

Deploying the HP GbE2c Ethernet Blade Switch for HP c-Class BladeSystem 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.....	7
Multiple spanning tree groups.....	7
VLAN and STG configuration guidelines.....	8
Trunking.....	8
Load balancing.....	8
Trunking and spanning tree.....	9
Trunking configuration guidelines.....	9
Uplink Failure Detection.....	10
Spanning Tree Protocol with UFD.....	10
UFD configuration guidelines.....	10
Common topological examples.....	11
Topology 1: Fully meshed with BL480c Blade Server.....	13
VLAN configuration.....	13
Spanning tree configuration.....	14

Port configuration	15
Trunking and EtherChannel.....	15
Topology 2: Partial mesh	16
VLAN configuration	16
Spanning tree configuration.....	17
Port configuration	17
Trunking and EtherChannel.....	18
Topology 3: Straight-through	18
VLAN configuration	19
Spanning tree configuration.....	19
Port configuration	20
Trunking and EtherChannel.....	20
Uplink Failure Detection configuration.....	21
Topology summary	21
Securing the GbE2c Ethernet Blade Switch	22
Management interfaces	22
Command line interface	22
Browser based interface.....	22
Setting source IP address range.....	22
SNMP management.....	22
RADIUS.....	23
TACACS+	23
Passwords	23
Additional best practices	23
Appendix A: GbE2c Ethernet Blade Switch Architecture	25
Appendix B: GbE2c Ethernet Blade Switch default settings	27
For more information.....	32

Abstract

This HOWTO provides best practice guidelines and configuration examples for installation of the HP GbE2c Ethernet Blade Switch into a Cisco-based network. This guide is meant to be a tool to help direct decisions in planning, optimizing and securing the GbE2c Ethernet Blade 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 GbE2c Ethernet Blade 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 GbE2c Ethernet Blade Switch for HP c-Class BladeSystem. For readers not familiar with GbE2c Ethernet Blade Switch, please see the *HP GbE2c Ethernet Blade Switch for c-Class BladeSystem-overview* as well as the user documentation that shipped with the GbE2c Ethernet Blade Switch. To obtain these documents, go to the HP website (<http://www.hp.com/support>), and search for **GbE2c**.

Introduction

This HOWTO identifies best practice guidelines and configuration examples for installation of the HP GbE2c Ethernet Blade 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 GbE2c Ethernet Blade 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). For additional information on the GbE2c Ethernet Blade Switch, please see the *HP GbE2c Ethernet Blade Switch for c-Class BladeSystem-overview* on the HP website.

Terminology

The terminology that differs between the Cisco Catalyst 6509 switch and the GbE2c Ethernet Blade Switch documentation is identified in Table 1.

Table 1. Network terminology cross reference

HP GbE2c Ethernet Blade 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, trunking	EtherChannel, channeling
spanning tree protocol group (STG)	spanning tree instance
IEEE 802.1d, Spanning Tree Protocol	per VLAN Spanning Tree Plus (PVST+)

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/12/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.	<i>/cfg/12/vlan <vlan number></i>

To distinguish between HP GbE2c Ethernet Blade Switch and Catalyst 6509 commands, each command will be preceded by a `GbE2c>>` and `6509#`, respectively.

Critical features for successful deployment

Understanding VLANs and VLAN tagging (VLAN trunking), spanning tree protocol, and trunking (channeling) is critical to the successful deployment of the GbE2c Ethernet Blade Switch. Each of these topics is covered providing a high-level primer inclusive of GbE2c Ethernet Blade 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 GbE2c Ethernet Blade Switch Application Guide* 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 GbE2c Ethernet Blade Switches and VLAN identification (VLANID) on Cisco Catalyst switches. The GbE2c Ethernet Blade Switch allows any PVID value from 2 to 4095 with PVID 1 reserved as the default VLAN. The default GbE2c Ethernet Blade Switch configuration has all ports assigned to PVID 1.

The IEEE industry standard for VLANs is 802.1Q. Each GbE2c Ethernet Blade Switch supports 255 port-based IEEE 802.1Q VLANs. The GbE2c Ethernet Blade Switch VLAN menu can be found under:

```
GbE2c>> /cfg/12/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 communicate across the network.

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.

GbE2c Ethernet Blade Switch ports may be individually configured as tagged or untagged using the following command:

```
GbE2c>> /cfg/port <port number>/tag ena
```

When implementing VLAN tagging on the GbE2c Ethernet Blade Switch, the PVID values must be established correctly between devices communicating in the VLAN. This option is found under:

```
GbE2c>> /cfg/port <port number>/pvid <PVID number>
```

IP management interface

The IP management interface provides management access to the GbE2c Ethernet Blade 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 GbE2c Ethernet Blade Switch to ensure remote communication to the switch remains possible. In order to access the GbE2c Ethernet Blade 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 GbE2c Ethernet Blade Switch management features are effectively disabled. To avoid these situations, it is suggested that all ports used for remote GbE2c Ethernet Blade Switch management remain on the default VLAN and that an IP management interface be assigned to the default VLAN.

On the GbE2c Ethernet Blade Switch, assign the IP management interface to a VLAN using the commands:

```
GbE2c>> /cfg/l3/if <number>/mask <mask>/addr <address>  
GbE2c>> /cfg/l3/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 GbE2c Ethernet Blade 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+.

NOTE: The GbE2c Ethernet Blade Switch does not support Cisco's Per VLAN Spanning Tree (PVST), only Per VLAN Spanning Tree Plus (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 GbE2c Ethernet Blade Switch using the following command:

```
GbE2c>> /cfg/l2/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. By default, the GbE2c switch assigns fixed costs to all ports regardless of the link speed (default costs are 4 in PVST+ mode and 20000 in RSTP/MSTP modes). If the path cost is set to 0, the cost is set to auto. The spanning tree protocol assigns lower values to high-bandwidth ports, such as Gigabit Ethernet, to encourage their use. In PVST+ mode, the path cost is automatically changed to 4 with 1 Gbps link, to 19 with a 100 Mbps link, and to 100 with a 10 Mbps link. In MSTP mode the path cost is automatically changed to 20000/200000/2000000. The path cost may also be set manually on the GbE2c Ethernet Blade Switch using the following command:

```
GbE2c>> /cfg/l2/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 GbE2c Ethernet Blade Switch using the following command:

```
GbE2c>> /cfg/l2/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 GbE2c Ethernet Blade Switch integrates into a PVST+ environment through the use of STGs. In the GbE2c implementation, similar to the Cisco implementation, an administrator creates a new VLAN on the GbE2c switch, and then a spanning tree instance (i.e. STG) is automatically assigned to it.

Note: When creating a new VLAN on GbE2c switch, it is automatically assigned to STP 1 in PVST+ /RSTP modes or to CIST in MSTP mode.

The PVST+ interoperability feature on the GbE2c Ethernet Blade 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 GbE2c Ethernet Blade Switch.
- The default STG 1 can hold multiple VLANs; all other STGs (groups 2–16) can hold one VLAN.

On each GbE2c Ethernet Blade Switch, the five external ports (ports 20–24) 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:

```
GbE2c>> /cfg/l2/stp <stg number>/port <port number>
```

VLAN and STG configuration guidelines

When creating a VLAN on the GbE2c Ethernet Blade 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 tagged ports can belong to multiple STGs; if ports are untagged, they can belong to only one 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.
- 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 any port in a trunk is set to forwarding (STP), the remaining ports in the trunk will also be set to forwarding.

Trunking

Trunking, 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 trunking is IEEE 802.3ad. Cisco has developed a similar trunking method known as EtherChannel. The GbE2c Ethernet Blade Switch supports twelve IEEE 802.3ad (without LACP¹) trunks per switch interoperable with EtherChannel. Each trunk may contain one to five ports, providing a 10-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 GbE2c Ethernet Blade 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 MAC addresses. For traffic that contains IP addresses, the GbE2c Ethernet Blade 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.

¹ Link aggregation control protocol (LACP) is an enhancement over EtherChannel and other static trunking methods. LACP dynamically learns about the link status and makes decisions on which links to use for load balancing and failback in case of link failure. As a result, IEEE 802.3ad with LACP is often called dynamic trunking.

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. Trunking can be used to provide redundant links while ensuring that STP does not block this redundancy. Within a trunk, all the individual ports are seen as one logical by the spanning tree protocol.

Trunking configuration guidelines

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

- Confirm the GbE2c Ethernet Blade Switch ports to be trunked are set to enable.
- 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 GbE2c Ethernet Blade Switch port is configured in a trunk, the port becomes a trunk member using the following trunk command:

```
GbE2c>> /cfg/l2/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.

Uplink Failure Detection

Uplink Failure Detection (UFD) is designed to provide High Availability in “straight-through” topologies. A straight through topology is one that does not provide any redundancy either through STP or Virtual Router Redundancy Protocol (VRRP). Uplink Failure detection is designed to work with Network Adapter Teaming on HP server blades.

For details about Network Adapter Teaming on HP ProLiant server blades, refer to the white paper at the following location: <http://h18004.www1.hp.com/products/servers/networking/whitepapers.html>

The main components of UFD are as follows:

- Uplinks (external ports)
- Downlinks (internal ports)
- Server network adapters (NICs)

When UFD is configured, it enables the switch to monitor uplink ports. Once the switch detects an uplink failure or state change to blocking, it automatically disables the corresponding downlink ports. The Network Adaptor Teaming driver detects that the downlink port has been disabled and triggers a network-adaptor failover to another port on the switch, or another switch in the chassis. This provides an alternate path for the traffic.

The switch automatically enables the disabled downlink port/s when the failed or blocked uplink returns to service.

You must first configure a Failure Detection Pair and then turn on UFD. A Failure Detection Pair consists of the following groups of ports:

- Link to Monitor (LtM)
The Link to Monitor group consists of one uplink port (20-24), or one trunk group that contains only uplink ports. The switch monitors the LtM for link failure.
- Link to Disable (LtD)
The Link to Disable group consists of one or more downlink ports (1-16) and trunk groups that contain only downlink ports. When the switch detects a link failure on the LtM, it automatically disables all ports in the LtD.

When the LtM returns to service, the switch automatically enables all ports in the LtD.

Spanning Tree Protocol with UFD

If Spanning Tree Protocol (STP) is enabled on ports in the LtM, then the switch monitors the STP state and the link status on ports in the LtM. The switch automatically disables the ports in the LtD when it detects a link failure or STP Blocking state.

When the switch determines that ports in the LtM are in STP Forwarding State, then it automatically enables the ports in the LtD, to fall back to normal operation.

UFD configuration guidelines

This section provides important information about configuring UFD:

- UFD is required only when uplink-path redundancy is not available on the blade switches.
- Only one Failure Detection pair (one group of Links to Monitor and one group of Links to Disable) is supported on each switch (all VLANs and Spanning Tree Groups).
- An LtM can be either one uplink port or one trunk group of uplink ports.
- Ports that are already members of a trunk group are not allowed to be assigned to an LtM.
- An uplink port cannot be added to a trunk group if it already belongs to an LtM.

- A trunk group configured as an LtM can contain multiple uplink ports (20-24), but no downlink ports (1-16) or crossconnect ports (17-18).
- A trunk group configured as an LtD can contain multiple downlink ports (1-16), but no uplink ports (20-24) or crossconnect ports (17-18).
- Ports that are already members of a trunk group are not allowed to be assigned to an LtD
- A downlink port cannot be added to a trunk group if it already belongs to an LtD.
- A trunk cannot be disabled once it is added into LtM or LtD.
- Ports cannot be added to a trunk already assigned to an active LtM , unless they are of the proper type (uplink ports).
- Ports cannot be added to a trunk already assigned to an active LtD , unless they are of the proper type (downlink ports).

Common topological examples

Three common topological configurations are provided that highlight the main requirements for a successful configuration of the GbE2c Ethernet Blade 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, 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 GbE2c Ethernet Blade Switch in any specific environment. This information is intended to supplement the user documentation included with the GbE2c Ethernet Blade Switch and the Catalyst 6509 switches.

The configuration steps assume the GbE2c Ethernet Blade Switch configuration is set to factory default. For a list of GbE2c Ethernet Blade Switch default settings, see the *HP GbE2c Ethernet Blade Switch User Guide*. To set the GbE2c Ethernet Blade Switch configuration settings to the factory default, use the following procedure:

1. Select the configuration block per the command:

```
GbE2c>> /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 GbE2c Ethernet Blade Switch must be reset per the command:

```
GbE2c>> /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 GbE2c Ethernet Blade Switch for c-Class BladeSystem Command Reference Guide*.

NOTE: Resetting the GbE2c Ethernet Blade Switch causes the spanning tree protocol to restart. This process can be lengthy depending on the topology of the network.

All topologies assume a c-Class server blade enclosure with enhanced backplane components, eight c-Class blade servers, and two GbE2c Ethernet Blade Switches operating at layer 2. The two GbE2c Ethernet Blade 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 requires use of the default VLAN (VLAN 1) for 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 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 four of the five exterior Gigabit Ethernet uplink ports (uplinks ports 20-23), but in a different design configuration. Each topology could be modified by connecting the 5th uplink (port 24) to the network for increased bandwidth, added availability, creation of a remote port diagnostic network, etc.

Table 3: Topology overview

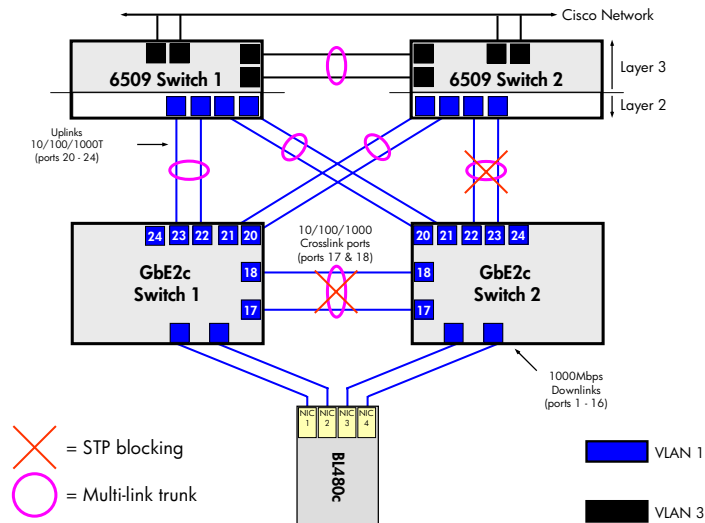
Topology	Benefits	Drawbacks
1	<ul style="list-style-type: none"> • Fully meshed design optimizing availability: <ul style="list-style-type: none"> – Maximum resiliency even with a dual link or dual non-like switch failure. – Loss of up to eight physical links without an interruption in service. 	<ul 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 GbE2c Ethernet Blade 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. 	<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 GbE2c Ethernet Blade 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 GbE2c Ethernet Blade Switches and the Cisco network. • Spanning tree protocol is not required. • Topology simplification decreasing support requirements. • From any single GbE2c Ethernet Blade Switch uplink, communicate to all c-Class BladeSystem 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 with BL480c Blade Server

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 GbE2c Ethernet Blade 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 with BL480c Blade Server



VLAN configuration

Per Figure 1, this configuration utilizes the default VLAN (VLAN 1) for GbE2c server ports, and uplink ports connecting the GbE2c switches to the 6509 switches. Any remaining Catalyst 6509 ports must be separate from VLAN 1, 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 trunking protocol mode:

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

2. On both Cisco 6509 switches, configure VLAN per the command:

```
6509# set vlan 1
```

3. On both Cisco 6509 switches, enable 802.1Q tagging (VLAN trunking) on all ports connected to both GbE2c Ethernet Blade Switches per the command:

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

4. On both GbE2c Ethernet Blade Switches (GbE2c Switch 1 and GbE2c Switch 2), enable the two crosslink ports (17-18).

```
GbE2c>> /cfg/port <number> ena
```

5. On both GbE2c Ethernet Blade Switches (GbE2c Switch 1 and GbE2c Switch 2), enable tagging (VLAN encapsulation) on the four uplink ports used in this configuration (20-23) and the two crosslink ports (17-18) per the command:

```
GbE2c>> /cfg/port <number>/tag ena
```

Spanning tree configuration

To eliminate the physical loops in this topology, STP is required on the ports connecting the GbE2c Ethernet Blade Switches and the Cisco Catalyst 6509 switches and on the GbE2c Ethernet Blade 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 1, set the bridge priority to 0 for VLAN 1, per the command:

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

2. On 6509 Switch 2, set the bridge priority to a slightly higher value than 6509 Switch 1 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
```

NOTE: Do not alter the default bridge priority for either GbE2c Ethernet Blade Switch. This will ensure that one of the Catalyst 6509 switches always becomes the root bridge.

3. Modify the port cost to 100 on GbE2c Switch 2, ports 17, 18, 21 and 22 using the command:

```
GbE2c>> /cfg/l2/stp 2/port <port number>/cost 100
```

This ensures that STP will behave in a predictable manner by blocking the GbE2c Switch 2 crosslink ports 17 and 18 and the uplink ports 21 and 22, and by placing all uplinks on GbE2c Switch 1 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 GbE2c Ethernet Blade Switches per the commands:

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

2. On both GbE2c Ethernet Blade Switches, enable ports 17 and 18 per the command:

```
GbE2c>> /cfg/port <port_number>/enable
GbE2c>> apply
GbE2c>> save
```

3. Configure the port speed and negotiation for GbE2c uplink ports 20-23 (on both GbE2c Ethernet Blade Switches), per the commands:

```
GbE2c>> /cfg/port <port number>/gig/auto on
GbE2c>> / cfg/port <port number>/gig/mode full
```

NOTE: This step is only necessary if the GbE2c Ethernet Blade Switch default configuration has been modified.

3. Utilizing Figure 1 as a reference, connect:

- GbE2c Switch 1 ports 21 and 22 to 6509 Switch 1
- GbE2c Switch 1 ports 20 and 21 to 6509 Switch 2.
- GbE2c Switch 2 ports 22 and 23 to 6509 Switch 2
- GbE2c Switch 2 ports 20 and 21 to 6509 Switch 1.

NOTE: Ports 17 and 18 on each GbE2c Ethernet Blade Switch are already connected across the server blade enclosure backplane and do not require any further physical connectivity.

Trunking and EtherChannel

This topology requires the creation of two trunks on each GbE2c Ethernet Blade 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 GbE2c Ethernet Blade Switches. Other trunks may also be present, as shown for VLAN 3 in Figure 1.

NOTE: For this topology, GbE2c uplink ports 20 and 21 represent trunk 2, and ports 22 and 23 represent trunk 3.

1. On both Cisco Catalyst 6509 switches, configure EtherChannel on the ports connected to both GbE2c Ethernet Blade Switches per the command:

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

2. Configure trunk 2 on each GbE2c Ethernet Blade Switch per the commands:

```
GbE2c>> /cfg/l2/trunk 2
GbE2c>> ena
GbE2c>> add 20
GbE2c>> add 21
```

3. On each GbE2c Ethernet Blade Switch repeat the above steps, but for trunk 3 using ports 22 and 23.

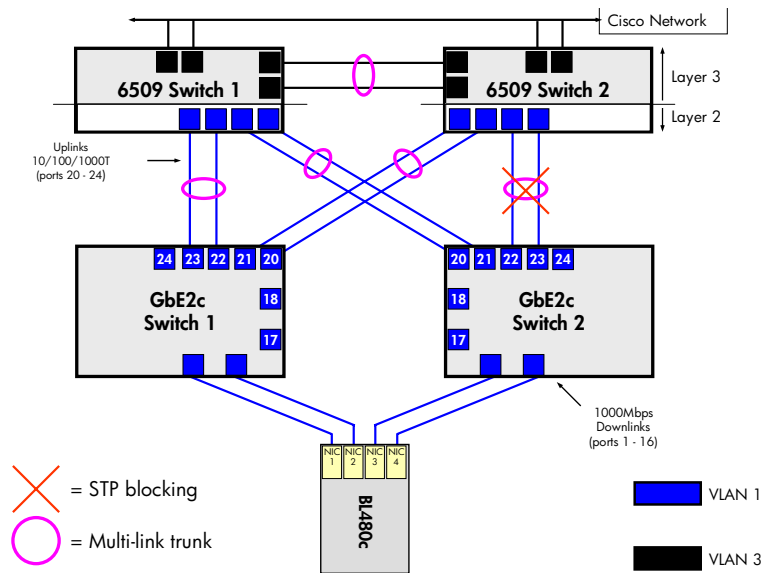
NOTE: It should not be necessary to configure trunk 1 as it is part of the GbE2c Ethernet Blade Switch default configuration. However, if the default configuration has been modified, configure trunk 1 using ports 17 and 18.

Topology 2: Partial mesh

Topology 2 is very similar in design to topology 1 except the GbE2c Ethernet Blade Switch crosslink ports are left in their default condition of disabled (Figure 2). This configuration maintains a high level of availability to ensure the loss of multiple physical links without an interruption in service. Leaving the crosslink ports or switch-to-switch links between the GbE2c Ethernet Blade Switches disabled by default 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 with BL480c Blade Server



VLAN configuration

Like topology 1, this configuration utilizes the default VLAN for GbE2c server ports and the uplink ports connecting the GbE2c switches to the 6509 switches. Any remaining Catalyst 6509 ports must be separate from VLAN 1 and are collectively represented in Figure 3 as VLAN 3.

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

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 VLANs 1 per the command:

```
6509# set vlan 1
```


3. On both Cisco switches, enable 802.1Q tagging (VLAN trunking) on all ports connected to both of the GbE2c Ethernet Blade Switches per the command:

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

4. On both GbE2c Ethernet Blade Switches (GbE2c Switch 1 and GbE2c Switch 2), enable tagging (VLAN encapsulation) on the four uplink ports used in this configuration (20-23) per the command:

```
GbE2c>> /cfg/port <port number>/tag ena
```

Spanning tree configuration

To eliminate the physical loops in this topology, STP is required on the ports connecting the GbE2c Ethernet Blade 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 1, set the bridge priority to 0 for VLAN1 per the command:

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

2. On 6509 Switch 2, set the bridge priority to a slightly higher value than 6509 Switch 1 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
```

NOTE: Do not alter the default bridge priority for either GbE2c Ethernet Blade Switch. This will ensure that one of the Catalyst 6509 switches always becomes the root bridge.

4. Modify the port cost to 100 on GbE2c Switch 2, ports 21 and 22 using the command:

```
GbE2c>> /cfg/l2/stp 2/port <port number>/cost 100
```

This ensures that STP will behave in a predictable manner by blocking the GbE2c Switch 2 uplink ports 22 and 23, and by placing all uplinks on GbE2c Switch 1 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 GbE2c Ethernet Blade 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 GbE2c uplink ports 20-23 (on both GbE2c Ethernet Blade Switches), per the commands:

```
GbE2c>> /cfg/port <port number>/gig/auto on  
GbE2c>> / cfg/port <port number>/gig/mode full
```

NOTE: This step is only necessary if the GbE2c Ethernet Blade Switch default configuration has been modified.

3. Utilizing Figure 2 as a reference, connect:
 - GbE2c Switch 1 ports 22 and 23 to 6509 Switch 1
 - GbE2c Switch 1 ports 20 and 21 to 6509 Switch 2.
 - GbE2c Switch 2 ports 22 and 23 to 6509 Switch 2
 - GbE2c Switch 2 ports 20 and 21 to 6509 Switch 1.

Trunking and EtherChannel

This topology requires the creation of two trunks on each GbE2c Ethernet Blade Switch and three EtherChannel groups on each Catalyst 6509 switch. Other trunks may also be present, as shown for VLAN 3 in Figure 3.

NOTE: For this topology, GbE2c uplink ports 20 and 21 represent trunk 2, and ports 22 and 23 represent trunk 3.

1. On both Cisco Catalyst 6509 switches, configure EtherChannel on the ports connected to both GbE2c Ethernet Blade Switches per the command:

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

2. Configure trunk 2 on each GbE2c Ethernet Blade Switch per the commands:

```
GbE2c>> /cfg/l2/trunk 2
GbE2c>> ena
GbE2c>> add 21
GbE2c>> add 22
```

3. On each GbE2c Ethernet Blade Switch repeat the above steps, but for trunk 3 using ports 22 and 23.

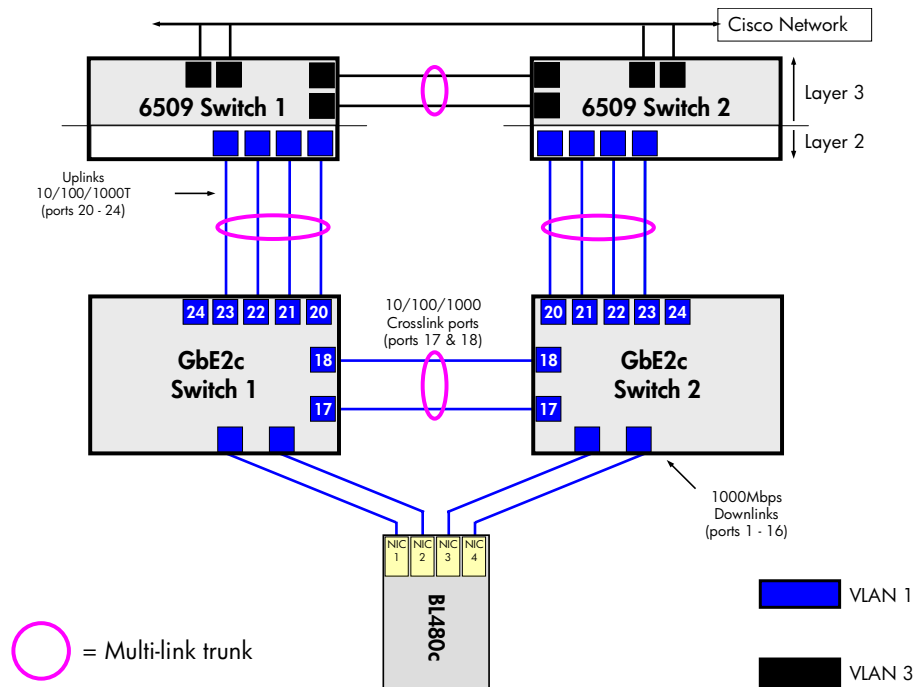
Topology 3: Straight-through

Topology 3 is a “straight-through” design providing a simplified architecture with maximum throughput (Figure 4). 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 GbE2c Ethernet Blade Switch crosslinks are enabled in this configuration. These Gigabit Ethernet links permit management of both switches and access to all c-Class BladeSystem server NICs from any single GbE2c Ethernet Blade Switch uplink. The crosslinks may be left disabled in this configuration, but it is not advised. In this case, a failure of any one switch (whether GbE2c Ethernet Blade 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 with BL480cBlade Server



VLAN configuration

Consistent with the other two topologies, this configuration utilizes the default VLAN for the GbE2c server ports and the uplink ports connecting the GbE2c switches to the 6509 switches. Any remaining Catalyst 6509 ports must be separate from VLAN 1 and are collectively represented in Figure 3 as VLAN 3.

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

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 VLAN per the command:

```
6509# set vlan 1
```

3. On both Cisco switches, enable 802.1Q tagging (VLAN trunking) on all ports connected to both of the GbE2c Ethernet Blade Switches per the command:

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

4. On both GbE2c Ethernet Blade Switches (GbE2c Switch 1 and GbE2c Switch 2), enable tagging (VLAN encapsulation) on the four uplink ports used in this configuration (20-23) per the command:

```
GbE2c>> /cfg/port <port number>/tag ena
```

Spanning tree configuration

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

1. On both GbE2c Interconnect Switches, disable spanning tree, per the commands:

```
GbE2c>> /cfg/12/stp 1/off
GbE2c>> save
```

CAUTION: If the GbE2c Ethernet Blade 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 GbE2c Ethernet Blade 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 GbE2c uplink ports 20-23 (on both GbE2c Ethernet Blade Switches), per the commands:

```
GbE2c>> /cfg/port <port number>/gig/auto on
GbE2c>> / cfg/port <port number>/gig/mode full
```

NOTE: This step is only necessary if the GbE2c Ethernet Blade Switch default configuration has been modified.

3. On both GbE2c Ethernet Blade Switches, enable ports 17 and 18 per the command:

```
GbE2c>> /cfg/port <port_number>/enable
GbE2c>> apply
GbE2c>> save
```

4. Utilizing Figure 4 as a reference, connect:

- GbE2c Switch 1 ports 20, 21, 22 and 23 to 6509 Switch 1
- GbE2c Switch 2 ports 20, 21, 22 and 23 to 6509 Switch 2

Trunking and EtherChannel

This topology requires the creation of one trunk on each GbE2c Ethernet Blade Switch and two EtherChannel groups on each the Catalyst 6509 switches. Additionally, a default preconfigured trunk (trunk 1) exists for the crosslink ports between the GbE2c Ethernet Blade Switches. Other trunks may also be present, as shown for VLAN 3 in Figure 4.

NOTE: For this topology, GbE2c uplink ports 20-23 represent trunk 2.

1. On both Cisco Catalyst 6509 switches, configure EtherChannel on the ports connected to both GbE2c Ethernet Blade Switches per the command:

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

2. Configure trunk 2 on each GbE2c Ethernet Blade Switch per the commands:

```
GbE2c>> /cfg/12/trunk 2
GbE2c>> ena
GbE2c>> add 20
GbE2c>> add 21
GbE2c>> add 22
GbE2c>> add 23
```

Uplink Failure Detection configuration

Topology 3 supports Uplink Failure Detection (UFD), but the crosslink ports (17-18) must be left in their default state, disabled.

1. On both GbE2c Ethernet Blade Switches assign the trunk group to be monitored for communication failure:

```
GbE2c>> /cfg/ufd/fdp ena
GbE2c>> ltm
GbE2c>> addtrnk 2
```

2. Disable crosslink ports (17-18) on both GbE2c Ethernet Blade Switch 1 and B (By default, ports 17 and 18 are disabled. Perform this step if you have enabled these ports for another configuration).

```
GbE2c>> /cfg/port 17 dis
GbE2c>> /cfg/port 18 dis
```

3. Assign downlink ports (1-16) to disable when an uplink failure occurs:

```
GbE2c>> /cfg/ufd/fdp/ltd
GbE2c>> addport 1
GbE2c>> addport 2
```

NOTE: Add only port 1 on GbE2c Switch 2.

4. Turn UFD on:

```
GbE2c>> /cfg/ufd/on
GbE2c>> apply
GbE2c>> save
```

When a link failure or Spanning Tree blocking occurs on trunk group 2, Switch 1 disables port 1 and port 2.

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
GbE2c 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 / trunking	Yes	Yes	Yes
Manage both GbE2c Ethernet Blade Switches from any of its uplink ports	No	No	Yes
Communicate to all NICs from any GbE2c Ethernet Blade Switch uplink port	No	No	Yes

* May be increased by an additional 2 Gbps full duplex by utilizing Gigabit Ethernet port 24 on the front of each GbE2c Ethernet Blade Switch

Securing the GbE2c Ethernet Blade Switch

HP recommends a variety of best practices to ensure the security of the network is maintained when deploying GbE2c Ethernet Blade Switches. The suggestions provided here are applicable to the GbE2c Ethernet Blade Switch independent of specific vendor used for the network infrastructure components.

Management interfaces

The GbE2c Ethernet Blade 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 GbE2c Ethernet Blade 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:

```
GbE2c>> /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:

```
GbE2c>> /cfg/sys/idle <idle time in minutes>
```

Browser based interface

The GbE2c Ethernet Blade 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:

```
GbE2c>> /cfg/sys/ access/http/wport <TCP connection port number>
```

Setting source IP address range

To limit management access to the GbE2c Ethernet Blade Switch without having to configure filters for each switch port, HP recommends the source IP address range be configured using the commands:

```
GbE2c>> /cfg/sys/access/mnet <management network, such as 192.192.192.0>  
GbE2c>> /cfg/sys/access/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 GbE2c Ethernet Blade Switch IP interface with a source IP address outside the range of 192.192.192.1 to 192.192.192.127.

SNMP management

The GbE2c Ethernet Blade Switch software provides simple network management protocol (SNMP) v3.0 support for access through network management software, such as HP OpenView and HP Systems Insight Manager. For improved security, HP recommends the default read and write community strings (public and private, respectively) be changed using the commands:

```
GbE2c>> /cfg/snmp/rcomm <SNMP read community string>  
GbE2c>> /cfg/snmp/wcomm <SNMP write community string>
```

RADIUS

The GbE2c Ethernet Blade 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 GbE2c Ethernet Blade Switch does not natively possess. For configuration procedures, refer to the *HP GbE2c Ethernet Blade Switch for c-Class BladeSystem Application Guide* chapter 1 and the *HP GbE2c Ethernet Blade Switch for c-Class BladeSystem Command Reference Guide* chapter 6.

TACACS+

The GbE2c switch software supports authentication, authorization, and accounting with networks using the Cisco Systems TACACS+ protocol. The switch functions as the Network Access Server (NAS) by interacting with the remote client and initiating authentication and authorization sessions with the TACACS+ access server. The remote user is defined as someone requiring management access to the switch either through a data or management port.

- TACACS+ offers the following advantages over RADIUS:
- TACACS+ uses TCP-based connection-oriented transport; whereas RADIUS is UDP based. TCP offers a connection-oriented transport, while UDP offers best-effort delivery. RADIUS requires additional programmable variables such as re-transmit attempts and time-outs to compensate for best-effort transport, but it lacks the level of built-in support that a TCP transport offers.
- TACACS+ offers full packet encryption whereas RADIUS offers password-only encryption in authentication requests.
- TACACS+ separates authentication, authorization, and accounting.

For configuration procedures, refer to the *HP GbE2c Ethernet Blade Switch for c-Class BladeSystem Application Guide* chapter 1 and the *HP GbE2c Ethernet Blade Switch for c-Class BladeSystem Command Reference Guide* chapter 6

Passwords

HP recommends all GbE2c Ethernet Blade 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 GbE2c Ethernet Blade Switch for c-Class BladeSystem Command Reference Guide* chapter 3.
- To change the SCP administrator password, see the *HP GbE2c Ethernet Blade Switch for c-Class BladeSystem Application Guide* chapter 1.

Additional best practices

Additional steps are recommended to ensure that the GbE2c Ethernet Blade Switch is easily serviceable and readily available within the network environment. The suggestions provided here are applicable to the GbE2c Ethernet Blade Switch independent of specific vendor used for the network infrastructure components.

- Record the GbE2c Ethernet Blade 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:

```
GbE2c >> /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 file using TFTP or SCP by entering one of the following commands:

```
GbE2c>> /cfg/ptcfg <TFTP server IP address> <config. file name>
```

```
GbE2c>> 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 GbE2c Ethernet Blade Switch includes two flash regions to store firmware. For more information, refer to the *HP GbE2c Ethernet Blade Switch for c-Class BladeSystem Command Reference Guide* chapter 8.
- Configure redundant syslog servers. For more information, refer to the *HP GbE2c Ethernet Blade Switch for c-Class BladeSystem Command Reference Guide* chapter 6.
- 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 GbE2c Ethernet Blade Switch deployment. Given that most implementations will be similar between the two GbE2c Ethernet Blade Switches within a c-Class BladeSystem 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.
 2. Verify spanning tree settings including STGs, especially if the configuration includes tagging (trunking) different VLANs between the GbE2c Ethernet Blade Switches.
 3. Change all IP interfaces 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: GbE2c Ethernet Blade Switch Architecture

Two GbE2c Ethernet Blade Switches can be installed into an HP c-Class BladeSystem to provide 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 Ethernet Blade switch (four NICs total per server bay) creating a fully meshed topology to the external Ethernet network (Figure 4). The architecture diagram assumes that GbE2c switches 1 and 2 are inserted into bays 1 and 2, respectively, and that servers in server bays 1-8 are full height servers (BL480c).

Figure 4. HP GbE2c Ethernet Blade Switch Architecture



Port No.	Port Type	GbE2c
1 – 16	Downlink	1000
17, 18	Crosslink	1000
20 – 24	Uplink	10/100/1000T

The GbE2c Ethernet Blade Switch supports all c-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 GbE2c Ethernet Blade 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 GbE2c Ethernet Blade Switch port, use the switch ping command:

```
GbE2c>> ping <NIC IP address>
```

To get detailed information about the default settings for the GbE2c Ethernet Blade Switch, refer to the *HP GbE2c Ethernet Blade Switch User Guide*, available at <http://www.hp.com/support>.

Appendix B: GbE2c Ethernet Blade Switch default settings

This section provides the default settings for the GbE2c Ethernet Blade Switch.

Table B-1. Switch 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	Factory
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	Both (TX/RX)
STP	STG1—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	4 for ports 1-18, 20-24
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
Port Trunking	Trunk Group 1 enabled, with Port 17 and 18 disabled
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	Enabled
SNMP System Name	None
SNMP System Location	None
SNMP System Contact	None
SNMP Community String/Access Right	Public = read-only Private = read/write
SNMP Authentication Traps	Disabled
SNMP Uplink Failure Detection Trap	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 = 5
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	1440 minutes
GMT Timezone Offset	-08:00
Daylight Savings Time State	Disabled
System Up Time	0 days 00 :00 :00
Current Time	RTC or NTP (00 :00 :00)
Date	12/31/2069
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	Enabled
Log	console—Enabled system—Enabled mgmt—Enabled cli—Enabled stg—Enabled vlan—Enabled ssh—Enabled ntp—Enabled ip—Enabled web—Enabled rmon—Enabled ufd—Enabled cfg—Enabled
RSA Server Key Autogen Interval	0
RSA Server Key Autogen	Disabled
SSH Server	Off
SCP-only Administrator Password	admin
SSH Server Port	22

SCP Apply and Save	Disabled
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
TACACS+ Server	Off
TACACS+ Secret	None
Primary TACACS+ Server	0.0.0.0
Secondary TACACS+ Server	0.0.0.0
TACACS+ Server Port	49
TACACS+ Server Retires	3
TACACS+ Server Timeout	5
TACACS+ Backdoor for Telnet Access	Disabled
Re-ARP Period in Minutes	10

Table B-2 contains port name, VLAN, and trunking default settings for Switch 1

Table B-2. Switch 1 port name, VLAN, and trunking default settings

Port Type	Port No.	Speed	VID	VLAN Membership	VLAN Name	Port Name	STP	Trunk Group
Server	1	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink1	Disabled	
Server	2	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink2	Disabled	
Server	3	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink3	Disabled	
Server	4	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink4	Disabled	
Server	5	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink5	Disabled	
Server	6	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink6	Disabled	
Server	7	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink7	Disabled	
Server	8	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink8	Disabled	
Server	9	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink9	Disabled	
Server	10	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink10	Disabled	
Server	11	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink11	Disabled	
Server	12	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink12	Disabled	

Server	13	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink13	Disabled	
Server	14	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink14	Disabled	
Server	15	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink15	Disabled	
Server	16	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink16	Disabled	
X-Connect	17	1000 (Auto)	1	Egress/Untag	Default VLAN	Xconnect1	Enabled	1
X-Connect	18	1000 (Auto)	1	Egress/Untag	Default VLAN	Xconnect2	Enabled	1
Management	19	100 (Auto off)	1	Egress/Untag	VLAN 4095	4095 Mgmt	Enabled	
Uplink	20	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Uplink1	Enabled	
Uplink	21	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Uplink2	Enabled	
Uplink	22	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Uplink3	Enabled	
Uplink	23	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Uplink4	Enabled	
Uplink	24	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Uplink5	Enabled	

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.

Table B-3. Switch 2 port name, VLAN, and trunking default settings

Port Type	Port No.	Speed	VID	VLAN Membership	VLAN Name	Port Name	STP	Trunk Group
Server	1	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink1	Disabled	
Server	2	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink2	Disabled	
Server	3	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink3	Disabled	
Server	4	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink4	Disabled	
Server	5	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink5	Disabled	
Server	6	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink6	Disabled	
Server	7	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink7	Disabled	
Server	8	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink8	Disabled	

Server	9	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink9	Disabled	
Server	10	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink10	Disabled	
Server	11	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink11	Disabled	
Server	12	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink12	Disabled	
Server	13	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink13	Disabled	
Server	14	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink14	Disabled	
Server	15	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink15	Disabled	
*Downlink Server	16	1000 (Auto)	1	Egress/Untag	Default VLAN	Downlink16	Disabled	
X-Connect	17	1000 (Auto)	1	Egress/Untag	Default VLAN	Xconnect 1	Enabled	1
X-Connect	18	1000 (Auto)	1	Egress/Untag	Default VLAN	Xconnect2	Enabled	1
Management	19	100 (Auto off)	1	Egress/Untag	Default VLAN	4095 Mgmt	Enabled	
Uplink	20	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Uplink1	Enabled	
Uplink	21	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Uplink2	Enabled	
Uplink	22	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Uplink3	Enabled	
Uplink	23	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Uplink4	Enabled	
Uplink	24	10/100/1000 (Auto)	1	Egress/Untag	Default VLAN	Uplink5	Enabled	

For more information

For additional information, refer to the resources detailed below.

Resource description	Web address
<i>GbE2c Ethernet Blade Switch</i> homepage	http://h18000.www1.hp.com/products/blades/components/ethernet/GbE2c/index.html
<i>GbE2c Ethernet Blade Switch</i> documentation	http://h20000.www2.hp.com/bizsupport/TechSupport/DocumentIndex.jsp?contentType=SupportManual&lang=en&cc=us&docIndexId=179111&taskId=101&prodTypeId=329290&prodSeriesId=1845925
<i>GbE2c Ethernet Blade Switch</i> key benefits	http://h18000.www1.hp.com/products/blades/components/ethernet/GbE2c/benefits.html#4
<i>Network Adapter Teaming</i> white paper	http://h18000.www1.hp.com/products/servers/networking/teaming.html

© 2006 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.

4AA1-0779ENW

