HP ProLiant BL p-Class GbE2 Interconnect Switch Application Guide

August 2003 (First Edition)
Part Number 331403-001

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for HP products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

HP ProLiant BL p-Class GbE2 Interconnect Switch Application Guide

August 2003 (First Edition)
Part Number 331403-001

Contents

About This Guide

Chapter 1
Accessing the GbE2 Interconnect Switch

Chapter 2
Ports and Trunking

Chapter 3
VLANs

Chapter 4
Spanning Tree Protocol

Appendix A
Troubleshooting Tools

Index

List of Figures

List of Tables

About This Guide

Use this guide as reference when configuring and maintaining the HP ProLiant BL p-Class GbE2 Interconnect Switch.

WARNING: To reduce the risk of personal injury from electric shock and hazardous energy levels, only authorized service technicians should attempt to repair this equipment. Improper repairs can create conditions that are hazardous.

Technician Notes

WARNING: Only authorized technicians trained by HP should attempt to repair this equipment. All troubleshooting and repair procedures are detailed to allow only subassembly/module-level repair. Because of the complexity of the individual boards and subassemblies, no one should attempt to make repairs at the component level or to make modifications to any printed wiring board. Improper repairs can create a safety hazard.

WARNING: To reduce the risk of personal injury from electric shock and hazardous energy levels, do not exceed the level of repairs specified in these procedures. Because of the complexity of the individual boards and subassemblies, do not attempt to make repairs at the component level or to make modifications to any printed wiring board. Improper repairs can create conditions that are hazardous.

WARNING: To reduce the risk of electric shock or damage to the equipment:

- **Disconnect power from the system by unplugging all power cords from the power supplies.**
- **Do not disable the power cord grounding plug. The grounding plug is an important safety feature.**
- **Plug the power cord into a grounded (earthed) electrical outlet that is easily accessible at all times.**

CAUTION: To properly ventilate the system, you must provide at least 7.6 cm (3.0 in.) of clearance at the front and back of the switch.

CAUTION: The computer is designed to be electrically grounded (earthed). To ensure proper operation, plug the AC power cord into a properly grounded AC outlet only.

NOTE: Any indications of component replacement or printed wiring board modifications may void any warranty.

Where to Go for Additional Help

In addition to this guide, the following information sources are available:

- *HP ProLiant BL p-Class GbE2 Interconnect Switch User Guide*
- *HP ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide*
- *HP ProLiant BL p-Class GbE2 Interconnect Switch Browser-based Interface Reference Guide*
- *HP ProLiant BL p-Class C-GbE2 Interconnect Kit Quick Setup Instructions*
- *HP ProLiant BL p-Class F-GbE2 Interconnect Kit Quick Setup Instructions*
- Service training guides
- Service advisories and bulletins
- QuickFind information services
- Insight Manager software

Telephone Numbers

For the name of your nearest HP authorized reseller:

- In the United States, call 1-800-345-1518.
- In Canada, call 1-800-263-5868.

For HP technical support:

- In the United States and Canada, call 1-800-652-6672.
- Outside the United States and Canada, refer to

  www.hp.com

1

Accessing the GbE2 Interconnect Switch

Introduction

This guide will help you plan, implement, and administer the HP ProLiant BL p-Class GbE2 Interconnect Switch software. Where possible, each section provides feature overviews, usage examples, and configuration instructions.

- Chapter 1, "Accessing the GbE2 Interconnect Switch," describes how to configure and view information and statistics on the GbE2 Interconnect Switch over an IP network. This chapter also discusses different methods to manage the GbE2 Interconnect Switch for remote administrators, such as setting specific IP addresses and using Remote Authentication Dial-in User Service (RADIUS) authentication, Secure Shell (SSH), and Secure Copy (SCP) for secure access to the switch.
- Chapter 2, "Ports and Trunking," describes how to group multiple physical ports together to aggregate the bandwidth between large-scale network devices.
- Chapter 3, "VLANs," describes how to configure Virtual Local Area Networks (VLANs) for creating separate network segments, including how to use VLAN tagging for devices that use multiple VLANs.
- Chapter 4, "Spanning Tree Protocol," discusses how spanning trees configure the network so that the GbE2 Interconnect Switch uses the most efficient path when multiple paths exist.
- Appendix A, "Troubleshooting Tools," describes port mirroring and other troubleshooting techniques.

Additional References

Additional information about installing and configuring the GbE2 Interconnect Switch is available in the following guides, which are located on the ProLiant BL p-Class GbE2 Interconnect Switch Management Utilities and User Documentation CD:

- *HP ProLiant BL p-Class GbE2 Interconnect Switch User Guide*
- *HP ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide*
- *HP ProLiant BL p-Class GbE2 Interconnect Switch Browser-based Interface Reference Guide*

Typographical Conventions

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

Table 1-1: Typographic Conventions

Typeface or Symbol	Meaning	Example
`AaBbCc123`	This type depicts onscreen computer output and prompts.	`Main#`
`AaBbCc123`	This type displays in command examples and shows text that must be typed in exactly as shown.	`Main#` **`sys`**
`<AaBbCc123>`	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. This also shows book titles, special terms, or words to be emphasized.	To establish a Telnet session, enter: `host#` **`telnet`** ***`<IP address>`*** Read your *User's Guide* thoroughly.
`[ ]`	Command items shown inside brackets are optional and can be used or excluded as the situation demands. Do not type the brackets.	`host#` **`ls [-a]`**

Using the Command Line Interface

The command line interface (CLI) can be accessed via local terminal connection or a remote session using Telnet or SSH. The CLI is the most direct method for collecting GbE2 Interconnect Switch information and performing switch configuration. The **Main Menu** of the CLI with administrator privileges is displayed in the following table:

```
[Main Menu]
      info    - Information Menu
      stats   - Statistics Menu
      cfg     - Configuration Menu
      oper    - Operations Command Menu
      boot    - Boot Options Menu
      maint   - Maintenance Menu
      diff    - Show pending config changes  [global command]
      apply   - Apply pending config changes [global command]
      save    - Save updated config to FLASH [global command]
      revert  - Revert pending or applied changes [global command]
      exit    - Exit  [global command, always available]
```

Connecting through the Console Port

Using a null modem cable, you can directly connect to the GbE2 Interconnect Switch through the console port. A console connection is required in order to configure Telnet or other remote access applications. For more information on establishing console connectivity to the GbE2 Interconnect Switch, refer to the *HP ProLiant BL p-Class GbE2 Interconnect Switch User Guide*.

Configuring an IP Interface

An IP interface address must be set on the GbE2 Interconnect Switch to provide management access to the switch over an IP network. By default, the management interface is set up to request its IP address from a Bootstrap Protocol (BOOTP) server.

If you have a BOOTP server on your network, add the Media Access Control (MAC) address of the GbE2 Interconnect Switch to the BOOTP configuration file located on the BOOTP server. The MAC address can be found on a small white label on the back panel of the GbE2 Interconnect Switch. The MAC address can also be found in the System Information menu (Refer to the “System Information” section in Chapter 4 of the *HP ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide*.) If you are using a DHCP server that also does BOOTP, you do not have to configure the MAC address.

If you do not have a BOOT server, you must manually configure an IP address. The BOOTP client will be disabled if you manually configure an IP address.

The following example shows how to manually configure an IP address on the GbE2 Interconnect Switch:

1. Configure an IP interface for the Telnet connection, using the sample IP address of 205.21.17.3.

   The pending subnet mask address and broadcast address are automatically calculated.

```
>> # /cfg/ip/if 1                           (Select IP interface 1)
>> IP Interface 1# addr 205.21.17.3         (Assign IP address for the interface)
Current IP address:     0.0.0.0
New pending IP address: 205.21.17.3
Pending new subnet mask:        255.255.255.0
Pending new broadcast address:  205.21.17.255
. . . . . . . . . . . . .
>> IP Interface 1# ena                      (Enable IP interface 1)
```

2. If necessary, configure up to two default gateways.

   Configuring the default gateways allows the GbE2 Interconnect Switch to send outbound traffic to the routers.

```
>> IP Interface 5# ../gw 1                 (Select primary default gateway)
>> Default gateway 1# addr 205.21.17.1     (Assign IP address for primary router)
>> Default gateway 1# ena                  (Enable primary default gateway)
>> Default gateway 1# ../gw 2              (Select secondary default gateway)
>> Default gateway 2# addr 205.21.17.2     (Assign address for secondary router)
>> Default gateway 2# ena                  (Enable secondary default gateway)
```

3. Apply, verify, and save the configuration.

```
>> Default gateway 2# apply                (Apply the configuration)
>> Default gateway 2# diff flash           (Verify the configuration)
>> Default gateway 2# save                 (Save the configuration)
```

Connecting via Telnet

By default, Telnet is enabled on the GbE2 Interconnect Switch. Once the IP parameters are configured, you can access the CLI from any workstation connected to the network using a Telnet connection. Telnet access provides the same options for a user and an administrator as those available through the console port, minus certain commands. The GbE2 Interconnect Switch supports four concurrent Telnet connections.

To establish a Telnet connection with the GbE2 Interconnect Switch, run the Telnet program on your workstation and issue the **telnet** command, followed by the switch IP address:

```
telnet <switch IP address>
```

Connecting via Secure Shell

The Secure Shell (SSH) protocol enables you to securely log into another computer over a network to execute commands remotely. As a secure alternative to using Telnet to manage switch configuration, SSH ensures that all data sent over the network is encrypted and secure. For more information, refer to the “Secure Shell and Secure Copy” section later in this chapter. For additional information on the CLI, refer to the *HP ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide.*

Using the Browser-Based Interface

The browser-based interface (BBI) provides access to the common configuration, management, and operation features of the GbE2 Interconnect Switch through your Web browser.

To use the browser-based interface, you need

- Installed GbE2 Interconnect Switch with Internet Protocol (IP) address
- PC or workstation with network access to the GbE2 Interconnect Switch
- Frame-capable Web-browser software, such as:
    - Netscape Navigator 4.7x or higher
    - Microsoft® Internet Explorer 6.0x or higher
- JavaScript™ enabled in your Web browser

For more information, refer to the *HP ProLiant BL p-Class GbE2 Interconnect Switch Browser-based Interface Reference Guide*.

Using Simple Network Management Protocol

The GbE2 Interconnect Switch software provides Simple Network Management Protocol (SNMP) v1.0 support for access through any network management software, such as HP OpenView Network Node Manager and Insight Manager 7.

To configure SNMP on the GbE2 Interconnect Switch, use the **`/cfg/sys/snmp`** command, and choose snmp access as **disabled, read-only,** or **read-write [d/r/w].**

To access the SNMP agent on the GbE2 Interconnect Switch, the read and write community strings on the SNMP manager should be configured to match those on the GbE2 Interconnect Switch. The default read community string on the GbE2 Interconnect Switch is **public** and the default write community string is **private.**

The read and write community strings on the GbE2 Interconnect Switch can be changed using the following commands on the CLI:

```
>> /cfg/snmp/rcomm
```

and

```
>> /cfg/snmp/wcomm
```

The SNMP manager should be able to reach any one of the IP interfaces on the GbE2 Interconnect Switch to communicate to the management interface.

For the SNMP manager to receive the traps sent out by the SNMP agent on the GbE2 Interconnect Switch, the trap host on the switch should be configured with the following command:

```
>> /cfg/snmp/<trap1|trap2>
```

To configure the community string for the first trap host use the following command:

```
>> /cfg/snmp/t1comm
```

To configure the community string for the second trap host use the following command:

```
>> /cfg/snmp/t2comm
```

For more information on using SNMP, refer to the *HP ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide*.

Refer to the HP ProLiant BL p-Class GbE2 Interconnect Switch Management Utilities for a complete list of supported MIBs. The utilities package is located on the utilities and user documentation CD included in the interconnect kit and at the following website:

www.compaq.com/support/servers

Secure Access to the Switch

Secure switch management is needed for environments that perform significant management functions across the Internet. The following are some of the functions for secured management:

- Limiting management users to a specific IP address range. Refer to the “Setting Allowable Source IP Address Ranges” section, later in this chapter.
- Authentication and authorization of remote administrators. Refer to the “RADIUS Authentication and Authorization” section, later in this chapter.
- Encryption of management information exchanged between the remote administrator and the switch. Refer to the “Secure Shell and Secure Copy” section, later in this chapter.

Setting Allowable Source IP Address Ranges

To limit access to the GbE2 Interconnect Switch without having to configure filters for each switch port, you can set a source IP address (or range) that will be allowed to connect to the GbE2 Interconnect Switch IP interface through Telnet, SSH, SNMP, or the GbE2 Interconnect Switch browser-based interface (BBI).

When an IP packet reaches the application switch, the source IP address is checked against the range of addresses defined by the management network and mask (**mnet** and **mmask**). If the source IP address of the host or hosts is within this range, it is allowed to attempt to log in. Any packet addressed to a GbE2 Interconnect Switch IP interface with a source IP address outside this range is discarded.

Configuring an IP Address Range for the Management Network

Configure the management network IP address and mask from the **System Menu** in the CLI. For example:

```
>> Main# /cfg/sys/mnet 192.192.192.0
Current management network:     0.0.0.0
New pending management network: 192.192.192.0
>> System# mmask 255.255.255.128
Current management netmask:     0.0.0.0
New pending management netmask: 255.255.255.128
```

In this example, the management network is set to 192.192.192.0 and management mask is set to 255.255.255.128. This defines the following range of allowed IP addresses: 192.192.192.1 to 192.192.192.127.

The following source IP addresses are granted or not granted access to the GbE2 Interconnect Switch:

- A host with a source IP address of 192.192.192.21 falls within the defined range and would be allowed to access the GbE2 Interconnect Switch.
- A host with a source IP address of 192.192.192.192 falls outside the defined range and is not granted access. To make this source IP address valid, you would need to shift the host to an IP address within the valid range specified by the **mnet** and **mmask** or modify the **mnet** to be 192.192.192.128 and the **mmask** to be 255.255.255.128. This would put the 192.192.192.192 host within the valid range allowed by the **mnet** and **mmask** (192.192.192.128-255).

RADIUS Authentication and Authorization

The GbE2 Interconnect Switch supports the Remote Authentication Dial-in User Service (RADIUS) method to authenticate and authorize remote administrators for managing the GbE2 Interconnect Switch. This method is based on a client/server model. The Remote Access Server (RAS)—the switch—is a client to the back-end database server. A remote user (the remote administrator) interacts only with the RAS, not the back-end server and database.

RADIUS authentication consists of the following components:

- A protocol with a frame format that utilizes User Datagram Protocol (UDP) over IP, based on Request For Comments (RFC) 2138 and 2866
- A centralized server that stores all the user authorization information
- A client, in this case, the GbE2 Interconnect Switch

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. Transactions between the client and the RADIUS server are authenticated using a shared key that is not sent over the network. In addition, the remote administrator passwords are sent encrypted between the RADIUS client (the switch) and the back-end RADIUS server.

How RADIUS Authentication Works

RADIUS authentication works as follows:

1. A remote administrator connects to the GbE2 Interconnect Switch and provides the user name and password.
2. Using Authentication/Authorization protocol, the GbE2 Interconnect Switch sends the request to the authentication server.
3. The authentication server checks the request against the user ID database.
4. Using RADIUS protocol, the authentication server instructs the GbE2 Interconnect Switch to grant or deny administrative access.

Configuring RADIUS on the Switch

To configure RADIUS on the GbE2 Interconnect Switch, do the following:

1. Turn RADIUS authentication on, then configure the Primary and Secondary RADIUS servers. For example:

```
>> Main# /cfg/sys/radius                           (Select the RADIUS Server menu)
>> RADIUS Server# on                               (Turn RADIUS on)
Current status: OFF
New status:     ON
>> RADIUS Server# prisrv 10.10.1.1                (Enter primary server IP)
Current primary RADIUS server:     0.0.0.0
New pending primary RADIUS server: 10.10.1.1
>> RADIUS Server# secsrv 10.10.1.2                (Enter secondary server IP)
Current secondary RADIUS server:     0.0.0.0
New pending secondary RADIUS server: 10.10.1.2
```

2. Configure the RADIUS secret.

```
>> RADIUS Server# secret
Enter new RADIUS secret: <1-32 character secret>
```

CAUTION: If you configure the RADIUS secret using any method other than a direct console connection, the secret may be transmitted over the network as clear text.

3. If desired, you may change the default Transmission Control Protocol (TCP) port number used to listen to RADIUS.

   The well-known port for RADIUS is 1645.

```
>> RADIUS Server# port
Current RADIUS port: 1645
Enter new RADIUS port [1500-3000]: <port number>
```

4. Configure the number of retry attempts for contacting the RADIUS server, and the timeout period.

```
>> RADIUS Server# retries
Current RADIUS server retries: 3
Enter new RADIUS server retries [1-3]:  <server retries>
>> RADIUS Server# time
Current RADIUS server timeout: 3
Enter new RADIUS server timeout [1-10]: 10   (Enter the timeout period in minutes)
```

RADIUS Authentication Features

The GbE2 Interconnect Switch supports the following RADIUS authentication features:

- Supports RADIUS client on the GbE2 Interconnect Switch, based on the protocol definitions in RFC 2138 and RFC 2866.
- Allows RADIUS secret password up to 32 bytes.
- Supports secondary authentication server so that when the primary authentication server is unreachable, the GbE2 Interconnect Switch can send client authentication requests to the secondary authentication server. Use the **/cfg/sys/radius/cur** command to show the currently active RADIUS authentication server.
- Supports user-configurable RADIUS server retry and time-out values:

  — Time-out value = 1-10 seconds

  — Retries = 1-3

  The GbE2 Interconnect Switch will time out if it does not receive a response from the RADIUS server in one to three retries. The GbE2 Interconnect Switch will also automatically retry connecting to the RADIUS server before it declares the server down.
- Supports user-configurable RADIUS application port. The default is 1645/User Datagram Protocol (UDP)-based on RFC 2138. Port 1812 is also supported.
- Allows network administrator to define privileges for one or more specific users to access the GbE2 Interconnect Switch at the RADIUS user database.
- Supports SecurID if the RADIUS server can do an ACE/Server client proxy. The password is the PIN number, plus the token code of the SecurID card.

User Accounts for RADIUS Users

The user accounts listed in Table 1-2 can be defined in the RADIUS server dictionary file.

Table 1-2: User Access Levels

User Account	Description and Tasks Performed	Password
User	User interaction with the GbE2 Interconnect Switch is completely passive; nothing can be changed on the GbE2 Interconnect Switch. Users may display information that has no security or privacy implications, such as GbE2 Interconnect Switch statistics and current operational state information.	**`user`**
Operator	Operators can only effect temporary changes on the GbE2 Interconnect Switch. These changes will be lost when the GbE2 Interconnect Switch is rebooted/reset. Operators have access to the GbE2 Interconnect Switch management features used for daily switch operations. Because any changes an operator makes are undone by a reset of the GbE2 Interconnect Switch, operators cannot severely impact switch operation. By default, the operator account is disabled and has no password.	
Administrator	Administrators are the only ones that can make permanent changes to the switch configuration—changes that are persistent across a reboot/reset of the GbE2 Interconnect Switch. Administrators can access GbE2 Interconnect Switch functions to configure and troubleshoot problems on the switch level. Because administrators can also make temporary (operator-level) changes as well, they must be aware of the interactions between temporary and permanent changes.	**`admin`**

RADIUS Attributes for User Privileges

When the user logs in, the GbE2 Interconnect Switch authenticates his/her level of access by sending the RADIUS access request, that is, the client authentication request, to the RADIUS authentication server.

If the authentication server successfully authenticates the remote user, the GbE2 Interconnect Switch verifies the privileges of the remote user and authorizes the appropriate access. When both the primary and secondary authentication servers are not reachable, the administrator has an option to allow backdoor access via the console only or console and Telnet access. The default is **disable** for Telnet access and **enable** for console access.

All user privileges, other than those assigned to the administrator, must be defined in the RADIUS dictionary. RADIUS attribute 6, which is built into all RADIUS servers, defines the administrator. The file name of the dictionary is RADIUS vendor-dependent. The RADIUS attributes shown in the following table are defined for user privileges levels.

Table 1-3: Proprietary Attributes for RADIUS

User Name/Access	User-Service-Type	Value
User	Vendor-supplied	255
Operator	Vendor-supplied	252

Secure Shell and Secure Copy

Secure Shell (SSH) and Secure Copy (SCP) use secure tunnels to encrypt and secure messages between a remote administrator and the GbE2 Interconnect Switch. Telnet does not provide this level of security. The Telnet method of managing a GbE2 Interconnect Switch does not provide a secure connection.

SSH is a protocol that enables remote administrators to log securely into the GbE2 Interconnect Switch over a network to execute management commands. By default, SSH is disabled (off) on the GbE2 Interconnect Switch.

SCP is typically used to copy files securely from one machine to another. SCP uses SSH for encryption of data on the network. On a GbE2 Interconnect Switch, SCP is used to download and upload the switch configuration via secure channels. By default, SCP is disabled on the GbE2 Interconnect Switch.

The GbE2 Interconnect Switch implementation of SSH is based on version 1.5, and supports SSH clients from version 1.0 through version 1.5. The following SSH clients are supported:

- SSH 1.2.23 and SSH 1.2.27 for Linux (freeware)
- SecureCRT® 3.0.2 and SecureCRT 3.0.3 (VanDyke Technologies, Inc.)
- F-Secure® SSH 1.1 for Windows® (F-Secure Corporation)
- OpenSSH_3.4 for Linux (RH 8.1)
- PuTTY Release 0.51 (Simon Tatham) for Windows

Configuring SSH and SCP Features

Before you can use SSH commands, use the following commands to turn on SSH and SCP.

Enabling or Disabling SSH

To enable or disable the SSH feature, connect to the GbE2 Interconnect Switch CLI and enter the following commands:

```
>> # /cfg/sys/sshd/on                               (Turn SSH on)
Current status: OFF
New status: ON
>> # /cfg/sys/sshd/off                              (Turn SSH off)
Current status: OFF
New status: ON
```

Enabling or Disabling SCP Apply and Save

Enter the following commands from the GbE2 Interconnect Switch CLI to enable the SCP **putcfg_apply** and **putcfg_apply_save** commands:

```
>> # /cfg/sys/sshd/ena           (Enable SCP apply and save)
SSHD# apply                      (Apply the changes to start generating RSA host and server keys)
RSA host key generation starts
. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .
. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .
RSA host key generation completes (lasts 212549 ms)
RSA host key is being saved to Flash ROM, please don’t reboot the box
immediately.
RSA server key generation starts
. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .
RSA server key generation completes (lasts 75503 ms)
RSA server key is being saved to Flash ROM, please don’t reboot the box
immediately.
------------------------------------------------------------------------
Apply complete; don’t forget to “save” updated configuration.


>> # /cfg/sys/sshd/dis          (Disable SSH/SCP apply and save)
```

Configuring the SCP Administrator Password

To configure the **scpadm** (SCP administrator) password, first connect to the GbE2 Interconnect Switch via the RS-232 management console. For security reasons, the **scpadm** password may only be configured when connected directly to the GbE2 Interconnect Switch console.

To configure the password, enter the following command via the CLI. At factory default settings, the current SCP administrator password is **admin**.

```
>> # /cfg/sys/sshd/scpadm
Changing SCP-only Administrator password; validation required. . .
Enter current administrator password: <password>
Enter new SCP-only administrator password: <new password>
Re-enter new SCP-only administrator password: <new password>
New SCP-only administrator password accepted.
```

Using SSH and SCP Client Commands

The following shows the format for using some client commands. The examples below use 205.178.15.157 as the IP address of a sample GbE2 Interconnect Switch.

Logging in to the GbE2 Interconnect Switch

Enter the following command to log in to the GbE2 Interconnect Switch:

```
ssh <switch IP address>
```

For example**:**

```
>> # ssh 205.178.15.157
```

Downloading configuration from the GbE2 Interconnect Switch using SCP

Enter the following command to download the GbE2 Interconnect Switch configuration using SCP. You will be prompted for a password.

```
scp <switch IP address>:getcfg <local filename>
```

For example:

```
>> # scp 205.178.15.157:getcfg ad4.cfg
```

Uploading Configuration to the GbE2 Interconnect Switch using SCP

Enter the following command to upload configuration to the GbE2 Interconnect Switch. You will be prompted for a password.

```
scp <local filename> <switch IP address>:putcfg
```

For example:

```
>> # scp ad4.cfg 205.178.15.157:putcfg
```

Applying and Saving Configuration

Enter the **apply** and **save** commands after the command above (**scp ad4.cfg 205.178.15.157:putcfg**), or use the following commands. You will be prompted for a password.

```
>> # scp <local_filename><switch IP addr>:putcfg_apply
>> # scp <local_filename><switch IP addr>:putcfg_apply_save
```

For example:

```
>> # scp ad4.cfg 205.178.15.157:putcfg_apply
>> # scp ad4.cfg 205.178.15.157:putcfg_apply_save
```

Note the following:

- The **diff** command is automatically executed at the end of **putcfg** to notify the remote client of the difference between the new and the current configurations.
- **putcfg_apply** runs the **apply** command after the **putcfg** is done.
- **putcfg_apply_save** saves the new configuration to the flash after **putcfg_apply** is done.
- The **putcfg_apply** and **putcfg_apply_save** commands are provided because extra **apply** and **save** commands are usually required after a **putcfg.**

SSH and SCP Encryption of Management Messages

The following encryption and authentication methods are supported for SSH and SCP:

- Server Host Authentication—Client RSA authenticates the switch at the beginning of every connection
- Key Exchange—RSA
- Encryption—3DES-CBC, Data Encryption Standard (DES)
- User Authentication—Local password authentication, RADIUS, SecurID (via RADIUS, for SSH only; does not apply to SCP)

Generating RSA Host and Server Keys for SSH Access

To support the SSH server feature, two sets of RSA keys (host and server keys) are required. The host key is 1024 bits and is used to identify the GbE2 Interconnect Switch. The server key is 768 bits and is used to make it impossible to decipher a captured session by breaking into the GbE2 Interconnect Switch at a later time.

When the SSH server is first enabled and applied, the GbE2 Interconnect Switch automatically generates the RSA host and server keys and is stored in the flash memory.

To configure RSA host and server keys, first connect to the GbE2 Interconnect Switch console connection (commands are not available via Telnet connection), and enter the following commands to generate them manually:

```
>> # /cfg/sys/sshd/hkeygen                    (Generates the host key)
>> # /cfg/sys/sshd/skeygen                    (Generates the server key)
```

These two commands take effect immediately without the need of an **apply** command.

When the GbE2 Interconnect Switch reboots, it will retrieve the host and server keys from the flash memory. If these two keys are not available in the flash memory and if the SSH server feature is enabled, the GbE2 Interconnect Switch automatically generates them during the system reboot. This process may take several minutes to complete.

The GbE2 Interconnect Switch can also automatically regenerate the RSA server key. To set the interval of RSA server key autogeneration, use the following command:

```
>> # /cfg/sys/sshd/intrval <number of hours (0-24)>
```

A value of 0 denotes that RSA server key autogeneration is disabled. When greater than 0, the GbE2 Interconnect Switch will auto generate the RSA server key every specified interval; however, RSA server key generation is skipped if the GbE2 Interconnect Switch is busy doing other key or cipher generation when the timer expires.

The GbE2 Interconnect Switch will perform only one session of key/cipher generation at a time. Thus, an SSH/SCP client will not be able to log in if the GbE2 Interconnect Switch is performing key generation at that time, or if another client has logged in immediately prior. Also, key generation will fail if an SSH/SCP client is logging in at that time.

SSH/SCP Integration with Radius Authentication

SSH/SCP is integrated with RADIUS authentication. After the RADIUS server is enabled on the GbE2 Interconnect Switch, all subsequent SSH authentication requests will be redirected to the specified RADIUS servers for authentication. The redirection is transparent to the SSH clients.

SecurID Support

SSH/SCP can also work with SecurID, a token card-based authentication method. The use of SecurID requires the interactive mode during login, which is not provided by the SSH connection.

There is no SNMP or browser-based interface (BBI) support for SecurID, because the SecurID server, ACE, is a one-time password authentication and requires an interactive session.

Using SecurID with SSH

Using SecurID with SSH involves the following tasks:

1. To log in using SSH, enter a special username, "**ace**," to bypass the SSH authentication.
2. After an SSH connection is established, you are prompted to enter the username and password (the SecurID authentication is being performed now).
3. Provide your username and the token in your SecurID card as a regular Telnet user.

Using SecurID with SCP

Using SecurID with SCP can be accomplished in two ways:

- Using a RADIUS server to store an administrator password.

  You can configure a regular administrator with a fixed password in the RADIUS server if it can be supported. A regular administrator with a fixed password in the RADIUS server can perform both SSH and SCP with no additional authentication required.

- Using an SCP-only administrator password.

  Use the command, **/cfg/sys/sshd/scpadm** to bypass the checking of SecurID.

  An SCP-only administrator password is typically used when SecurID is used. For example, it can be used in an automation program (in which the tokens of SecurID are not available) to back up (download) the GbE2 Interconnect Switch configurations each day.

IMPORTANT: The SCP-only administrator password must be different from the regular administrator password. If the two passwords are the same, the administrator using that password will not be allowed to log in as an SSH user because the switch will recognize him as the SCP-only administrator. The GbE2 Interconnect Switch will only allow the administrator access to SCP commands.

2

Ports and Trunking

Introduction

The first part of this chapter describes the different types of ports used on the GbE2 Interconnect Switch and how they map to the servers in the HP blade server enclosure. This information is useful in understanding other applications described in this guide, from the context of the embedded switch/server environment.

For specific information on how to configure ports for speed, auto-negotiation, and duplex modes, refer to the port commands in the *HP ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide*.

The second part of this chapter provides configuration background and examples for trunking multiple ports together. Trunk groups can provide super-bandwidth, multi-link connections between GbE2 Interconnect Switches or other trunk-capable devices. A trunk group is a group of links that act together, combining their bandwidth to create a single, larger virtual link. The GbE2 Interconnect Switch provides trunking support for the six external ports, two inter-switch links, and 16 server ports.

Ports on the GbE2 Interconnect Switch

On the GbE2 Interconnect Switch, port information differs depending on whether the GbE2 Interconnect Switch resides in Slot A or Slot B. The following table describes the mapping of the Ethernet ports of the GbE2 Interconnect Switch to a specific back-end server or network uplink, as well as each LED and the corresponding port.

Table 2-1: Mapping Ethernet Switch Ports to Blade Server Ports

Port Number	Port Origin	LED	Switch A Port Alias	Switch B Port Alias
1	NIC1/iLO, Server 1	P1	Server1_Port1	Server1_iLO
2	NIC2/NIC3, Server 1	P2	Server1_Port2	Server1_Port3
3	NIC1/iLO, Server 2	P3	Server2_Port1	Server2_iLO
4	NIC2/NIC3, Server 2	P4	Server2_Port2	Server2_Port3
5	NIC1/iLO, Server 3	P5	Server3_Port1	Server3_iLO
6	NIC2/NIC3, Server 3	P6	Server3_Port2	Server3_Port3
7	NIC1/iLO, Server 4	P7	Server4_Port1	Server4_iLO
8	NIC2/NIC3, Server 4	P8	Server4_Port2	Server4_Port3
9	NIC1/iLO, Server 5	P9	Server5_Port1	Server5_iLO
10	NIC2/NIC3, Server 5	P10	Server5_Port2	Server5_Port3
11	NIC1/iLO, Server 6	P11	Server6_Port1	Server6_iLO
12	NIC2/NIC3, Server 6	P12	Server6_Port2	Server6_Port3
13	NIC1/iLO, Server 7	P13	Server7_Port1	Server7_iLO
14	NIC2/NIC3, Server 7	P14	Server7_Port2	Server7_Port3
15	NIC1/iLO, Server 8	P15	Server8_Port1	Server8_iLO
16	NIC2/NIC3, Server 8	P16	Server8_Port2	Server8_Port3
17	Switch Port 17	P17	XConnect_1	XConnect_1
18	Switch Port 18	P18	XConnect_2	XConnect_2
19	Interconnect Module Port 19	P19	U1_Port_19	U1_Port_19
20	Interconnect Module Port 20	P20	U1_Port_20	U1_Port_20
21	Interconnect Module Port 21	P21	U2_Port_21	U2_Port_21
22	Interconnect Module Port 22	P22	U2_Port_22	U2_Port_22
23	Front Panel Port 23	P23	FrontPanel1	FrontPanel1
24	Front Panel Port 24	P24	FrontPanel2	FrontPanel2

NOTE: Refer to the server Setup and Installation Guide for more information about Integrated Lights-Out (iLO) management.

NOTE: Refer to Appendix D of the *HP ProLiant BL p-Class GbE2 Interconnect Switch User Guide* for a more detailed table listing the switch default settings for port names, VLANs, and port trunking.

Conceptual View

The following figure shows a conceptual diagram of the GbE2 Interconnect Switches and server links. If you were to view the enclosure from the front, Switch A will be on your left and Switch B will be on your right.

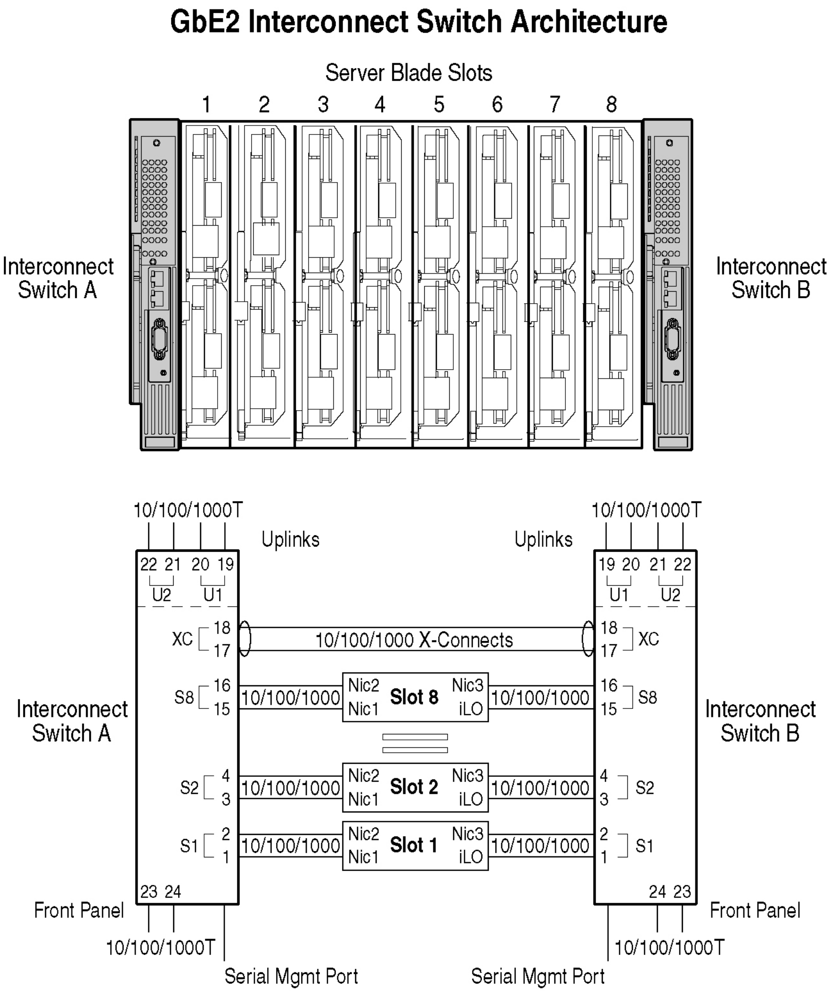

Figure 2-1: Conceptual view of GbE2 Interconnect Switch with QuadT2 Interconnect Module port-to-server mapping

Port Trunk Groups

When using port trunk groups between two GbE2 Interconnect Switches as shown in Figure 2-2, you can create an aggregate link operating at up to four Gigabits per second, depending on how many physical ports are combined. The GbE2 Interconnect Switch supports up to 12 trunk groups per switch, each with up to six ports per trunk group.

The trunking software detects misconfigured (or broken) trunk links and redirects traffic on the misconfigured or broken trunk link to other trunk members within that trunk. You can only use trunking if the port speeds are the same for each link.

Statistical Load Distribution

In a configured trunk group containing more than one port, the load distribution is determined by information embedded within the data frame. For traffic that does not contain IP information, the lowest port number in the trunk group is elected to be the designated port for forwarding traffic. For traffic that contains IP addresses, the GbE2 Interconnect Switch will calculate the designated trunk port to use by hashing the packet's source and destination IP addresses.

Built-In Fault Tolerance

Since each trunk group is composed of multiple physical links, the trunk group is inherently fault tolerant. As long as even one physical link between the switches is available, the trunk remains active.

Statistical load distribution is maintained whenever a link in a trunk group is lost or returned to service.

Before You Configure Trunks

When you create and enable a trunk, the trunk members (GbE2 Interconnect Switch ports) take on certain settings necessary for correct operation of the trunking feature.

Before you configure your trunk, you must consider these settings, along with specific configuration rules, as follows:

1. Read the configuration rules provided in the “Trunk Group Configuration Rules” section.
2. Determine which GbE2 Interconnect Switch ports (up to six) are to become trunk members (the specific ports making up the trunk).

   Ensure that the chosen GbE2 Interconnect Switch ports are set to enabled, using the **/cfg/port** command.

   Trunk member ports must have the same VLAN configuration.
3. Consider how the existing spanning tree will react to the new trunk configuration. Refer to Chapter 4, “Spanning Tree Protocol,” for spanning tree group configuration guidelines.
4. Consider how existing VLANs will be affected by the addition of a trunk.

Trunk Group Configuration Rules

The trunking feature operates according to specific configuration rules. When creating trunks, consider the following rules that determine how a trunk group reacts in any network topology:

- All trunks must originate from one device, and lead to one destination device. For example, you cannot combine a link from Server 1 and a link from Server 2 into one trunk group.
- Any physical switch port can belong to only one trunk group.
- Trunking from non-HP devices must comply with Cisco® EtherChannel® technology.
- All trunk member ports must be assigned to the same VLAN configuration before the trunk can be enabled.
- If you change the VLAN settings of any trunk member, you cannot apply the change until you change the VLAN settings of all trunk members.
- When an active port is configured in a trunk, the port becomes a trunk member when you enable the trunk using the **/cfg/trunk/ena** command. The spanning tree parameters for the port then change to reflect the new trunk settings.
- All trunk members must be in the same spanning tree group and can belong to only one spanning tree group. However if all ports are tagged, then all trunk ports can belong to multiple spanning tree groups.
- If you change the spanning tree participation of any trunk member to enabled or disabled, the spanning tree participation of all members of that trunk changes similarly.
- When a trunk is enabled, the trunk spanning tree participation setting takes precedence over that of any trunk member.
- You cannot configure a trunk member as a monitor port in a port mirroring configuration.
- Trunks cannot be monitored by a monitor port; however, trunk members can be monitored.

Port Trunking Example

In this example, the Gigabit uplink ports on each GbE2 Interconnect Switch, and the inter-switch link ports are configured into a total of five trunk groups: two on each GbE2 Interconnect Switch, and one trunk group at the inter-switch link between the two GbE2 Interconnect Switches. All ports operate at Gigabit Ethernet speed.

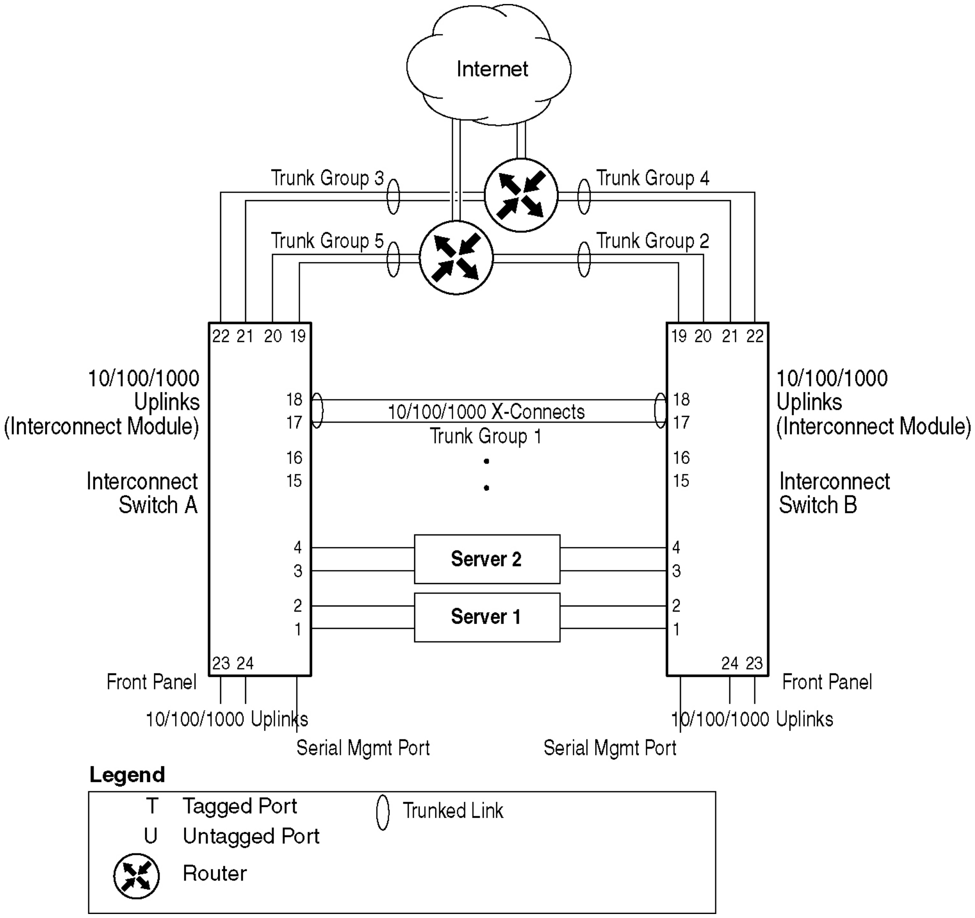

Figure 2-2: Port trunk group configuration example

The trunk groups are configured as follows:

- Trunk group 1 is configured by default on the inter-switch ports 17 and 18, which connect the GbE2 Interconnect Switches A and B together. Since this is the default configuration, you do not need to configure trunk group 1 on either GbE2 Interconnect Switch.
- Trunk groups 2-5 consist of two Gigabit uplink ports each, configured to act as a single link to the upstream routers. The trunk groups on each GbE2 Interconnect Switch are configured so that there is a link to each router for redundancy.

Prior to configuring each GbE2 Interconnect Switch in this example, you must connect to the appropriate switch CLI as the administrator. For details about accessing and using any of the commands described in this example, refer to the *HP ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide*.

1. On Switch A, configure trunk groups 5 and 3:

```
>> # /cfg/trunk 5                        (Select trunk group 5)
>> Trunk group 5# add 19                 (Add port 19 to trunk group 5)
>> Trunk group 5# add 20                 (Add port 20 to trunk group 5)
>> Trunk group 5# ena                    (Enable trunk group 5)
>> Trunk group 5# cur                    (View trunking configuration)
>> # /cfg/trunk 3                        (Select trunk group 3)
>> Trunk group 3# add 21                 (Add port 21 to trunk group 3)
>> Trunk group 3# add 22                 (Add port 22 to trunk group 3)
>> Trunk group 3# ena                    (Enable trunk group 3)
>> Trunk group 3# cur                    (View trunking configuration)
>> Trunk group 3# apply                  (Make your changes active)
>> Trunk group 3# save                   (Save for restore after reboot)
```

2. On Switch B, configure trunk groups 4 and 2:

```
>> # /cfg/trunk 4                        (Select trunk group 4)
>> Trunk group 4# add 21                 (Add port 21 to trunk group 4)
>> Trunk group 4# add 22                 (Add port 22 to trunk group 4)
>> Trunk group 4# ena                    (Enable trunk group 4)
>> Trunk group 4# cur                    (View trunking configuration)
>> # /cfg/trunk 2                        (Select trunk group 2 for inter-switch links)
>> Trunk group 2# add 19                 (Add port 19 to trunk group 2)
>> Trunk group 2# add 20                 (Add port 20 to trunk group 2)
>> Trunk group 2# ena                    (Make your changes active)
>> Trunk group 2# cur                    (View trunking configuration)
>> Trunk group 2# apply                  (Make your changes active)
>> Trunk group 2# save                   (Save for restore after reboot)
```

NOTE: In this example, two GbE2 Interconnect Switches are used. Any third-party device supporting link aggregation should be configured manually. Connection problems could arise when using automatic trunk group negotiation on the third-party device.

3. Examine the trunking information on each GbE2 Interconnect Switch using the following command:

```
>> /info/trunk                                    (View trunking information)
```

Information about each port in each configured trunk group will be displayed. Make sure that trunk groups consist of the expected ports and that each port is in the expected state.

3

VLANs

Introduction

This chapter describes network design and topology considerations for using Virtual Local Area Networks (VLANs). VLANs are commonly used to split up groups of network users into manageable broadcast domains, to create logical segmentation of workgroups, and to enforce security policies among logical segments.

The following topics are discussed in this chapter:

- VLANs and Port VLAN ID Numbers
- VLAN Tagging
- VLANs and IP Interfaces
- VLAN Topologies and Design Considerations

NOTE: Basic VLANs can be configured during initial switch configuration. Refer to the "Using the Setup Utility" section in the *HP ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide*.

More comprehensive VLAN configuration can be done from the command line interface. Refer to the "VLAN Configuration" and "Port Configuration" sections in the *HP ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide*.

Overview

Setting up VLANs is a way to segment networks to increase network flexibility without changing the physical network topology. With network segmentation, each GbE2 Interconnect Switch port connects to a segment that is a single broadcast domain. When a GbE2 Interconnect Switch port is configured to be a member of a VLAN, it is added to a group of ports (workgroup) that belongs to one broadcast domain.

Ports are grouped into broadcast domains by assigning them to the same VLAN. Frames received in one VLAN can only be forwarded within that VLAN, and multicast, broadcast, and unknown unicast frames are flooded only to ports in the same VLAN.

VLANs and Port VLAN ID Numbers

VLAN Numbers

The GbE2 Interconnect Switch supports up to 255 VLANs per switch. Even though the maximum number of VLANs supported at any given time is 255, each can be identified with any number between 1 and 4095. VLAN 1 is the default VLAN, and all ports are assigned to it.

Viewing VLANs

The VLAN information menu displays all configured VLANs and all member ports that have an active link state, for example:

```
>> Information# vlan
VLAN                Name                Status Ports
----  -------------------------------- ------ ----------------------
1     Default VLAN                       ena    1 4-22 24
2     VLAN 2                             ena    2 3
4092  VLAN 4092                          ena    23
```

PVID Numbers

Each port in the GbE2 Interconnect Switch has a configurable default VLAN number, known as its PVID. This places all ports on the same VLAN initially, although each port PVID is configurable to any VLAN number between 1 and 4095.

The default configuration settings for GbE2 Interconnect Switches have all ports set as untagged members of VLAN 1 with all ports configured as PVID = 1. In the default configuration example shown in Figure 3-1, all incoming packets are assigned to VLAN 1 by the default port VLAN identifier (PVID =1).

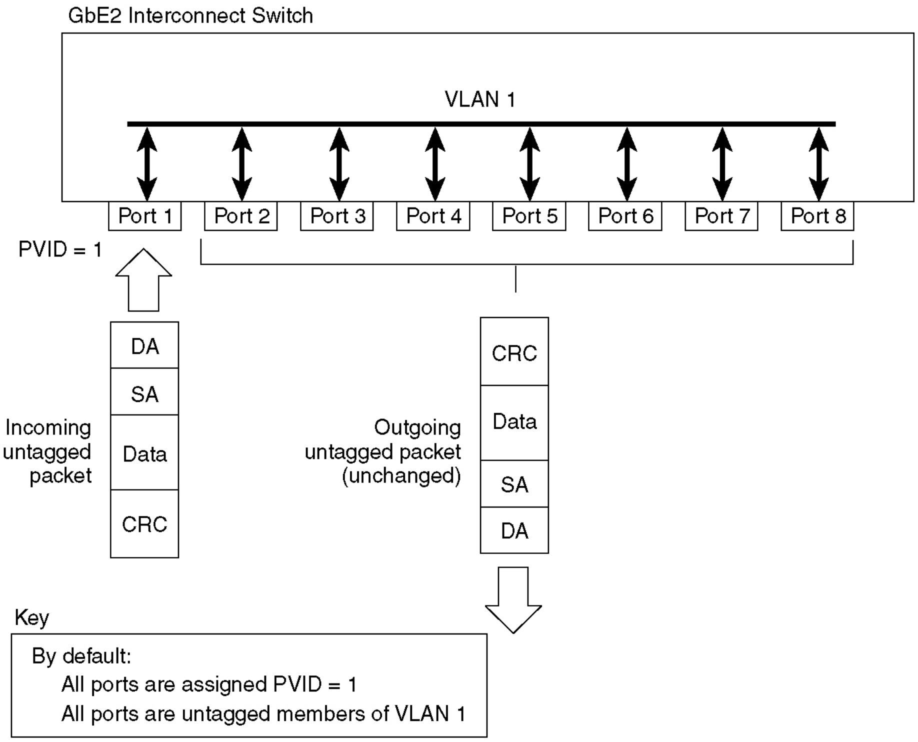

By default:
All ports are assigned PVID = 1
All ports are untagged members of VLAN 1

Figure 3-1: Default VLAN settings

Viewing and Configuring PVIDs

You can view PVIDs from the following CLI commands:

Port Information

```
>> /info/port
Port  Tag   PVID       NAME             VLAN(s)
----  ---   ----  --------------  ------------------------
  1    n      1  Server1_Port1       1
  2    n      1  Server1_Port2       1
  3    n      1  Server2_Port1       1
  4    n      1  Server2_Port2       1
  5    n      1  Server3_Port1       1
  6    n      1  Server3_Port2       1
  7    n      1  Server4_Port1       1
  :
  :
```

Port Configuration

```
>>  /cfg/port 22/pvid 22
Current port VLAN ID:     1
New pending port VLAN ID: 22

>> Port 22#
```

Each port on the GbE2 Interconnect Switch can belong to one or more VLANs, and each VLAN can have any number of switch ports in its membership. Any port that belongs to multiple VLANs, however, must have VLAN tagging enabled. Refer to the “VLAN Tagging” section.

Any untagged frames (those with no VLAN specified) are classified with the PVID of the sending port.

VLAN Tagging

The GbE2 Interconnect Switch supports IEEE 802.1Q VLAN tagging, providing standards-based VLAN support for Ethernet systems.

Tagging places the VLAN identifier in the frame header, allowing each port to belong to multiple VLANs. When you configure multiple VLANs on a port, you must also enable tagging on that port.

Since tagging fundamentally changes the format of frames transmitted on a tagged port, you must carefully plan network designs to prevent tagged frames from being transmitted to devices that do not support 802.1Q VLAN tags, or devices where tagging is not enabled.

Important terms used with the 802.1Q tagging feature are:

- VLAN identifier (VID)—the 12-bit portion of the VLAN tag in the frame header that identifies an explicit VLAN.
- Port VLAN identifier (PVID)—a classification mechanism that associates a port with a specific VLAN. For example, a port with a PVID of 3 (PVID =3) assigns all untagged frames received on this port to VLAN 3.
- Tagged frame—a frame that carries VLAN tagging information in the header. The VLAN tagging information is a 32-bit field (VLAN tag) in the frame header that identifies the frame as belonging to a specific VLAN. Untagged frames are marked (tagged) with this classification as they leave the GbE2 Interconnect Switch through a port that is configured as a tagged port.
- Untagged frame—a frame that does not carry any VLAN tagging information in the frame header.

- Untagged member—a port that has been configured as an untagged member of a specific VLAN. When an untagged frame exits the GbE2 Interconnect Switch through an untagged member port, the frame header remains unchanged. When a tagged frame exits the GbE2 Interconnect Switch through an untagged member port, the tag is stripped and the tagged frame is changed to an untagged frame.
- Tagged member—a port that has been configured as a tagged member of a specific VLAN. When an untagged frame exits the GbE2 Interconnect Switch through a tagged member port, the frame header is modified to include the 32-bit tag associated with the PVID. When a tagged frame exits the GbE2 Interconnect Switch through a tagged member port, the frame header remains unchanged (original VID remains).

NOTE: If an 802.1Q tagged frame is sent to a port that has VLAN-tagging disabled, then the frames are dropped at the ingress port.

NOTE: The port numbers specified in these illustrations may not directly correspond to the physical port configuration of your GbE2 Interconnect Switch model.

When you configure VLANs, you configure the GbE2 Interconnect Switch ports as tagged or untagged members of specific VLANs. Refer to Figure 3-2 through Figure 3-5.

In Figure 3-2, the untagged incoming packet is assigned directly to VLAN 2 (PVID = 2). Port 5 is configured as a tagged member of VLAN 2, and port 7 is configured as an untagged member of VLAN 2.

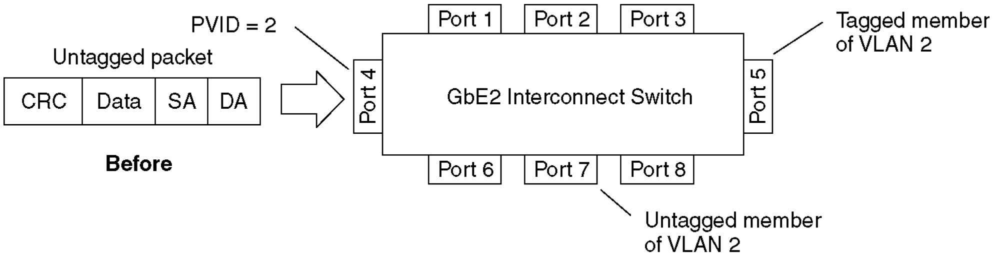

Figure 3-2: Port-based VLAN assignment

As shown in Figure 3-3, the untagged packet is marked (tagged) as it leaves the GbE2 Interconnect Switch through port 5, which is configured as a tagged member of VLAN 2. The untagged packet remains unchanged as it leaves the GbE2 Interconnect Switch through port 7, which is configured as an untagged member of VLAN 2.

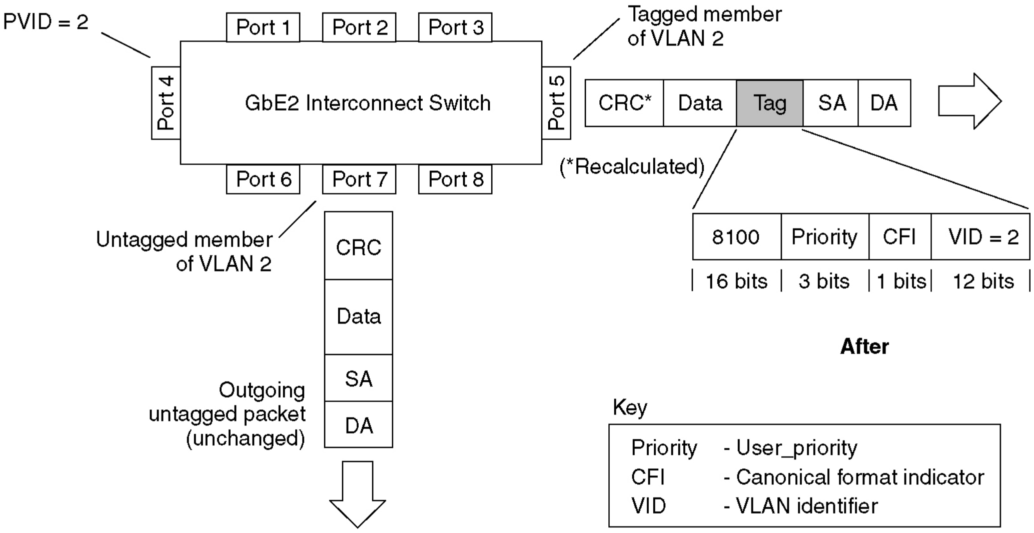

Figure 3-3: 802.1Q tagging after port-based VLAN assignment

In Figure 3-4, the tagged incoming packet is assigned directly to VLAN 2 because of the tag assignment in the packet. Port 5 is configured as a tagged member of VLAN 2, and port 7 is configured as an untagged member of VLAN 2.

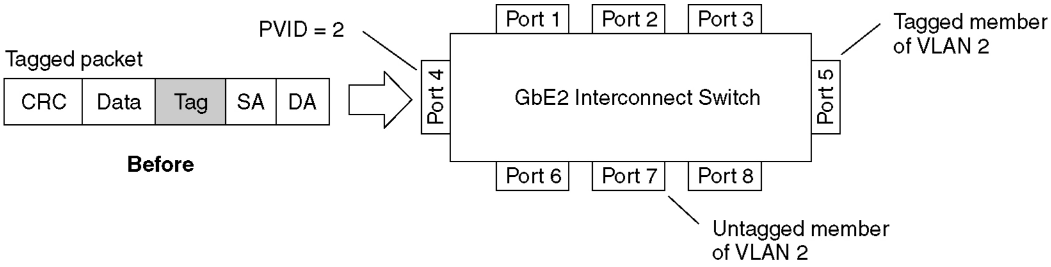

Figure 3-4: 802.1Q tag assignment

As shown in Figure 3-5, the tagged packet remains unchanged as it leaves the GbE2 Interconnect Switch through port 5, which is configured as a tagged member of VLAN 2. However, the tagged packet is stripped (untagged) as it leaves the GbE2 Interconnect Switch through port 7, which is configured as an untagged member of VLAN 2.

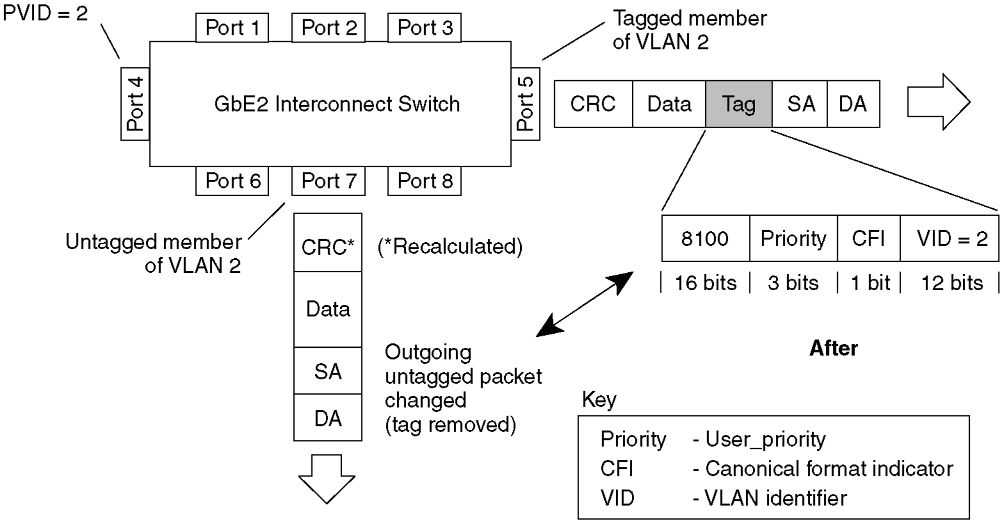

Figure 3-5: 802.1Q tagging after 802.1Q tag assignment

NOTE: Using the **/boot/conf factory** command resets all ports to VLAN 1 and all other settings to the factory defaults at the next reboot.

VLANs and IP Interfaces

Carefully consider how you create VLANs within the GbE2 Interconnect Switch, so that communication with the GbE2 Interconnect Switch remains possible. In order to access the GbE2 Interconnect Switch for remote configuration, trap messages, and other management functions, make sure that at least one IP interface on the GbE2 Interconnect Switch has a VLAN defined.

You can also inadvertently cut off access to management functions if you exclude the ports from the VLAN membership. For example, if all IP interfaces are left on VLAN 1 (the default), and all ports are configured for VLAN 2, then GbE2 Interconnect Switch management features are effectively cut off.

To remedy this, keep all ports used for remote GbE2 Interconnect Switch management on the default VLAN and assign an IP interface to the default VLAN.

For more information on configuring IP interfaces, refer to the “Configuring an IP Interface” section in Chapter 1.

VLAN Topologies and Design Considerations

By default, all GbE2 Interconnect Switch ports are configured to the default VLAN 1. This configuration groups all ports into the same broadcast domain. The VLAN has an 802.1Q VLAN ID of 1. VLAN tagging is turned off, because by default all ports are members of a single VLAN only.

If configuring Spanning Tree Protocol, note that each of spanning tree groups 2-16 may contain only one VLAN.

VLAN Configuration Rules

VLANs operate according to specific configuration rules which must be considered when creating VLANs. For example, HP recommends that all ports involved in trunking and port mirroring have the same VLAN configuration. If a port is on a trunk with a mirroring port, the VLAN configuration cannot be changed.

- For more information on port trunking, refer to the “Port Trunking Example” section in Chapter 2.
- For more information on configuring port mirroring, refer to the “Port Mirroring” section in Appendix A.

Multiple VLANS with Tagging

The following figure shows only those switch port to server links that must be configured for the example. While not shown, all other server links remain set at their default settings.

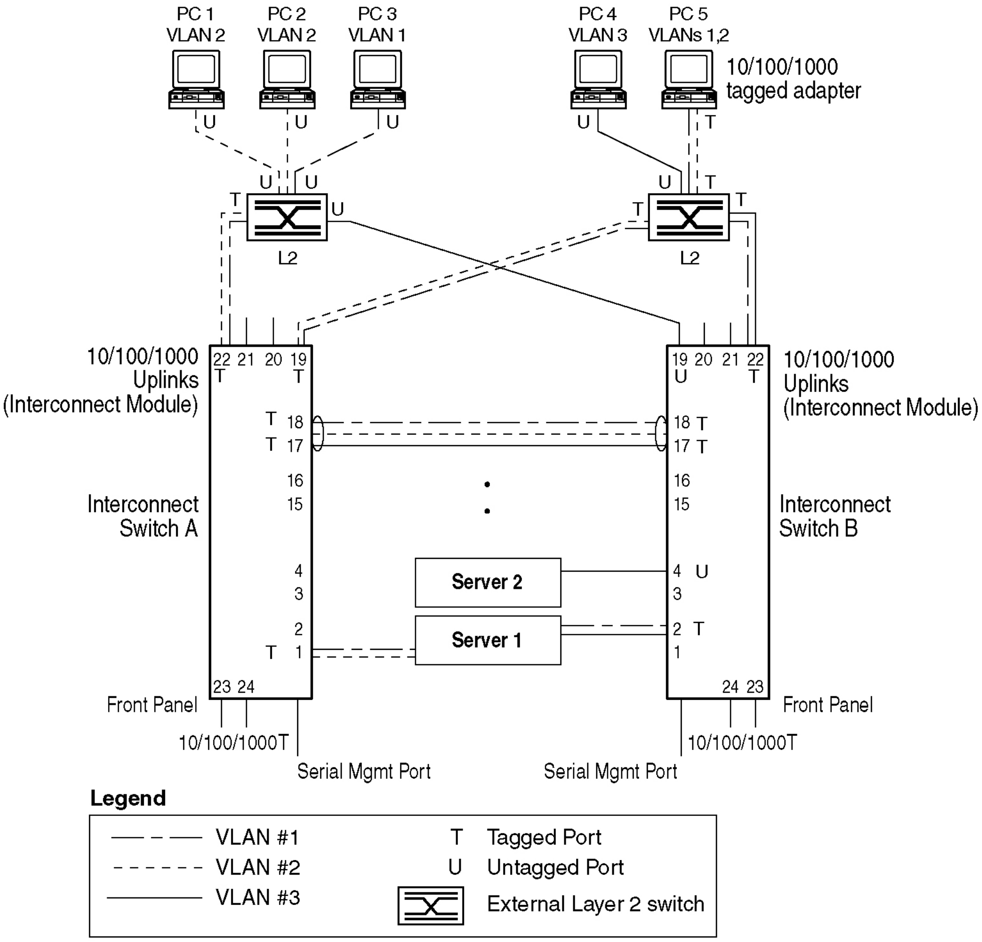

Spanning Tree is Disabled in this example

Figure 3-6: Multiple VLANs with VLAN-tagging

The features of this VLAN are described in the following table:

Table 3-1: Multiple VLANs with Tagging

Component	Description
Switch A	Switch A is configured for VLANS 1 and 2.
Switch B	Switch B is configured for VLANS 1 and 3. Port 2 is tagged to accept traffic from VLANS 1 and 3. Port 4 is configured only for VLAN 3, so VLAN tagging is off.
Blade Server #1	This high-use blade server needs to be accessed from all VLANs and IP subnets. The server has a VLAN-tagging adapter installed with VLAN tagging turned on. The adapter is attached to one of the switch 10/100/1000 Mb/s ports that is configured for VLANs 1, 2, and 3. Because of the VLAN tagging capabilities of both the adapter and the switch, the server is able to communicate on all three VLANs in this network, while maintaining broadcast separation among all three VLANs and subnets.
Blade Server #2	This blade server belongs to VLAN 3. The port that the VLAN is attached to is configured only for VLAN 3, so VLAN tagging is off.
PC #1	A member of VLAN 2, this PC can only communicate with Server 1, PC 2, and PC 5.
PC #2	A member of VLAN 2, this PC can only communicate with Server 1, PC 1, and PC 5.
PC #3	A member of VLAN 1, this PC can communicate with Server 1 and PC5.
PC #4	A member of VLAN 3, this PC can communicate with Server 1and Server 2.
PC #5	A member of both VLAN 1 and VLAN 2, this PC has VLAN-tagging Gigabit Ethernet adapter installed. It can communicate with Server 1 via VLAN 1, and both PC 1 and PC 2 via VLAN 2. The Layer 2 switch port to which it is connected is configured for both VLAN 1 and VLAN 2 and has tagging enabled.

Configuring the Example Network

This example describes how to configure ports and VLANs on Switch A and Switch B.

Configuring Ports and VLANs on Switch A

To configure ports and VLANs on Switch A, do the following:

1. On Switch A, enable VLAN tagging on the necessary ports.

```
Main# /cfg/port 1
>> Port 1# tag e                            (Select port 1: connection to server 1)
Current VLAN tag support: disabled
New VLAN tag support:     enabled           (Enable tagging)
Port 1 changed to tagged.

Main# /cfg/port 17                          (Select inter-switch link port 17)
>> Port 17# tag e                           (Enable tagging)
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 17 changed to tagged.

Main# /cfg/port 18                          (Select inter-switch link port 18)
>> Port 18# tag e                           (Enable tagging)
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 18 changed to tagged.

Main# /cfg/port 19                          (Select uplink port 19)
>> Port 19# tag e                           (Enable tagging)
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 19 changed to tagged.

Main# /cfg/port 22                          (Select uplink port 22)
>> Port 22# tag e                           (Enable tagging)
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 22 changed to tagged.
>> Port 22# apply                           (Apply the port configurations)
```

2. Configure the VLANs and their member ports. Since all ports are by default configured for VLAN 1, configure only those ports that belong to VLAN 2. Inter-switch link ports 17 and 18 must belong to VLANs 1 and 2.

```
>> /cfg/vlan 2
>> VLAN 2# add 1                                       (Add port 1 to VLAN  2)
Current ports for VLAN 2:        empty
Pending new ports for VLAN 2:    1
>> VLAN 2# add 17                                      (Add port 17  to VLAN  2)
Current ports for VLAN 2:        1
Pending new ports for VLAN 2:    17
>> VLAN 2# add 18                                      (Add port 18  to VLAN 2)
Current ports for VLAN 2:        1, 17
Pending new ports for VLAN 2:    18
>> VLAN 2# add 19                                      (Add port 19  to VLAN 2)
Current ports for VLAN 2:        1, 17, 18
Pending new ports for VLAN 2:    19
>> VLAN 2# add 22                                      (Add port 22  to VLAN 2)
Current ports for VLAN 2:        1, 17, 18, 19
Pending new ports for VLAN 2:    22

>> /cfg/vlan 3
>> VLAN 3# add 17                                       (Add port 17 to VLAN  3)
Current ports for VLAN 3:        empty
Pending new ports for VLAN 3:    17
>> VLAN 3# add 18                                      (Add port 18  to VLAN  3)
Current ports for VLAN 3:        17
Pending new ports for VLAN 3:    18

>> apply                                               (Apply the port configurations)
>> save                                                (Save the port configurations)
```

Configuring Ports and VLANs on Switch B

1. On Switch B, enable VLAN tagging on the necessary ports. Port 4 (connection to server 2) remains untagged, so it is not configured below.

```
Main# /cfg/port 2                          (Select port 2: connection to server 1)
>> Port 1# tag e
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 1 changed to tagged.

Main# /cfg/port 17                         (Select inter-switch link port 17)
>> Port 17# tag e                          (Enable tagging)
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 17 changed to tagged.

Main# /cfg/port 18                         (Select inter-switch link port 18)
>> Port 18# tag e                          (Enable tagging)
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 18 changed to tagged.

Main# /cfg/port 22                         (Select uplink port 22)
>> Port 22# tag e                          (Enable tagging)
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 22 changed to tagged.
>> Port 22# apply                          (Apply the port configurations)
```

2. Configure the VLANs and their member ports. Since all ports are by default configured for VLAN 1, configure only those ports that belong to other VLANs.

```
>> /cfg/vlan 2
>> VLAN 2# add 17
Current ports for VLAN 2:        empty
Pending new ports for VLAN 2:    17
>> VLAN 2# add 18
Current ports for VLAN 2:        17
Pending new ports for VLAN 2:    18

>> VLAN 3#/cfg/vlan 3
>> VLAN 3# add 2
Current ports for VLAN 3:        empty
Pending new ports for VLAN 3:    2
>> VLAN 3# add 4
Current ports for VLAN 3:        2
Pending new ports for VLAN 3:    4
>> VLAN 3# add 17
Current ports for VLAN 3:        2, 4
Pending new ports for VLAN 3:    17
>> VLAN 3# add 18
Current ports for VLAN 3:        2, 4, 17
Pending new ports for VLAN 3:    18
>> VLAN 3# add 19
Current ports for VLAN 3:        2, 4, 17, 18
Pending new ports for VLAN 3:    19
>> VLAN 3# add 22
Current ports for VLAN 3:        2, 4, 17, 18, 19
Pending new ports for VLAN 3:    22

>> apply                                         (Apply the configurations)
>> save                                          (Save the port configurations)
```

The external Layer 2 switches should also be configured for VLANs and tagging as shown in Figure 3-6, “Multiple VLANs with VLAN-tagging.”

4

Spanning Tree Protocol

Introduction

When multiple paths exist on a network, Spanning Tree Protocol (STP) configures the network so that a switch uses only the most efficient path. The following topics are discussed in this chapter:

- Overview
- Bridge Protocol Data Units (BPDUs)
- Spanning Tree Group (STG) Configuration Guidelines
- Multiple Spanning Trees

Overview

Spanning Tree Protocol (STP) detects and eliminates logical loops in a bridged or switched network. STP forces redundant data paths into a standby (blocked) state. When multiple paths exist, STP configures the network so that a switch uses only the most efficient path. If that path fails, STP automatically sets up another active path on the network to sustain network operations.

Bridge Protocol Data Units

To create a spanning tree, the application switch generates a configuration Bridge Protocol Data Unit (BPDU), which it then forwards out of its ports. All switches in the Layer 2 network participating in the spanning tree gather information about other switches in the network through an exchange of BPDUs.

A BPDU is a 64-byte packet that is sent out at a configurable interval, which is typically set for two seconds. The BPDU is used to establish a path, much like a “hello” packet in IP routing. BPDUs contain information about the transmitting bridge and its ports, including bridge and MAC addresses, bridge priority, port priority, and port path cost. If the ports are tagged, each port sends out a special BPDU containing the tagged information.

The generic action of a switch on receiving a BPDU is to compare the received BPDU to its own BPDU that it will transmit. If the received BPDU has a priority value closer to zero than its own BPDU, it will replace its BPDU with the received BPDU. Then, the application switch adds its own bridge ID number and increments the path cost of the BPDU. The application switch uses this information to block any redundant paths.

Determining the Path for Forwarding BPDUs

When determining which port to use for forwarding and which port to block, the GbE2 Interconnect Switch uses information in the BPDU, including each bridge priority ID. A technique based on the "lowest root cost" is then computed to determine the most efficient path for forwarding.

Bridge Priority

The bridge priority parameter controls which bridge on the network is the STP root bridge. To make one switch the root bridge, configure the bridge priority lower than all other switches and bridges on your network. The lower the value, the higher the bridge priority. The bridge priority is configured using the **/cfg/stp/brg/prior** command in the CLI.

Port Priority

The port priority helps determine which bridge port becomes the designated port. In a network topology that has multiple bridge ports connected to a single segment, the port with the lowest port priority becomes the designated port for the segment. The port priority is configured using the **/cfg/stp/port/prior** command in the CLI.

Port Path Cost

The port path cost 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 "cost" of 10 in half-duplex mode, and a "cost" of 5 in full-duplex mode. The objective is to use the fastest links so that the route with the lowest cost is chosen. A value of 0 indicates that the default cost will be computed for an auto-negotiated link speed.

Spanning Tree Group Configuration Guidelines

This section provides important information on configuring spanning tree groups (STGs):

Adding a VLAN to a Spanning Tree Group

If no VLANs exist beyond the default VLAN 1, refer to the "Creating a VLAN" section for information on adding ports to VLANs.

Add the VLAN to the STG using the **/cfg/stp <*stg #*>/add <*vlan-number*>** command.

Creating a VLAN

When you create a VLAN, that VLAN automatically belongs to STG 1, the default STG. If you want the VLAN in another STG, you must assign it to another STG.

Keep the following rules in mind when creating a VLAN:

- You cannot delete or move VLAN 1 from STG 1.

  VLANs must be contained within a single STG; a VLAN cannot span multiple STGs. By confining VLANs within a single STG, you avoid problems with STP blocking ports and causing a loss of connectivity within the VLAN. When a VLAN spans multiple switches, the VLAN must be within the same spanning tree group (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 any STG. The port must be added to a VLAN, and that VLAN added to the desired STG.

Rules for VLAN Tagged Ports

Note the following rules for VLAN tagged ports:

- 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.

Adding and Removing Ports from STGs

Note the following when adding and removing ports from STGs:

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

  For example, assume that VLAN 1 belongs to STG 1. You add an untagged port, port 1, that does not belong to any STG to VLAN 1, and port 1 will become part of STG 1.

  If you add untagged port 5 (which is a member of STG 2) to STG 1, the switch will prompt you to change the PVID from 2 to 1:

  ```
  "Port 5 is an UNTAGGED port and its current PVID is 2.
  Confirm changing PVID from 2 to 1:"  y
  ```

- When you remove a port 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.

  As an example, assume that port 1 belongs to VLAN 1, and VLAN 1 belongs to STG 1. When you remove port 1 from VLAN 1, port 1 is also removed from STG 1.

  However, if port 1 belongs to both VLAN 1 and VLAN 2 and both VLANs belong to STG 1, removing port 1 from VLAN 1 does not remove port 1 from STG 1 because VLAN 2 is still a member of STG 1.

- 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.

The relationship between ports, trunk groups, VLANs, and spanning trees is shown in Table 4-1.

Table 4-1: Ports, Trunk Groups, and VLANs

Switch Element	Belongs to
Port	Trunk group, or One or more VLANs
Trunk group	One or more VLANs
VLAN (non-default)	One spanning tree group

Multiple Spanning Trees

Each GbE2 Interconnect Switch supports a maximum of 16 spanning tree groups (STGs). Multiple STGs provide multiple data paths, which can be used for load-balancing and redundancy.

You enable independent links on two GbE2 Interconnect Switches using multiple STGs by configuring each path with a different VLAN and then assigning each VLAN to a separate STG. Each STG is independent. Each STG sends its own Bridge Protocol Data Units (BPDUs), and each STG must be independently configured.

The STG, or bridge group, forms a loop-free topology that includes one or more virtual LANs (VLANs). The switch supports 16 STGs running simultaneously. The default STG 1 may contain an unlimited number of VLANs. All other STGs (2-16) may contain one VLAN each.

Default Spanning Tree Configuration

In the default configuration, STG 1 is enabled with all ports except 1-16 assigned to it. Although ports can be added to or deleted from the default STG, the default STG (STG 1) itself cannot be deleted from the system. Also, you cannot delete the default VLAN (VLAN 1) from STG 1.

All other STGs, except the default STG, are empty and VLANs must be added by the user. However, you cannot assign ports directly to an STG. Ports should be added to a VLAN and the VLAN can be added to the STG. Each STG will be enabled by default and assigned an ID number from 2 to 16.

Why Do We Need Multiple Spanning Trees?

Figure 4-1 shows a simple example of why we need multiple spanning trees. Two VLANs, VLAN 1 and VLAN 2, exist between Switch A and Switch B. If Spanning Tree Protocol is enabled on the switch and there is only a single spanning tree group, the switches see an apparent loop, and will block one of the VLANs.

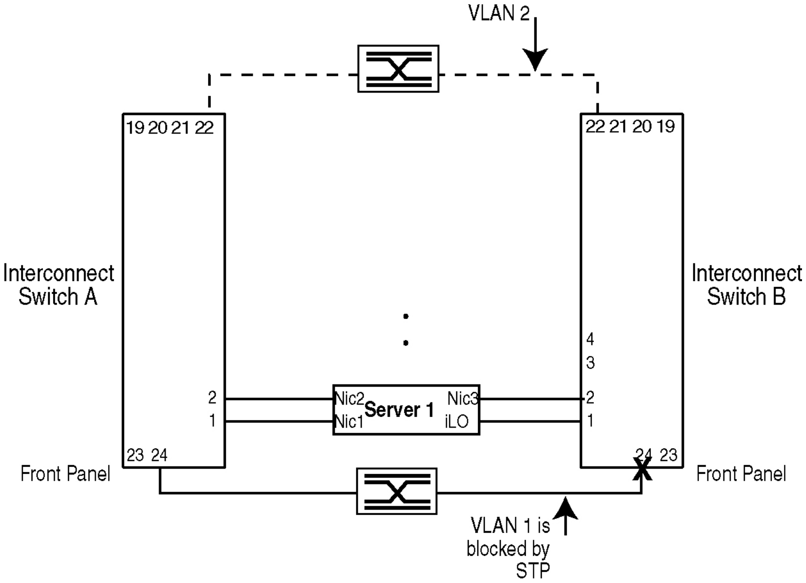

Figure 4-1: Two VLANs on one instance of Spanning Tree Protocol

In Figure 4-2, VLAN 1 and VLAN 2 belong to different spanning tree groups. The two instances of spanning tree separate the topology without forming a loop, so that both VLANs can forward packets between the switches without losing connectivity.

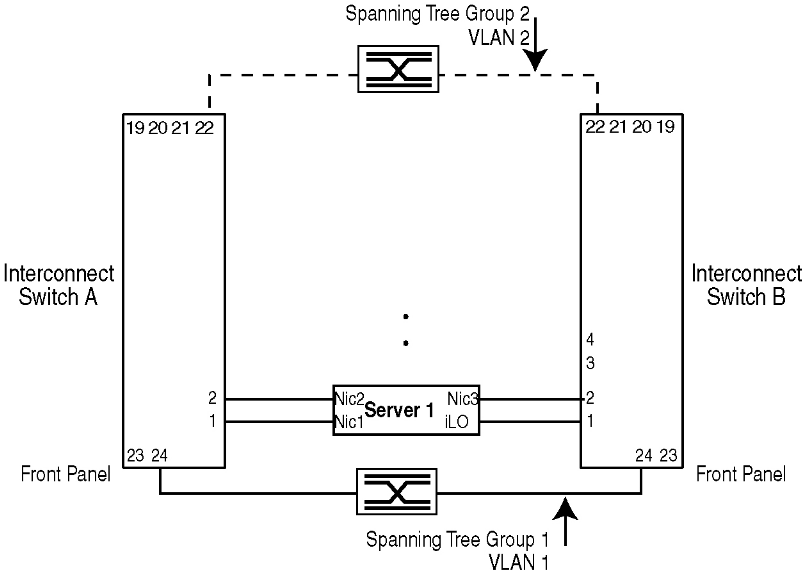

Figure 4-2: Two VLANs on separate instances of Spanning Tree Protocol

Assigning Each VLAN to its Own Spanning Tree Group

The network shown in the following figure can be fixed by assigning each VLAN to its own spanning tree group. Each spanning tree group will detect loops on its own network, without isolating any device.

In this example, VLANs 1, 2, and 3 are placed into spanning tree groups 1, 2, and 3 respectively.

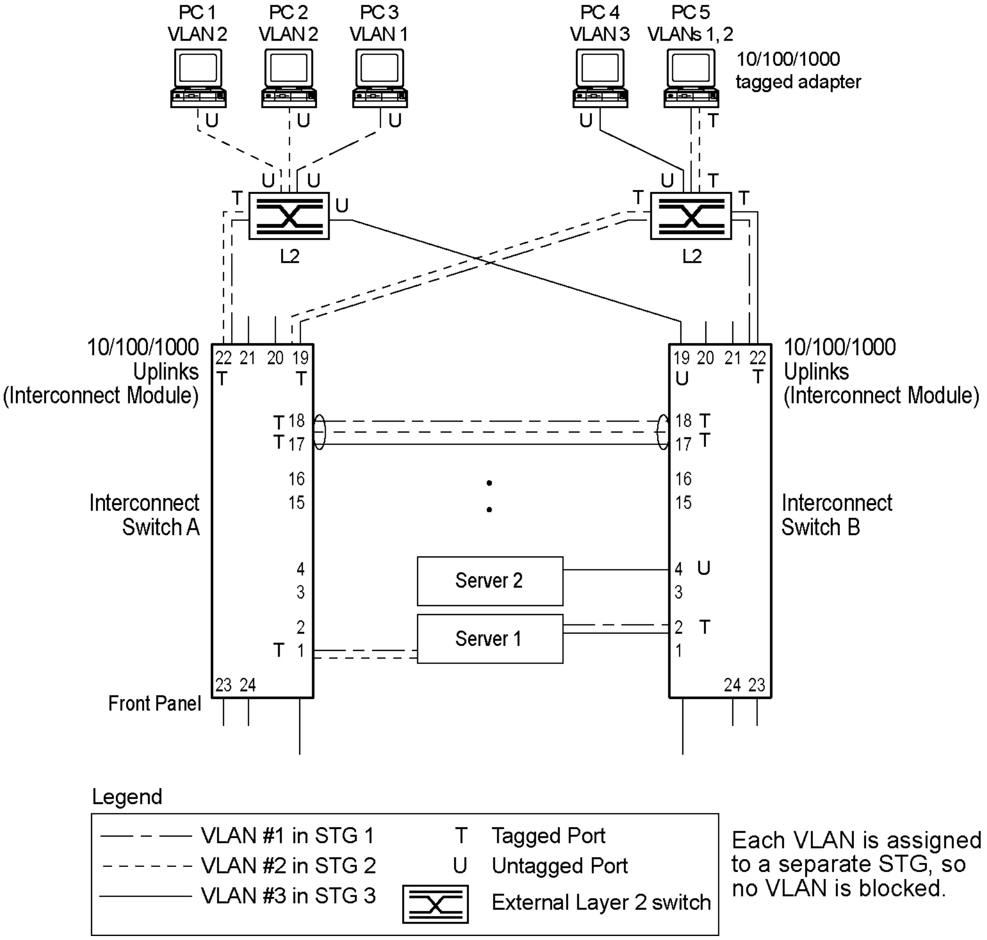

Each VLAN is assigned to a separate STG, so no VLAN is blocked.

Figure 4-3: Implementing multiple spanning tree groups

VLAN Participation in Spanning Tree Groups

The following table shows which switch ports in Figure 4-3 participate in each spanning tree group. Server ports (ports 1-16) do not participate in spanning tree even though they are members of their respective VLANs.

Table 4-2: VLAN Participation in Spanning Tree Groups

	VLAN 1	VLAN 2	VLAN 3
Switch A	Spanning tree group 1	Spanning tree group 2	Spanning tree group 3
	Ports 17, 18, 19, 22	Ports 17, 18, 19, 22	Port 17, 18
Switch B	Spanning tree group 1	Spanning tree group 2	Spanning tree group 3
	Ports 17, 18, 22	Ports 17, 18	Port 17, 18, 19, 22

Configuring Multiple Spanning Tree Groups

This configuration shows how to configure the three instances of spanning tree groups on Switch A and Switch B as illustrated in Figure 4-3.

By default spanning trees 2-16 are empty, and spanning tree group 1 contains all configured VLANs until individual VLANs are explicitly assigned to other spanning tree groups. Except for the default spanning tree group 1, which may contain more than one VLAN group, spanning tree groups 2-16 may contain only one VLAN each.

NOTE: Each instance of spanning tree group is enabled by default.

Configure Switch A

1. Configure port and VLAN membership on Switch A as described in the “Configuring Ports and VLANs on Switch A” section, in Chapter 3.
2. Add VLAN 2 to spanning tree group 2.

```
>> /cfg/stp 2                                 (Select spanning tree group 2)
>> Spanning Tree Group 2# add 2               (Add VLAN 2)
```

VLAN 2 is automatically removed from spanning tree group 1.

Configure Switch B

1. Configure port and VLAN membership as described in the "Configuring Ports and VLANs on Switch B" section in Chapter 3.
2. Add VLAN 3 to spanning tree group 3.

```
>> /cfg/stp 3                                    (Select spanning tree group 3)
>> Spanning Tree Group 3# add 3                  (Add VLAN 3)
```

   VLAN 3 is automatically removed from spanning tree group 1.

3. Apply and save.

```
>> apply
>> save
```

A

Troubleshooting Tools

Introduction

This appendix discusses some tools to help you use the port mirroring feature to troubleshoot common network problems on the HP ProLiant BL p-Class GbE2 Interconnect Switch.

Port Mirroring

The port mirroring feature on the GbE2 Interconnect Switch is very useful for troubleshooting any connection-oriented problem. Any traffic in or out of one or more ports can be mirrored to a single mirror port to which a network monitor can be attached.

Port mirroring can be used as a troubleshooting tool or to enhance the security of your network. For example, an Intrusion Detection Service (IDS) server can be connected to the monitor port to detect intruders attacking the network.

As shown in Figure A-1, port 19 is monitoring ingress traffic (traffic entering the GbE2 Interconnect Switch) on port 23 and egress traffic (traffic leaving the GbE2 Interconnect Switch) on port 1. You can attach a device to port 19 to monitor the traffic on ports 23 and 1.

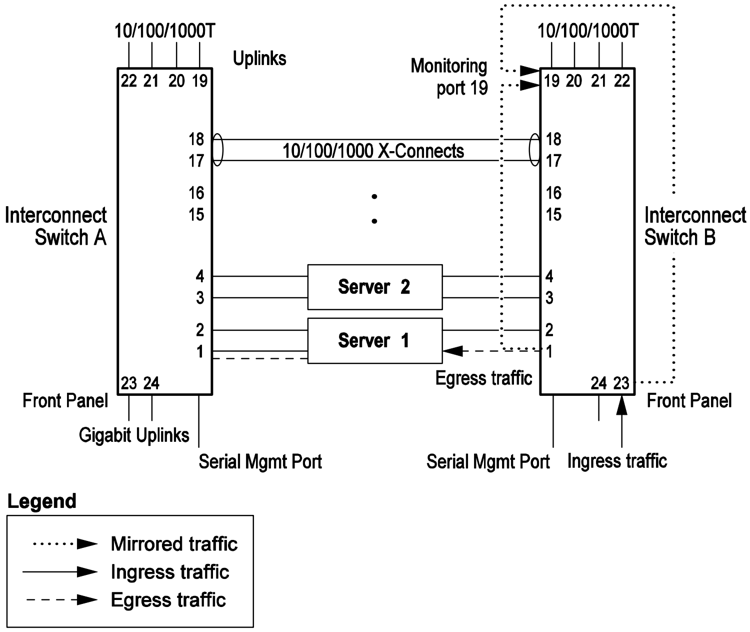

Figure A-1: Port mirroring

Figure A-1 shows two mirrored ports monitored by a single port. Similarly, you can have one mirrored port to one monitored port, or many mirrored ports to one monitored port. The GbE2 Interconnect Switch does not support a single port being monitored by multiple ports because it supports only one monitored port configured at a time.

Ingress traffic is duplicated and sent to the mirrored ports before processing, and egress traffic is duplicated and sent to the mirrored ports after processing.

To configure port mirroring for the example shown in Figure A-1:

1. Specify the monitoring port.

```
>> # /cfg/pmirr/monport 19                    (Select port 19 for monitoring)
```

2. Select the ports that you want to mirror.

```
>> Port 19 # add 23                           (Select port 23 to mirror)
>> Enter port mirror direction [in, out, or both]: in
                                              (Monitor ingress traffic on port 23)
>> Port 19 # add 1                            (Select port 11 to mirror)
>> Enter port mirror direction [in, out, or both]: out
                                              (Monitor egress traffic on port 1)
```

3. Enable port mirroring.

```
>> # /cfg/pmirr/mirr ena                                   (Enable port mirroring)
```

4. Apply and save the configuration.

```
>> PortMirroring# apply                                (Apply the configuration)
>> PortMirroring# save                                 (Save the configuration)
```

5. View the current configuration.

```
>> PortMirroring# cur                                  (Display the current settings)
Port mirroring is enabled
Monitoring Ports    Mirrored Ports
1      none
2      none
3      none
4      none
5      none
   :
   :
17     none
18     none
19     (23, in) (1, out)
20     none
   :
```

Other Network Troubleshooting Techniques

Other network troubleshooting techniques include the following.

Console and Syslog Messages

When a GbE2 Interconnect Switch experiences a problem, review the console and Syslog messages. The GbE2 Interconnect Switch displays these informative messages when state changes and system problems occur. Syslog messages can be viewed by using the **/info/syslog** command. For more information on interpreting syslog messages, refer to the *HP ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide*.

Ping

To verify station-to-station connectivity across the network, execute the following command:

ping <*host name*> | <*IP address*> [(*number of tries*)> [*msec delay*]]

The IP address is the hostname or IP address of the device. The number of tries (optional) is the number of attempts (1-32). Msec delay (optional) is the number of milliseconds between attempts.

Trace Route

To identify the route used for station-to-station connectivity across the network, execute the following command:

traceroute <*host name*> | <*IP address*> [<*max-hops*> [*msec delay*]]

The IP address is the hostname or IP address of the target station. Max-hops (optional) is the maximum distance to trace (1-16 devices). Msec delay (optional) is the number of milliseconds to wait for the response.

Statistics and State Information

The GbE2 Interconnect Switch keeps track of a large number of statistics and many of these are error condition counters. The statistics and state information can be very useful when troubleshooting a LAN or Real Server problem. For more information about available statistics, refer to the **/stats** commands in the *HP ProLiant BL p-Class GbE2 Interconnect Switch Command Reference Guide*.

Customer Support Tools

The following diagnostics tools are not user-configurable and should be performed through HP technical support.

- Offline Diagnostics—this tool is used for troubleshooting suspected GbE2 Interconnect Switch hardware issues. These tests verify that the selected hardware is performing within expected engineering specifications.
- Software Panics—if a fatal software condition is found during runtime, the GbE2 Interconnect Switch will capture the current hardware and software state information into a panic dump. This dump file can be analyzed post-mortem to determine the cause of the problem.

Index