 Interconnect Switches, disable ports 17 and 18 per the command:

```
GbE2>> /cfg/port <port_number>/disable  
GbE2>> apply  
GbE2>> save
```

4. Utilizing Figure 2 as a reference, connect:
 - GbE2 Switch A ports 19, 20, 21 and 22 to 6509 Switch A
 - GbE2 Switch B ports 19, 20, 21 and 22 to 6509 Switch B

Multi-link trunking and EtherChannel

This topology requires the creation of two trunks on each GbE2 Interconnect Switch and three EtherChannel groups on each the Catalyst 6509 switch. Additionally, a default preconfigured trunk (trunk 1) exists for the crosslink ports between the GbE2 Interconnect Switches. Other trunks may also be present, as shown for VLAN 3 in Figure 1.

NOTE: For this topology, GbE2 Interconnect Switch ports 19 and 20 represent trunk 2 and ports 21 and 22 represent trunk 3.

1. On both Cisco Catalyst 6509 switches, configure EtherChannel on the ports connected to both GbE2 Interconnect Switches per the command:

```
6509# set port channel <module number>/<port number> mode on
```

2. Configure trunk 2 on each GbE2 Interconnect Switch per the commands:

```
GbE2>> /cfg/trunk 2
GbE2>> ena
GbE2>> add 19
GbE2>> add 20
GbE2>> add 21
GbE2>> add 22
```

NOTE: It should not be necessary to configure trunk 1 as it is part of the GbE2 Interconnect Switch default configuration. However, if the default configuration has been modified, configure trunk 1 using ports 17 and 18.

Topology summary

In summary, the three provided topologies differ primarily in their data throughput, the need for spanning tree, design complexity, and level of availability (Table 4).

Table 4. Topology summary

	Topology		
	1	2	3
Level of availability	Optimal	High	Good
GbE2 Interconnect Kit total uplink throughput (full duplex)	High (12 Gbps)*	High (12 Gbps)*	Optimal (16 Gbps)*
Spanning tree protocol required	Yes	Yes	No
Design complexity	High	Medium	Low
Utilizes EtherChannel / multi-link trunking	Yes	Yes	Yes
Separate VLAN for isolating iLO management traffic	Yes	Yes	Yes
Manage both GbE2 Interconnect Switches from any of its uplink ports	No	No	Yes
Communicate to all NICs from any GbE2 Interconnect Switch uplink port	No	No	Yes

* May be increased by an additional 8 Gbps full duplex by utilizing the Gigabit Ethernet ports on the front of each GbE2 Interconnect Switch

Securing the GbE2 Interconnect Switch

HP recommends a variety of best practices to ensure the security of the network is maintained when deploying GbE2 Interconnect Switches. The suggestions provided here are applicable to the GbE2 Interconnect Switch independent of specific vendor used for the network infrastructure components.

Management interfaces

The GbE2 Interconnect Switch provides many standard management access features, some of which provide potential security risks within a given network environment. There are several recommended practices that can be applied to decrease exposure and increase security.

Command line interface

The GbE2 Interconnect Switch command line interface (CLI) allows switch management locally via the serial port or remotely via Telnet and SSH. Since Telnet transmits data in clear text, HP recommends using only secure shell (SSH) for remote CLI management, unless the end-to-end path has no external access and there are no known means by which this traffic can be monitored.

It is recommended the default Telnet TCP port 23 be changed using the commands:

```
GbE2>> /cfg/sys/tnport <TCP port number>
```

Additionally, HP recommends modifying the default CLI idle timeout setting of five minutes to a value consistent with network security practices, per the command:

```
GbE2>> /cfg/sys/idle <idle time in minutes>
```

Browser based interface

The GbE2 Interconnect Switch browser based interface (BBI) allows remote switch management via a web console. Like Telnet, the HTTP interface can be vulnerable to security attacks. HP recommends the BBI only be used when the integrity and security of the connection cannot be compromised.

Additionally it is recommended that the default TCP connection port of 80 be changed using the command:

```
GbE2>> /cfg/sys/wport <TCP connection port number>
```

Setting source IP address range

To limit management access to the GbE2 Interconnect Switch without having to configure filters for each switch port, HP recommends the source IP address range be configured using the commands:

```
GbE2>> /cfg/sys/mnet <management network, such as 192.192.192.0>  
GbE2>> /cfg/sts/mmask <management mask, such as 255.255.255.128>
```

For the above example management network and mask addresses, any packet is discarded that is addressed to a GbE2 Interconnect Switch IP interface with a source IP address outside the range of 192.192.192.1 to 192.192.192.127.

SNMP management

The GbE2 Interconnect Switch software provides simple network management protocol (SNMP) v1.0 support for access through network management software, such as HP OpenView and Insight Manager 7. For improved security, HP recommends the default read and write community strings (public and private, respectively) and the trap host community strings be changed using the commands:

```
GbE2>> /cfg/snmp/rcomm <SNMP read community string>  
GbE2>> /cfg/snmp/wcomm <SNMP write community string>  
GbE2>> /cfg/snmp/t1comm <1st trap host community string>  
GbE2>> /cfg/snmp/t2comm <2nd trap host community string>
```

RADIUS

The GbE2 Interconnect Switch, acting as the RADIUS client, communicates to the RADIUS server to authenticate and authorize a remote administrator using the protocol definitions specified in RFC 2138 and 2866. The use of RADIUS is highly recommended as it allows for accounting and auditing of connections that the GbE2 Interconnect Switch does not natively possess. For configuration procedures, refer to the [HP ProLiant BL p-Class GbE2 Interconnect Switch Application Guide](#) chapter 1⁷ and the [HP ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide](#) chapter 6⁸.

Passwords

HP recommends all GbE2 Interconnect Switch default passwords be changed at initial configuration and as regularly as required under the network security policies.

- To change the user, operator, and administrator management passwords, see the [HP ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide](#) chapter 3⁸.
- To change the SCP administrator password, see the [HP ProLiant BL p-Class GbE2 Interconnect Switch Application Guide](#) chapter 1⁷.

Additional best practices

Additional steps are recommended to ensure that the GbE2 Interconnect Switch is easily serviceable and readily available within the network environment. The suggestions provided here are applicable to the GbE2 Interconnect Switch independent of specific vendor used for the network infrastructure components.

- Record the GbE2 Interconnect Switch MAC address located on the exterior switch label.
- Save a copy of the switch configuration setting by performing one or both of the following methods:
 1. Capture the configuration file a terminal screen using the dump command:

```
GbE2 >> /cfg/dump
```

This will print the text to the console screen which can be saved to a text file.

2. Save the configuration to a using TFTP or SCP using one of the following commands:

```
GbE2>> /cfg/ptcfg <TFTP server IP address> <config. file name>
```

```
GbE2>> scp <switch IP address>:getcfg <local file name>
```

NOTE: The output file is formatted with line-breaks, but no carriage returns. The file cannot be viewed with editors that require carriage returns (such as Microsoft Notepad).

- Install a second version of the operating system firmware. The GbE2 Interconnect Switch includes two flash regions to store firmware. For more information, refer to the [HP ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide](#) chapter 8⁹.
- Configure redundant syslog servers. For more information, refer to the [HP ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide](#) chapter 6⁹.

⁷ Available at http://wws1pro.compaq.com/support/reference_library/viewdocument.asp?source=331403-001.xml&dt=264

⁸ Available at http://wws1pro.compaq.com/support/reference_library/viewdocument.asp?source=331404-001%20.xml&dt=264.

⁹ Available at http://wws1pro.compaq.com/support/reference_library/viewdocument.asp?source=331404-001%20.xml&dt=264.

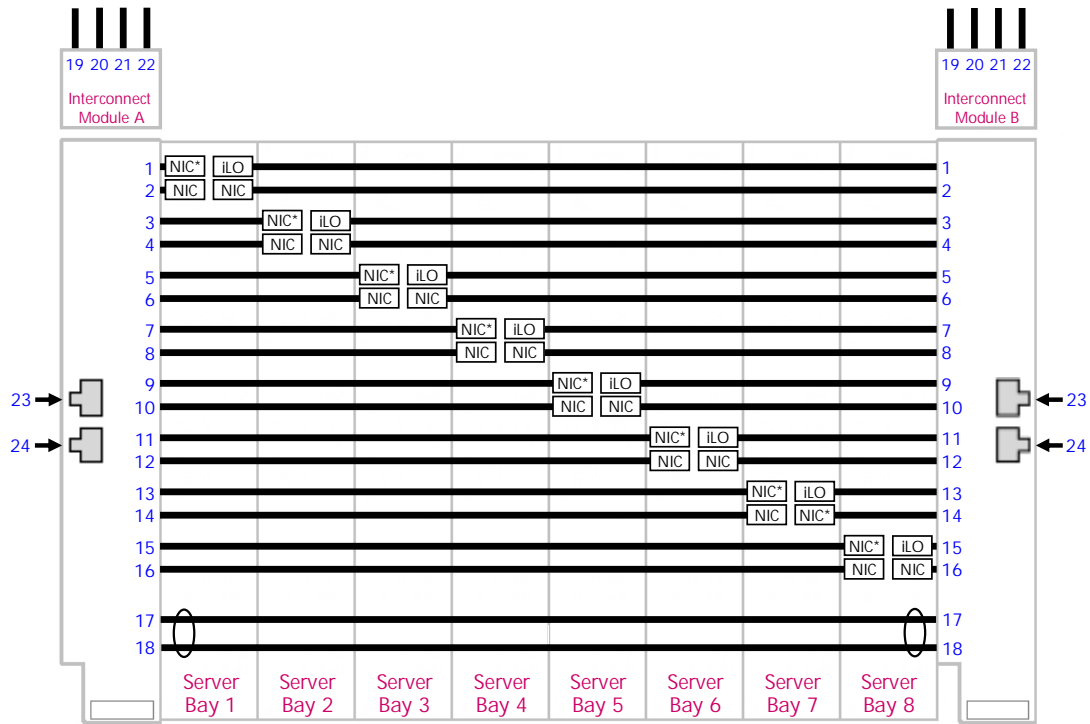
- To avoid compromising security, configure the interconnect switches and then physically connect them to the network infrastructure. Alternately, use the HP Diagnostic Station to configure the switches outside the rack environment prior to installation.
- Simplify GbE2 Interconnect Switch deployment. Given that most implementations will be similar between the two GbE2 Interconnect Switches within a ProLiant p-Class server blade enclosure, it is possible to configure one switch, download the existing configuration via TFTP or SCP, or capture the screen text from a /cfg/dump, and then modify the configuration for the other switch. When performing this action, consider the following for the switch the new configuration is being applied:
 1. Modify the VLAN settings as needed for items such as individual server VLAN requirements and the differences between the type of NIC port (iLO versus data).
 2. Verify spanning tree settings including STGs, especially if the configuration includes tagging (trunking) different VLANs between the GbE2 Interconnect switches.
 3. Change the management interface to avoid an IP conflict.

CAUTION: The Cutting and pasting of configuration settings can sometimes result in lost information due to the limited buffer size on some consoles. The recommended method is to use TFTP or SCP whenever possible.

Appendix A: GbE2 Interconnect Switch architecture

Two GbE2 Interconnect Switches are packaged into a GbE2 Interconnect Kit providing an end-to-end, fully redundant architecture that maximizes network availability. Redundant network adapters (NICs) are routed from each server blade bay to each hot-swappable interconnect switch (four NICs total per server bay) creating a fully meshed topology to the external Ethernet network (Figure 4). For more information about the network design within the server blade enclosure, see the [ProLiant BL p-Class Networking Overview](#)¹⁰ white paper.

Figure 4. ProLiant BL p-Class GbE2 Interconnect Switch architecture



* This is the default enabled PXE NIC. Using the ROM setup utility on the server, any other data NIC may be PXE enabled.

Port No.	Port Type	C-GbE2	F-GbE2
1 – 16	Downlink	10/100/1000	10/100/1000
17, 18	Crosslink	10/100/1000	10/100/1000
19, 20	Uplink	10/100/1000T	1000SX
21, 22	Uplink	10/100/1000T	1000SX
23, 24	Front Panel	10/100/1000T	10/100/1000T

The GbE2 Interconnect Switch supports all ProLiant BL p-Class server blades in mix and match combinations. Depending on the type and number of server blades installed in the server blade enclosure, not all GbE2 Interconnect Switch ports may be utilized. Also, the individual labeling of each NIC on the server blades is dependant on the server, the server operating system, and the NICs that are enabled on the server. To verify connectivity to a specific NIC from a GbE2 Interconnect Switch port, use the switch ping command:

```
GbE2>> ping <NIC IP address>
```

¹⁰ Available at http://wws1pro.compaq.com/support/reference_library/viewdocument.asp?source=5982-2202EN.xml&dt=21.

Appendix B: GbE2 Interconnect Switch default settings

This section provides the default settings for the GbE2 Interconnect Switch.

- Table B-1 contains general default settings for both switch A and switch B
- Table B-2 contains port name, VLAN, and trunking default settings for switch A
- Table B-3 contains port name, VLAN, and trunking default settings for switch B

Table B-1. Switch A and switch B general default settings

Setting	Value
Notice	None
Banner	None
User Names/Passwords	User Name: Password: user—Enabled user oper—Disabled None admin—Enabled (cannot be disabled) admin
BOOTP Service	Enabled
IP Address (if manual IP option is selected)	0.0.0.0
Subnet Mask (if manual IP option is selected)	0.0.0.0
Primary Default Gateway	0.0.0.0
Secondary Default Gateway	0.0.0.0
Primary DNS Server Address	0.0.0.0
Secondary DNS Server Address	0.0.0.0
Default Domain Name	None
Management Network/Mask	0.0.0.0/0.0.0.0
Switch Software Image on Next Boot	Image1
Switch Config File on Next Boot	Active
Display Hostname (sysName) in CLI Prompt	Disabled
Idle Timeout	5 minutes
Telnet Status	Enabled
Telnet Port	23
Web Status	Enabled
Web Port	80
Backpressure	Disabled
Port State	Enabled
Port Speed/Duplex	Auto
Flow Control	Off
STP	STG 1—Enabled with Default VLAN (VID=1) Port 1-16 (Server Ports) STP—Disabled at Port Level STG 2-16—Disabled
Bridge Max Age	20 seconds
Bridge Hello Time	2 seconds
Bridge Forward Delay	15 seconds
Bridge Priority	32768
MAC Address Aging Time	300 seconds
Port Priority	128
Path Cost	19 for ports 1-16 and 19-24 4 for ports 17-18
Static VLAN Entry	Default VLAN (VID = 1)
Port VID	1 for all ports
Default VLAN	Default VLAN (VID=1) with all ports assigned including CPU, STG=1

Table B-1. Switch A and switch B general default settings continued

Setting	Value
Port Trunking	Trunk Group 1, Enabled with Port 17 and 18
Port Trunking Load Sharing Algorithm	Src Mac Address
Port Mirroring—Mirror Status	Disabled
Port Mirroring—Mirror Port	None Selected
Port Mirroring—Mirror Port Traffic Direction	None Selected
Port Mirroring—Monitoring Port	None Selected
SNMP	Disabled
SNMP System Name	None
SNMP System Location	None
SNMP System Contact	None
SNMP Community String/Access Right	Public = read-only Private = read/write
SNMP Trap Host 1	0.0.0.0
SNMP Trap Host 1 Community String	Public
SNMP Trap Host 2	0.0.0.0
SNMP Trap Host 2 Community String	Public
SNMP Authentication Traps	Disabled
SNMP Link Up/Down Traps	Enabled
Security IP Network/Mask	0.0.0.0/0.0.0.0
TFTP Server IP Address	0.0.0.0
TFTP Port Number	69
Firmware Upgrade	File name = none
Configuration File from TFTP Server	File name = none
Configuration File to TFTP Server	File name = none
PING Tool	Target address = undefined Default tries = 4
TraceRoute Tools	Target address = undefined
Serial Port Baud Rate	9600
Serial Port Data Bit	8
Serial Port Parity Bit	None
Serial Port Stop Bit	1
Serial Port Flow Control	None
NTP State	Disabled
NTP Server	0.0.0.0
NTP Resync Interval	720 minutes
GMT Timezone Offset	-06:00
Daylight Savings Time State	Disabled
System Up Time	0 days 00 :00 :00
Current Time	RTC or NTP (00 :00 :00)
Date	None
Syslog Host	0.0.0.0
Syslog Host 2	0.0.0.0
Syslog Host Severity	7
Syslog Host 2 Severity	7
Syslog Console Output	Disabled

Table B-1. Switch A and switch B general default settings continued

Setting	Value
Log	console—Enabled system—Enabled mgmt—Enabled cli—Enabled stp—Enabled vlan—Enabled ssh—Enabled ntp—Enabled ip—Enabled web—Enabled
RSA Server Key Autogen Interval	0
RSA Server Key Autogen	Disabled
SSH Server	On
SCP-only Administrator Password	admin
SSH Server Port	22
SCP Apply and Save	Enabled
RADIUS Server	Off
RADIUS Secret	None
Primary RADIUS Server	0.0.0.0
Secondary RADIUS Server	0.0.0.0
RADIUS Server Port	1645
RADIUS Server Retries	3
RADIUS Server Timeout	3
RADIUS Backdoor for Telnet Access	Disabled
Re-ARP Period in Minutes	10

Table B-2. Switch A port name, VLAN, and trunking default settings

Note: The default port names are provided as one example and may or may not be completely applicable to the server configuration being deployed. It is recommended that the port names be modified as necessary to reflect the server configuration.

Port Type	Port No.	Speed	VID	VLAN Membership	VLAN Name	Port Name	STP	Trunk Group
Server	1	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server1_Port1	Disabled	
Server	2	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server1_Port2	Disabled	
Server	3	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server2_Port1	Disabled	
Server	4	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server2_Port2	Disabled	
Server	5	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server3_Port1	Disabled	
Server	6	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server3_Port2	Disabled	
Server	7	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server4_Port1	Disabled	
Server	8	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server4_Port2	Disabled	
Server	9	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server5_Port1	Disabled	
Server	10	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server5_Port2	Disabled	
Server	11	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server6_Port1	Disabled	
Server	12	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server6_Port2	Disabled	
Server	13	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server7_Port1	Disabled	
Server	14	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server7_Port2	Disabled	
Server	15	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server8_Port1	Disabled	
Server	16	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server8_Port2	Disabled	
Crosslink	17	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Xconnect_1	Enabled	1
Crosslink	18	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Xconnect_2	Enabled	1
Uplink	19	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	U1_Port_19	Enabled	
Uplink	20	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	U1_Port_20	Enabled	
Uplink	21	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	U2_Port_21	Enabled	
Uplink	22	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	U2_Port_22	Enabled	
Front Panel	23	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	FrontPanel1	Enabled	
Front Panel	24	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	FrontPanel2	Enabled	

Table B-3. Switch B port name, VLAN, and trunking default settings

Note: The default port names are provided as one example and may or may not be completely applicable to the server configuration being deployed. It is recommended that the port names be modified as necessary to reflect the server configuration.

Port Type	Port No.	Speed	VID	VLAN Membership	VLAN Name	Port Name	STP	Trunk Group
Server	1	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server1_iLO	Disabled	
Server	2	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server1_Port3	Disabled	
Server	3	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server2_iLO	Disabled	
Server	4	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server2_Port3	Disabled	
Server	5	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server3_iLO	Disabled	
Server	6	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server3_Port3	Disabled	
Server	7	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server4_iLO	Disabled	
Server	8	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server4_Port3	Disabled	
Server	9	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server5_iLO	Disabled	
Server	10	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server5_Port3	Disabled	
Server	11	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server6_iLO	Disabled	
Server	12	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server6_Port3	Disabled	
Server	13	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server7_iLO	Disabled	
Server	14	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server7_Port3	Disabled	
Server	15	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server8_iLO	Disabled	
Server	16	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Server8_Port3	Disabled	
X-Connect	17	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	XConnect_1	Enabled	1
X-Connect	18	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	XConnect_2	Enabled	1
M Uplink	19	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	U1_Port_19	Enabled	
M Uplink	20	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	U1_Port_20	Enabled	
D Uplink	21	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	U2_Port_21	Enabled	
D Uplink	22	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	U2_Port_22	Enabled	
Front Panel	23	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	FrontPanel1	Enabled	
Front Panel	24	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	FrontPanel2	Enabled	

For more information

For additional information, refer to the resources detailed below.

Resource description	Web address
ProLiant BL p-Class GbE2 Interconnect Switch home page	http://h18004.www1.hp.com/products/servers/proliant-bl/p-class/bl-p-interconnect-switch2.html
<i>ProLiant BL p-Class GbE2 Interconnect Switch Application Guide</i>	http://wws1pro.compaq.com/support/reference_library/viewdocument.asp?source=331403-001.xml&dt=264
<i>ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide</i>	http://wws1pro.compaq.com/support/reference_library/viewdocument.asp?source=331404-001%20.xml&dt=264
<i>ProLiant BL p-Class GbE2 Interconnect Switch Web-based Interface Reference Guide</i>	http://h71025.www7.hp.com/support/reference_library/viewdocument.asp?countrycode=1000&prodid=5726&source=331401-001.xml&dt=264&docid=20470
<i>ProLiant BL p-Class GbE2 Interconnect Switch User Guide</i>	http://h71025.www7.hp.com/support/reference_library/viewdocument.asp?countrycode=1000&prodid=5726&source=331399-001%20.xml&dt=264&docid=20467
<i>ProLiant BL p-Class GbE2 Interconnect Switch Overview white paper</i>	http://wws1pro.compaq.com/support/reference_library/viewdocument.asp?source=5982-2175EN.xml&dt=21
<i>ProLiant BL p-Class Networking Overview white paper</i>	http://wws1pro.compaq.com/support/reference_library/viewdocument.asp?source=5982-2202EN.xml&dt=21
<i>Using ProLiant Essentials Rapid Deployment Pack for scripted blade-based switch configuration white paper</i>	http://h71025.www7.hp.com/support/reference_library/viewdocument.asp?countrycode=1000&prodid=5728 ProLiant+BL+e:Class+System&source=TC030905HT.xml&dt=21&docid=20890
ProLiant network adapter teaming	http://www.compaq.com/products/servers/networking/teaming.html
<i>HP ProLiant BL System Best Practices Guide</i>	http://wws1pro.compaq.com/support/reference_library/viewdocument.asp?source=351359-001.xml&dt=264
<i>HP ProLiant BL System Common Procedures Guide</i>	http://wws1pro.compaq.com/support/reference_library/viewdocument.asp?source=351360-001.xml&dt=264

© 2004 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Catalyst, Cisco, Cisco Systems, EtherChannel, and IOS are registered trademarks of Cisco Systems, Inc. Notepad is a registered trademark of Microsoft Corporation. All other trademarks or registered trademarks mentioned in this document are the property of their respective owners.

5982-3160EN, 04/2004

