hp industry
standard servers

march 2003

invent

white paper
16VZ-1102A-WWEN

ProLiant BL p-Class GbE Interconnect Switch Overview

table of contents

abstract

This white paper provides an overview of the ProLiant BL p-Class GbE Interconnect Kit[1] options for the ProLiant BL p-Class system. Used instead of a pair of "edge" switches, the GbE Interconnect Kit includes a pair of industry-standard layer 2 Ethernet switches that dramatically reduce the number of Ethernet network cables attached to the rear of the ProLiant BL p-Class server blade. The HP ProLiant BL p-Class system includes dual-processor mid-tier and quad-processor back-end server blades that provide data centers and service providers high levels of flexibility and efficiency from their IT organizations and an overall better return on their IT investment.

The intended audience for this paper is engineers and system administrators familiar with the HP ProLiant BL p-Class system. For readers not familiar with the HP ProLiant BL p-Class system, more information is available at http://h18004.www1.hp.com/products/servers/platforms/index-bl.html.

introduction

The ProLiant BL product line is a portfolio of industry-standard server blades for use in multi-tiered environments. They can be quickly deployed in industry-standard racks along with legacy servers and storage devices. These server blades address the needs of space-constrained data centers and service providers for increased server density, rapid deployment and provisioning, and remote manageability. The ProLiant BL p-Class system includes high-performance, dual-processor mid-tier server blades and quad-processor back-end server blades designed for rapid deployment and provisioning, virtual presence and control, and rip-and-replace serviceability.

Since many ProLiant BL server blades are packaged in a small space, the number of Ethernet network cables within this space can quickly become overwhelming. A ProLiant BL p-Class server enclosure supports up to eight server blades, each with four Ethernet network adapters.

The ProLiant BL p-Class GbE Interconnect Kit is one of the connectivity options available with the ProLiant BL p-Class system. In a typical Ethernet network, edge switches such as the HP ProCurve switch 2524 are used to collect Ethernet signals from the servers and concentrate them to the Ethernet network backbone. Designed to replace a pair of edge switches, the GbE Interconnect Kit concentrates and significantly reduces the number of server blade network signals for direct connection to the Ethernet network backbone. The GbE Interconnect Kit uses non-blocking, managed layer 2 Ethernet switching technology to provide up to a 32-to-1 reduction in network cables at the back of the BL p-Class enclosure, while providing network redundancy and a full set of enterprise-level features. This cable reduction significantly reduces the time required to deploy, manage, and service ProLiant BL p-Class systems.

The GbE Interconnect Kit option includes a pair of redundant, hot-pluggable and hot-removable GbE Interconnect Switches that mount from the front of the server blade enclosure like bookends (Figure 1). Two kit options are available: the C-GbE Interconnect Kit for copper-based networks and the F-GbE Interconnect Kit for fiber-based networks. The different kit options are discussed later in this paper.

[1] GbE is the acronym for Gigabit Ethernet.

figure 1. ProLiant BL p-Class server blade enclosure (front view) with the C-GbE Interconnect Kit and BL20p G2 server blades

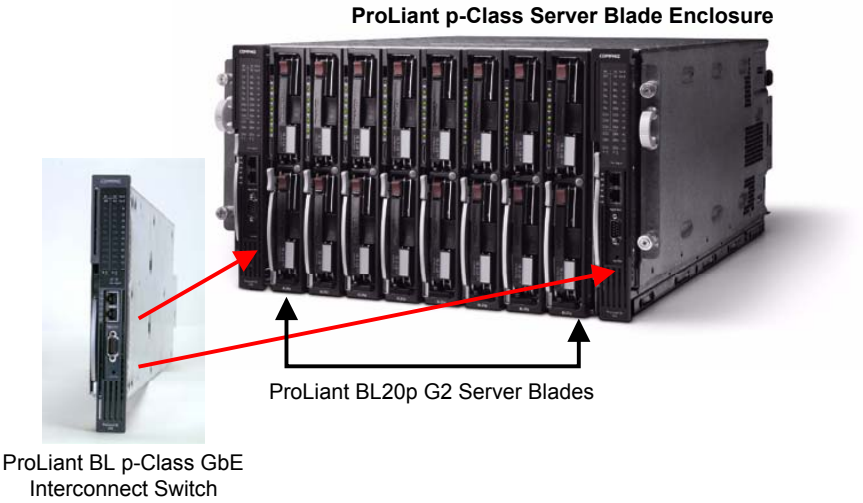

ProLiant BL p-Class network routing

The ProLiant BL p-Class server enclosure has eight server bays and two interconnect bays. Each server blade bay supports three general purpose network adapters and one additional 10/100 Ethernet adapter dedicated to Integrated Lights-Out (iLO) management. Therefore, each server blade enclosure can have up to 32 network adapters.

The signals from these 32 network adapters are routed from the servers, across category 5e specified signal traces on the passive backplane assembly, and to the interconnect blades (Figure 2). At the rear of the interconnect blades are hot-pluggable interconnect modules that plug into the interconnect blade and provide the network connectors for connection to the Ethernet network backbone.

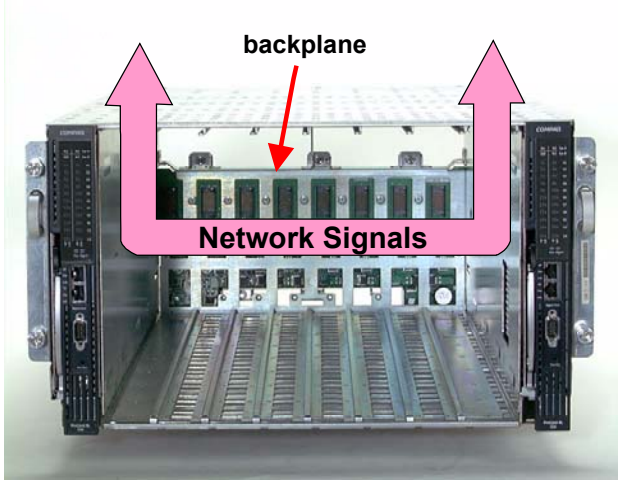

figure 2. ProLiant BL p-Class server enclosure

To provide redundant access to each server, half the network signals from each server go to the left interconnect blade and the other half go to the right interconnect blade (Figure 3). Each server blade enclosure requires a pair of interconnects. HP offers different types of interconnect kits for a choice of Ethernet network connectivity and level of network cable reduction.

figure 3: ProLiant BL p-Class network routing redundancy

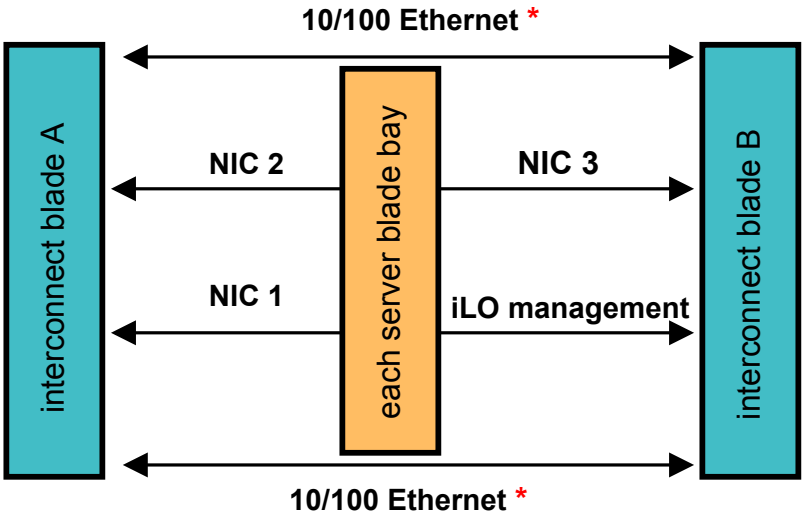

* C-GbE Interconnect Switch only

RJ-45 Patch Panel Kit

The RJ-45 Patch Panel Kit functions as an Ethernet pass through only; it is designed for customers who want to have flexibility in choosing Ethernet network switches. Consisting of a pair of patch panel interconnect blades, the kit brings all Ethernet signals out as separate RJ-45 connections (Figure 4). A server blade enclosure full of 2-processor server blades has 32 Ethernet signals, and a server blade enclosure full of 4-processor server blades has 12 Ethernet signals. A mixture of 2-processor and 4-processor server blades can have up to 22 RJ-45 Ethernet signals.

figure 4. ProLiant BL p-Class server blade enclosure with RJ-45 Patch Panel (rear view)

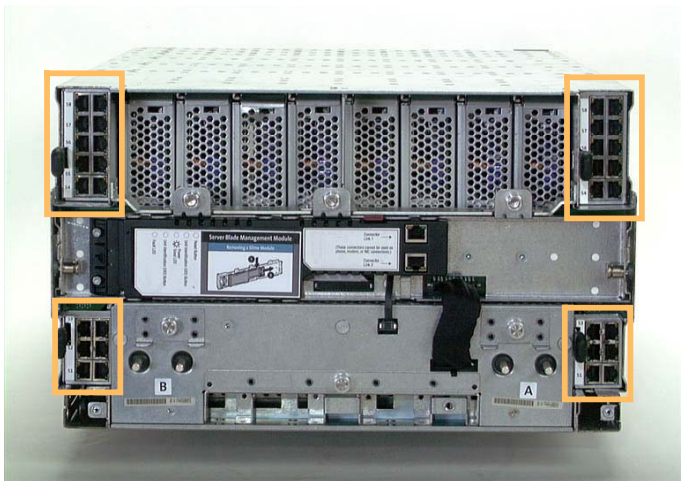

RJ-45 Patch Panel 2 Kit

For Ethernet signal pass through, the RJ-45 Patch Panel 2 Kit functions in exactly the same manner as the previously discussed RJ-45 Patch Panel Kit. In addition to Ethernet pass-through, the RJ-45 Patch Panel 2 Kit provides Fibre Channel pass through for up to eight ProLiant BL20p G2 servers with the Dual Port Fibre Channel Mezzanine Card installed. For additional information about the Patch Panel 2 and ProLiant BL20p G2 Fibre Channel SAN connectivity, please see the ProLiant BL p-Class web page at http://h18004.www1.hp.com/products/servers/proliant-bl/p-class/index.html.

GbE Interconnect Kits

As an alternative to a patch panel kit, the GbE Interconnect Kits are designed to dramatically reduce the number of Ethernet network cables attached to the rear of the server blade enclosure. The GbE Interconnect Kits include two GbE Interconnect Switches that concentrate the 32 server blade Ethernet signals to one to eight uplinks for connection to the Ethernet network backbone (Figure 5). The upcoming GbE2 Interconnect Switch (projected availability in 2Q03) will also provide optional pass-through of ProLiant BL20p G2 Fibre Channel signals and consolidation of Ethernet signals to 1000 Mbps gigabit speeds. More information on the GbE2 Interconnect Switch will be available in 2Q03.

Two different GbE Interconnect Kits are available: C-GbE for copper-based and F-GbE for fiber-based Gigabit Ethernet networks, respectively. The remaining portion of this paper is dedicated to discussing in detail the GbE Interconnect Kit options.

figure 5. ProLiant BL p-Class server blade enclosure with C-GbE Interconnect Kit (rear view)

same network technology, different form factor

In a typical server Ethernet network configuration designed with redundancy, two or more Ethernet network adapters (NICs) are used per server. The Ethernet signals from these NICs are typically routed to two separate Ethernet edge switches. One or more connections are commonly made between the edge switches for added availability. The edge switches are used to collect NIC signals from the servers for concentration to the Ethernet network backbone, typically via gigabit uplinks. The ProLiant BL p-Class GbE Interconnect Switch and blade architecture accomplish the same thing using the same technology but in an integrated, compact form factor designed for rapid deployment and serviceability (Figure 6).

Using a GbE Interconnect Kit, the edge switches and Ethernet network connections to the servers move inside the BL p-Class server blade enclosure. Using the uplinks on the GbE Interconnect Switches, the server blade enclosure is directly connected to the Ethernet network backbone. The interconnect switches are configured and administered in the same manner as any other industry-standard Ethernet switch. The GbE Interconnect Switch is an industry standard managed layer 2 Ethernet switch compatible with other industry standard Ethernet devices.

figure 6. typical redundant server Ethernet network configuration

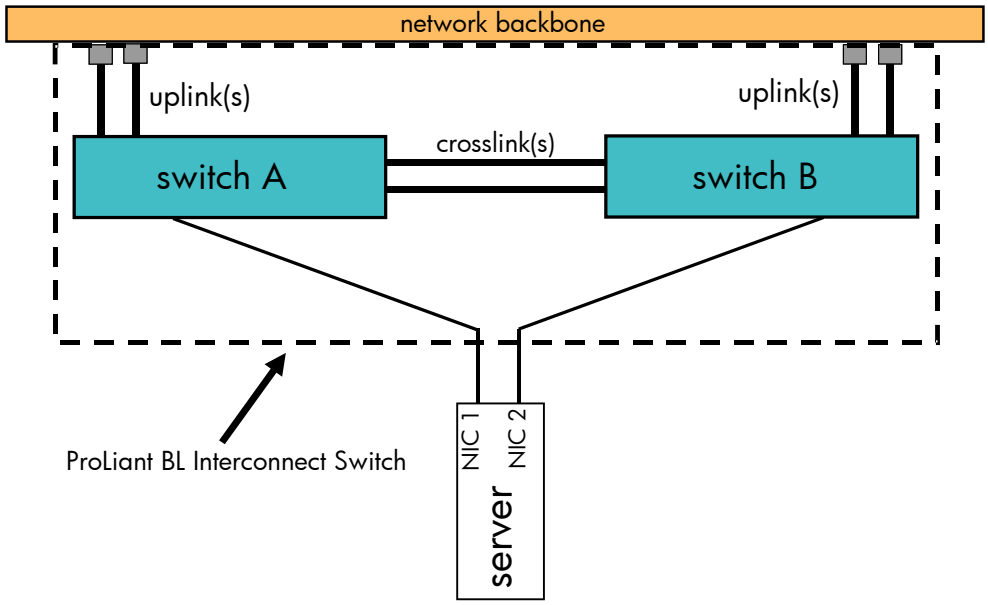

simplified network cabling

As mentioned earlier, each 6U ProLiant BL p-Class server blade enclosure supports up to 32 Ethernet connections. Therefore, a single, fully configured BL p-Class server blade enclosure would require up to 32 Ethernet cables; and a 42U rack fully configured with BL p-Class server blades would require 192 Ethernet cables. A pair of GbE Interconnect Switches collects up to 4 Ethernet signals (to 100 Mbps) per server blade bay (32 total signals per enclosure) and concentrates them for connection to the external Ethernet network. For maximum cable reduction (97 percent), the 32 signals can be concentrated into a single Gigabit Ethernet uplink port (Figure 7).

It possible to use just one uplink per server enclosure and a total of only 14 Ethernet connections for a 42U rack containing 192 network adapters. However, system administrators may use any combination from one to eight uplink ports to increase bandwidth, to separate network and management data onto physically isolated ports, or to add redundant connections to the Ethernet network backbone.

Since GbE Interconnect Switches use industry-standard managed layer 2 Ethernet switch technology, they may be linked together with interconnect switches from different server enclosures and other vendor Ethernet switches within the same rack or location.

figure 7. ProLiant BL p-Class GbE Interconnect Switch Ethernet network cable reduction

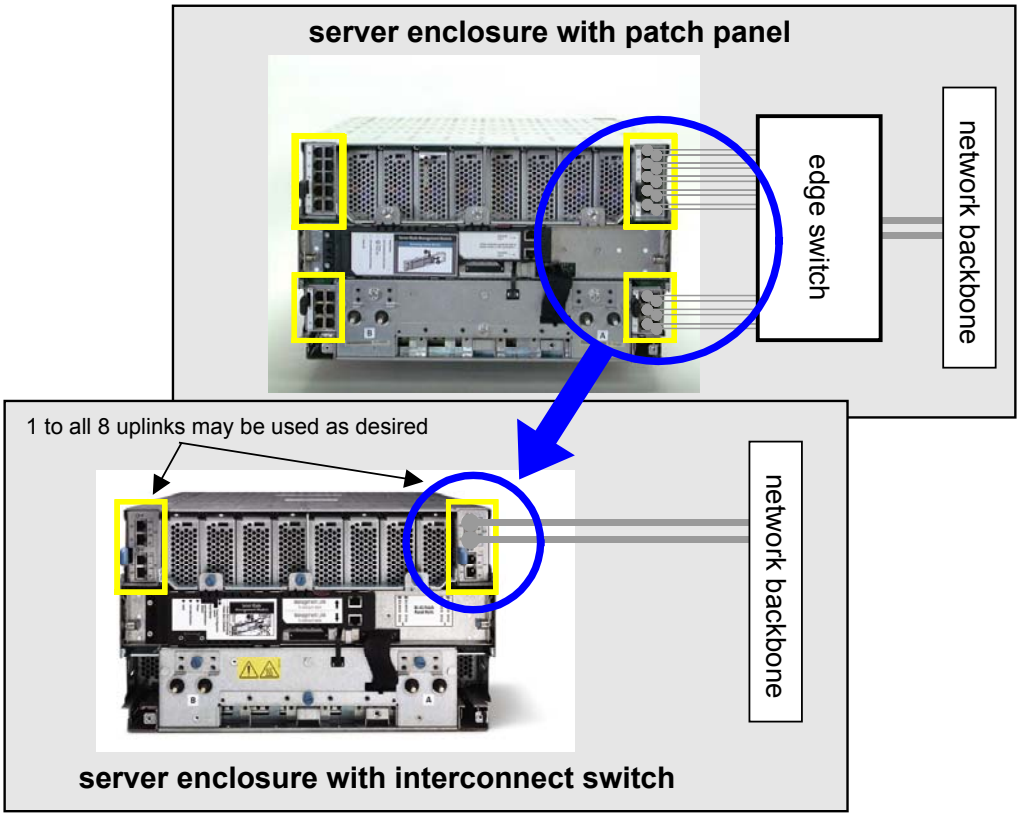

For example, Figure 8 shows a configuration with three fully populated BL p-Class server blade enclosures, each with the C-GbE Interconnect Kit installed. The interconnect switches are linked, or daisy changed, together in redundant configuration using the four gigabit uplinks that, in turn, are connected to the Ethernet network backbone. Each enclosure contains thirty-two 10/100 Ethernet network adapters with an aggregate bandwidth of 3.2 Gbps, for a total bandwidth of 9.6 Gbps for entire system of three enclosures. However, since uplinks are daisy changed together, the maximum system throughput for this configuration is 3.2 Gbps. This configuration creates a 3x blocking ratio (9.6 Gbps vs 3.2 Gbps); nevertheless, it reduces the total Ethernet network cables at the system or rack level from 96 to 4, a 96 percent cable reduction, while still maintaining redundant connections to the Ethernet network backbone.

Figure 8: GbE Interconnect Switch linking

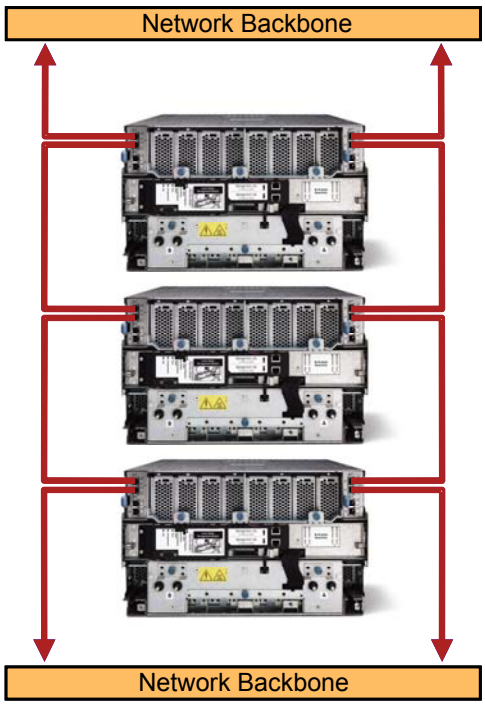

In the example shown in Figure 8, four gigabit uplinks per server blade enclosure were used. Since the two GbE Interconnect Switches provide a total of eight uplinks per enclosure, a wide variety of linking configurations are possible. Additionally, the two interconnect switches within each enclosure are connected by redundant crosslinks, which provide additional linking configuration possibilities.

interconnect switch architecture

The p-Class GbE Interconnect Kits feature a fully redundant architecture including two separate interconnect switches and redundant paths to the Ethernet network ports on the server blades. This architecture provides a highly available, fully meshed topology to the external Ethernet network (Figure 9).

Each interconnect switch (labeled Interconnect Switch A and B) directly connects to two of the network adapters (labeled as 1 through 4) on each server blade via the server blade enclosure's passive backblane. Each GbE Interconnect Switch includes two 10/100T and two Gigabit Ethernet uplink ports. Besides providing a completely non-blocking network configuration, these uplinks can be used to physically isolate data and provide redundant connections to the Ethernet network backbone. IEEE 802.1D spanning tree protocol is supported to eliminate potential network loops caused by redundant networking paths.

Figure 9: ProLiant BL p-Class GbE Interconnect Kit architecture

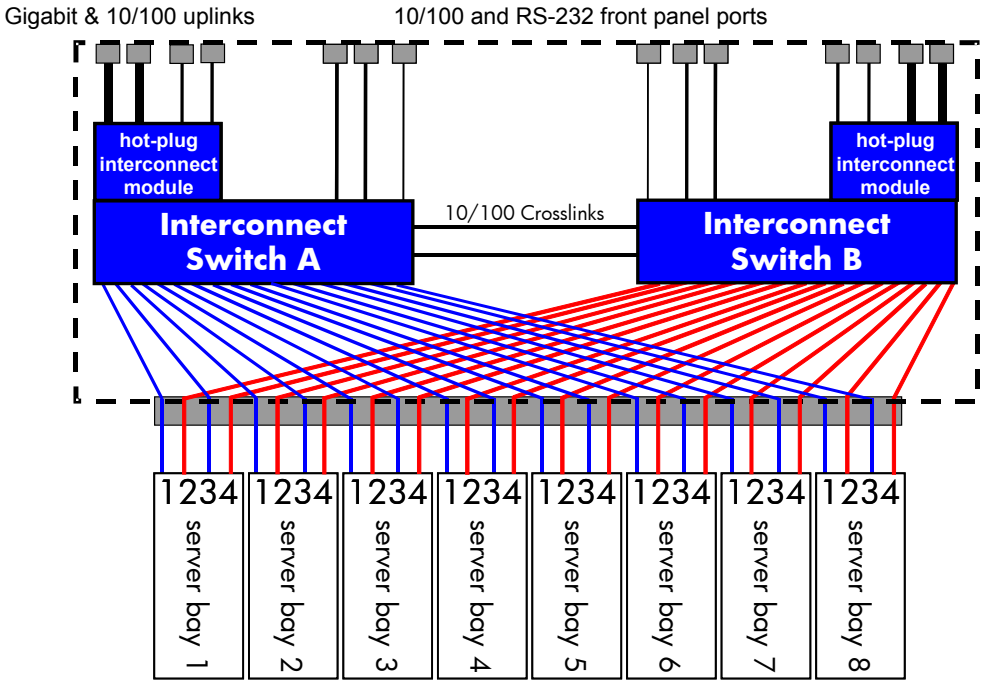

Interconnect switches A and B are further connected through redundant 10/100 Mbps crosslinks bundled into a 200 Mbps multi-link Ethernet trunk. The crosslinks permit communication between the switches, if desired. These crosslinks make it possible to see any server network adapter from any switch uplink port. As a result, any single uplink port may be used to connect to all 32 server Ethernet connections for the 32-to-1 network cable reduction. However, on a heavily utilized system, using a single uplink port for all 32 network adapters can cause a traffic bottleneck. For example, using only a single uplink from interconnect switch A requires the traffic from network adapters 2 and 4 to

travel over the crosslinks (a 200-Mb path). The crosslinks are intended primarily as a failover route and generally are not used as a primary path. For more optimal performance, at least one uplink port on each interconnect switch could be used.

Each interconnect switch includes redundant fans and is powered by the ProLiant BL p-Class redundant, hot-plug power enclosure. Like the server blades, the interconnect switch accepts redundant -48 VAC power from the p-Class power enclosure or directly from a facility that has DC power available.

The redundant architecture of the GbE Interconnect Kit allows system administrators to configure the Ethernet network for continued access to each server blade in case of the following failures:

- interconnect switch
- switch within the Ethernet network backbone
- server blade network adapter
- server blade NIC-to-interconnect switch port or connection
- uplink port and uplink port connection and cable
- interconnect switch cross-link port or connection
- power and fan failure

For example, if interconnect switch A were to fail for any reason, Ethernet network access to all servers would still be possible via the uplink ports on interconnect switch B. Similarly, if access to the uplinks on interconnect switch A were inaccessible, network traffic could be sent through interconnect switch B and then directly to the network adapters 2 and 4, and also to network adapters1 and 3 via the crosslinks. In this case, access to all 32 network adapters would be possible, but network adapter 1 and 3 traffic would be via the two 10/100 Mbps crosslinks. As explained earlier, these crosslinks are intended primarily as a failover route. Such a failover scenario should be viewed as a temporary situation if the crosslink bandwidth could cause a potential bottleneck.

The GbE Interconnect Switch supports server network adapter speeds to 100 Mbps. However, ProLiant BL p-Class server blades with 10/100/1000T gigabit network adapters may be used. In this case, the GbE Interconnect Switch downlink ports (server blade to switch) will auto-negotiate the network adapter speed to 100 Mbps. For example, the ProLiant BL20p G2 server blade includes three NC7781 10/100/1000 Mbps network adapters. The GbE Interconnect Switch will then auto-negotiate the NC7781 network adapter speed to 100 Mbps. The ProLiant BL p-Class GbE Interconnect Switch Kit is designed for applications requiring Ethernet network adapter speeds to 100 Mbps. The RJ-45 Patch Panel, RJ-45 Patch Panel G2, or upcoming GbE2 Interconnect Switch may be used if Ethernet network adapter speeds to 1000 Mbps are desired.

modular rip-n-replace design

Each p-Class GbE Interconnect Switch consists of a modular interconnect switch blade and an interconnect module. The interconnect switch blade slides into the front of the server blade enclosure, and the interconnect module installs from the rear (Figure 10). The interconnect module blind mates into the interconnect switch blade. It includes an independent guide pin to ensure insertion alignment and a robust connector assembly tot prevent mating of incompatible solutions.

figure 10. modular design of ProLiant BL p-Class GbE Interconnect Switch

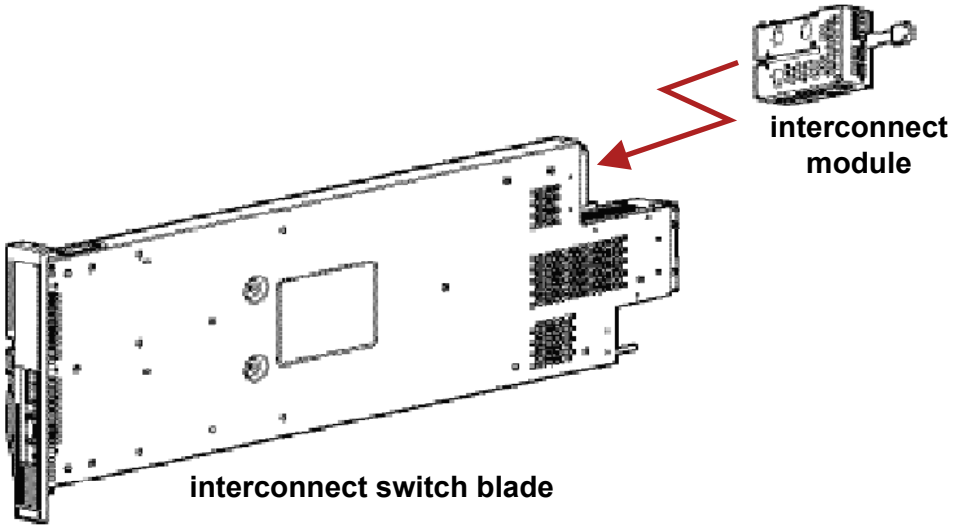

Like the interconnect switch blade, the interconnect module is hot pluggable and hot removable. The interconnect module includes the uplink port connectors and determines the supported Ethernet network cable media: copper-based with RJ-45 connector or fiber-based with LC connector (Figure 11).

This modular design provides two key benefits:

- The interconnect switch can be quickly and easily removed and replaced from the front of the rack without the need for re-cabling.
- Different types of Ethernet network connectors can be supported with no need for a completely new interconnect option.

figure 11. ProLiant BL p-Class QuadT Interconnect Module

Both p-Class GbE Interconnect Switch kits contain two GbE interconnect Switch blades and two interconnect modules. The kits are identical with exception to the interconnect modules (Table 1).

table 1. contents of ProLiant BL p-Class GbE Interconnect Switch Kits

components	C-GbE Interconnect Kit	F-GbE Interconnect Kit
interconnect switches	2x GbE	2x GbE
interconnect modules	2x QuadT providing: • 2x 10/100/1000 BASE-T/TX/T uplinks with RJ-45 connectors • 2x 10/100 BASE-T/TX uplinks with RJ-45 connectors	2x DualTSX providing: • 2x 1000 BASE-SX uplinks with LC connectors • 2x 10/100 BASE-T/TX uplinks with RJ-45 connectors

investment protection

The GbE Interconnect Switch Kits are designed to support any mix of ProLiant BL210p, BL20p G2, and BL40p servers. With the p-Class server blade enclosure, many different configurations of 2-processor and 4-processor servers are possible. For example, configurations may exist with only ProLiant BL20p or BL20p G2 servers, a mix of ProLiant BL20p and BL20p G2 servers, only BL40p servers, or a mix of BL40p plus BL20p and BL20p G2 servers. The GbE Interconnect Switch Kits support all these combinations.

interconnect switch management

Since the GbE Interconnect Switch option is an industry standard managed layer 2 Ethernet switch, it is configured and administered in the same manner as any other industry standard Ethernet switch. The interconnect switch includes a default configuration that that is fully operational at initial boot.

Web browser-based, command line (CLI) and menu-driven console user interfaces are provided to configure, manage, and monitor the interconnect switch on a per port basis. SNMP (simple network management protocol) and RMON (remote monitoring) manageability are further supported. Even an SNMP-based scripting utility and recommended best-case example scripts are included. Any combination of the downlink ports, uplink ports, and cross connection ports can be disabled, enabled, configured, and monitored on a per port basis as required. Out-of-band and in-band access to the switch management interfaces are supported locally and remotely from anywhere on the network. Administration of both switches is possible via any single uplink port. Additionally, switch administration is possible via the front panel serial port and the two 10/100T ports on each interconnect switch (Figure 12).

figure 12. ProLiant BL p-Class GbE Interconnect Switch front panel

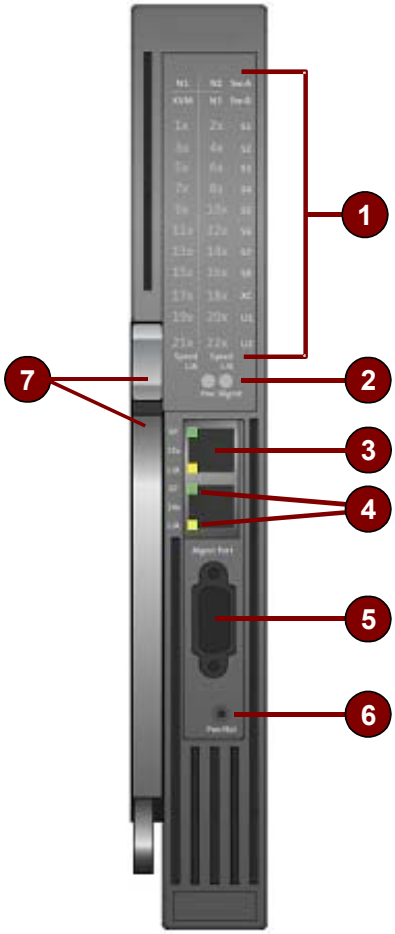

legend

item	description
1	LED panel for link speed and activity status per port
2	power and management status LEDs
3	two local-access 10/100T Ethernet switch ports
4	RJ-45 connector link speed and activity LEDs
5	DB-9 serial connector for access to the CLI and menu driven console
6	power/reset button
7	blade chassis latch and handle

user account security

To secure the management interfaces, the switch administrator can configure multiple password protected accounts with various levels of access. Access can be restricted to a limited number of management stations by configuring a list of IP addresses of those stations that can access the interconnect switch management console. The management interfaces can be configured with an idle time-out period. The maximum number of MAC (mandatory access control) addresses learned per port may also be restricted.

Web browser-based interface

The included Web browser-based interface is accessed using a web browser over a TCP/IP network (Figure 13). Thus, access is possible via any Ethernet switch port, such as the gigabit uplinks, and also through the front panel 10/100T Ethernet ports on each interconnect switch. In addition, to complete configuration and administration of the switch on a per port basis, the Web browser-based interface includes an active virtual graphical display of the front of each interconnect switch. The graphic provides real time status of the per port LED annunciation panel and provides the means to quickly view the statistics of any port. The GbE Interconnect Switch provides graphical and tabular monitoring screens of port utilization, data packets received and transmitted, port error packets, packet size, trunk utilization, and more.

figure 13. browser-based interface for ProLiant BL p-Class GbE Interconnect Switch

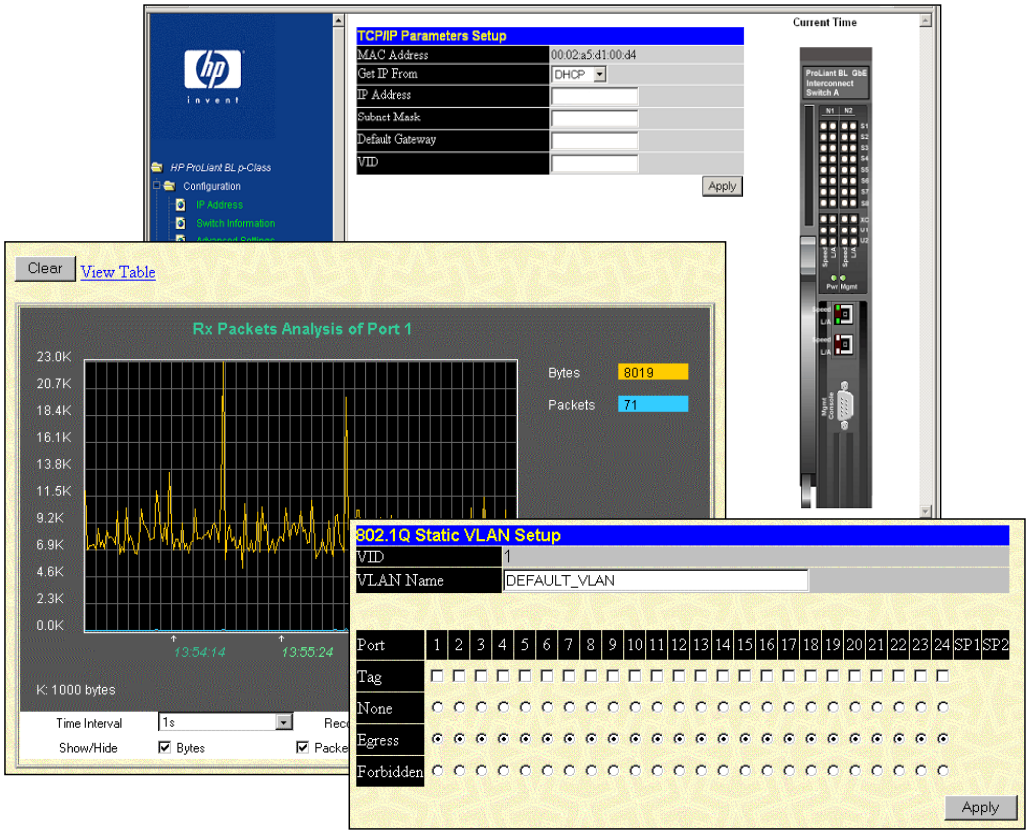

command line and menu-driven interfaces

The command line and the menu-driven console interfaces provide the same configuration options as the Web browser-based interface, plus the added benefits of an out-of-band connection and CLI scripting (Figure 14). There are three methods of accessing these interfaces:

- Locally via the front panel RS-232 console port on each interconnect switch
- Remotely using serial line IP (SLIP) via the front panel RS-232 console port on each interconnect switch
- Remotely using a console Telnet session over a TCP/IP network via the front panel 10/100T port or any other interconnect switch Ethernet port

figure 14. console interface for ProLiant BL p-Class C-GbE Interconnect Switch

simple network management protocol and remote monitoring

Each interconnect switch can be monitored remotely from an SNMP/RMON based network management station such as Insight Manager 7 and HP OpenView. An interconnect switch SNMP scripting utility and example scripts are included. The interconnect switch supports industry-standard SNMP management information bases (MIBs), HP enterprise switch MIBs, and RMON groups 1 (statistics), 2 (History), 3 (Alarm), and 9 (Event) for fault detection, configuration, and monitoring of switch functionality.

trivial file transfer protocol support

The trivial file transfer protocol (TFTP) feature of the interconnect switch permits a TFTP server to upload a copy of the interconnect switch configuration file that can be saved and downloaded to the original or to a different interconnect switch. This provides a means to rapidly deploy multiple systems with similar configurations and to provide backup and restore capabilities. In addition, a history log can be uploaded from the interconnect switch to the TFTP server. Firmware upgrades via TFTP can be performed as well.

switch diagnostics

As previously mentioned, the interconnect switch includes fault detection and monitoring of switch functionality through a SNMP manager or with a RMON probe. Each switch supports four configurable community strings for access via SNMP and SNMP trap managers.

Each interconnect switch uplink and front panel port connector has LEDs for link speed and activity (RJ-45) and for activity (LC). The front of each interconnect switch also has a LED panel for link speed and activity status per port (Figure 15).

figure 15. front port LED panel for ProLiant BL p-Class GbE Interconnect Switch

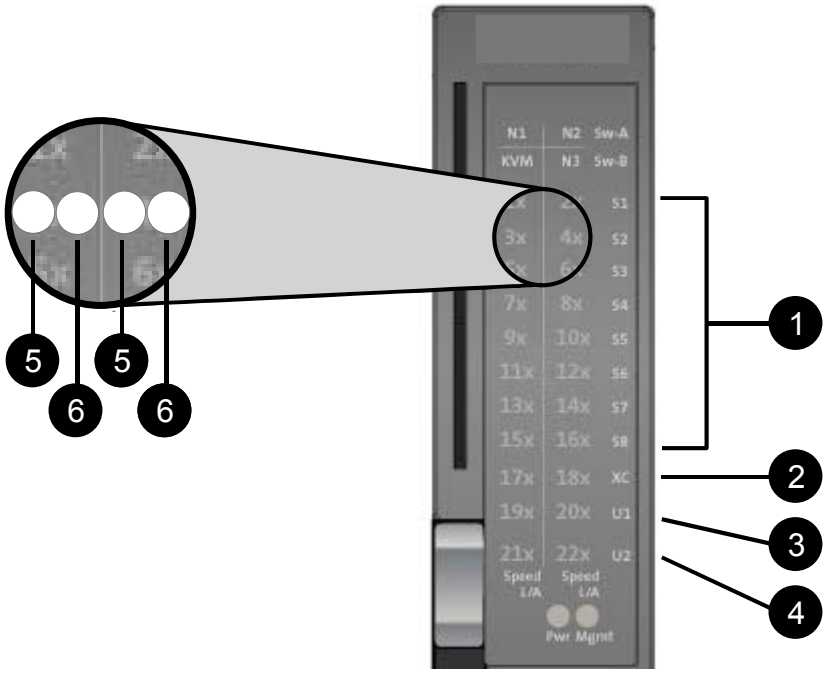

legend

item	description
1	server blade Ethernet port link and activity speed (ports 1 – 16)
2	crosslink port link and activity speed (ports 17 – 18)
3	10/100T uplink port link and activity speed (ports 19 – 20)
4	10/100/1000T (or 1000SX) uplink port link and activity speed (ports 21 – 22)
5	link speed LED: amber = 1000 Mb/s green = 100 Mb/s off = 10 Mb/s
6	link activity LED: green = link and no activity green flashing = link and activity amber = port disabled off = no link

Note: Uplink ports on the p-Class C-GbE Interconnect Kit are 10/100/1000T. Uplink ports on the p-Class F-GbE Interconnect Kit are 1000SX.

figure 16. Diagnostic Station

A Diagnostic Station option is available to diagnose a server blade or interconnect switch removed from the rack environment (Figure 16). Removing the server blade and interconnect switch from the server enclosure also removes them from their power source. The Diagnostic Station can be used to power up and build the initial interconnect switch configuration or to download a switch configuration file, test network adapter and interconnect switch port activity, diagnose interconnect switch status, and test an interconnect switch after installing an option or upgrade.

As previously mentioned, the Web browser-based user interface provides graphical and tabular monitoring screens of switch port utilization, data packets received and transmitted, port error packets, packet size, trunk utilization, and more.

Additional diagnostic features include support for port mirroring and a power-on self test (POST) built into the interconnect switch boot-up process for hardware verification. Port mirroring is a diagnostic and debugging tool to monitor network traffic on an interconnect switch port by sending a copy of the data to a different port that is connected to a network analyzer.

For more detailed information on the administration capabilities of the ProLiant BL p-Class GbE Interconnect Switch, see the *ProLiant BL p-Class GbE Interconnect Switch User Guide* available online at http://wwss1pro.compaq.com/support/reference_library/viewdocument.asp?source=263680-001.xml&dt=264.

network features and protocol support

The GbE Interconnect Switch options use industry-standard managed layer 2 Ethernet switches compatible with other industry-standard switches. They offer many enterprise-level Ethernet network features and support several industry standards.

enterprise class performance

The interconnect switch includes the following:

- nonblocking, full wire speed on all ports
- 6.4 Gbps switching fabric (3.2 Gbps per switch) full duplex
- frame forwarding throughput of 1.488 million packets per second (per gigabit port) and 0.1488 million packets per second (per 10/100 port) (64 byte packets)
- 8,191 MAC addresses per switch with MAC address based filtering and MAC address packet priority configurable on a per port and per VLAN basis, with eight levels of priority
- forwarding MAC address age, time settable to any value from 1 to 1,000,000 seconds

- 32 MB main memory, 8 MB flash memory, and 16 MB packet buffer memory per switch

IEEE 802.1Q virtual local area network (VLAN)

Each switch supports 255 port-based IEEE 802.1Q tagged VLANs with GVRP dynamic VLAN registration.[2] VLANs simplify administration and planning, enhance performance by conserving bandwidth, and improve security by limiting traffic to specific domains.

IEEE 802.1D spanning tree protocol

The support for IEEE 802.1D spanning tree protocol (STP) allows the blocking of links that form loops between switches in a network. When multiple links between switches are detected, a primary link is established. Duplicated links are blocked from use and become standby links. If the primary link fails, the standby link is activated. One instance of STP is supported per each switch, also known as mono-spanning tree. STP switch parameters can be configured at the switch and port level. Each switch can automatically find the STP root bridge on the network; if not, it will act as the root bridge for the STP domain. The bridge priority is configurable. Also, the switch supports STP fast forward or bypass mode, allowing a port to skip the standard STP modes and enable itself directly in the forwarding state.

IEEE 802.1p quality of service (QoS) packet prioritization

Quality of Service (QoS) for IEEE 802.1p packet prioritization allows switch administrators to set priority levels on each switch for forwarding packets. Each switch supports four classes of traffic (buffers or queues) for implementing priority based on the priority tag of the packet. The interconnect switch allows administrators to map eight priority levels to four classes. Traffic from a specific server port can be given priority over packets from other devices according to this range of priority levels. For example, with multiple packets in a buffer, the packet with the highest priority would be forwarded first, regardless of when it was received.

IEEE 802.3x flow control

The support for flow control enables each switch to signal network components to slow down or stop transmission of data during periods of peak network traffic. This provides a means to overcome switch overload by smoothing out traffic flow and eliminating risk of lost data.

port mirroring

The interconnect switch allows the user to mirror a switch port to another port for network monitoring and troubleshooting purposes. The desired type of data frames (egress, ingress, or both) may be mirrored. By providing a copy of traffic passing through the monitored port, this technology offers a way for network packet analyzers to view the traffic moving through a switch port.

EtherChannel compatible trunking

The interconnect switch IEEE 802.3ad static port trunking feature allows several ports to be grouped together and to act as a single logical link called a trunk. This feature provides a bandwidth that is a multiple of a single link's bandwidth. It also improves reliability, since a configurable way of load balancing is automatically applied to the ports in the trunked group. A link failure within the group causes the network traffic to be directed to the remaining links in the group. Six multi-link trunks (with up to eight ports per trunk) are supported per switch. The GbE Interconnect Switch trunking implementation is compatible with Cisco EtherChannel link aggregation (Fast EtherChannel, Gigabit EtherChannel).

[2] GVRP stands for GARP VLAN Registration Protocol.

internet group management protocol (IGMP) snooping v1 and v2

For efficient use of bandwidth, IGMP snooping manages multicast traffic in the interconnect switch by allowing directed switching of the IP multicast traffic. The interconnect switch can use IGMP snooping to configure switch ports dynamically, so that IP multicast traffic is forwarded only to those ports associated with IP multicast hosts.

IGMP snooping allows the switch to recognize IGMP queries and reports sent between network stations or devices and an IGMP host that belong to a specific multicast group. When enabled for IGMP snooping, the switch can open or close a port to a specific device based on IGMP messages passing through the switch. This feature further limits unnecessary broadcasts.

broadcast, multicast, and DA-storm prevention

The GbE Interconnect Switch allows thresholds (in packets per second) to prevent three types of packet storms: broadcast, multicast, and destination address (DA) unknown. If the threshold is exceeded, any additional packets received are dropped.

static unicast and multicast packet filtering

The switch uses a filtering database to segment the network, control communications between segments, and provide intrusion control. The switch allows manual entry of specific MAC addresses to be filtered from the network. Filtering of both unicast and multicast traffic is possible.

store and forward switching scheme

The interconnect switch provides a store and forward switching scheme that allows each packet to be buffered (stored) before it is forwarded to its destination. This scheme improves reliability in networks with high traffic. Packets that cannot be forwarded immediately saved, rather than dropped, and are forwarded at a later time.

simple network time protocol (SNTP)

Support for simple network time protocol (SNTP) allows the switch to obtain the accurate time through the network. When SNTP is enabled, the interconnect switch sends a request to a primary SNTP server in each period of a specified polling interval asking for the Greenwich Mean Time (GMT). If the primary SNTP server is not available, the request is sent to a secondary SNTP server.

XML formatted configuration file

The GbE Interconnect Switch firmware configuration file is specified in extensible markup language (XML). The XML configuration files are human readable and may be directly edited to create customized configurations.

DHCP and BOOTP protocols

An interconnect switch can be configured to obtain an IP Address automatically from a DHCP or BOOTP server.[3] The IP settings can be manually configured through the command line and menu-driven console interfaces. The IP settings are also configurable from the Web interface; but since the connection is based on an IP address for this interface, users would have to reconnect with the newly assigned IP address.

auto-negotiation of speed and duplex mode

The interconnect switch automatically identifies and matches its port speed and duplex mode of an attached device. This feature simplifies deployment in mixed 10, 100, and 1000 Mbps networks. Each port may also be enabled or disabled with manual setting of speed and duplex mode. The GbE Interconnect Switch further includes a bandwidth limitation feature that restricts the ingress (receiving) and egress (transmitting) packet rate configurable for each port. If the packet rate exceeds the allowed bandwidth rate, the excess packets will be dropped.

[3] DHCP stands for Dynamic Host Configuration Protocol. BOOTP stands for Bootstrap Protocol.

network adapter teaming

When teaming network adapters, switch considerations are the same with the interconnect switch options as with any other Ethernet switch. The ProLiant p-Class servers include standard HP NC series network adapters that support all three types of teaming:

- Network Fault Tolerance (NFT)
- Transmit Load Balancing (TLB)
- Switch-assisted Load Balancing (SLB), formerly known as Fast EtherChannel/ Gigabit EtherChannel (FEC/GEC) teaming

The three NC series network adapter signals per server blade bay are routed to the two interconnect switches for redundancy (two to one switch, one to the other), see Figure 3. Both NFT and TLB support teamed network adapters that are routed to separate switches, so both these types of teaming can be implemented on all three of the NC series network adapters. SLB teaming requires that the switch support trunking (link aggregation) and that all network adapters in the team be connected to the same switch. The interconnect switch does support trunking. However, only two of the three NC series network adapters per server bay are routed to the same switch. Therefore, these two network adapters can be set up as a SLB team.

For additional information about teaming ProLiant network adapters, see the teaming overview and white papers located at http://www.compaq.com/products/servers/networking/teaming.html.

ordering the p-Class GbE Interconnect Kits

Each ProLiant BL p-Class GbE Interconnect Kit includes the components necessary for one server blade enclosure. Order the p-Class C-GbE Kit (SKU 249655-B21) for copper-based Gigabit Ethernet networks and the F-GbE Kit (SKU 279720-B21) for fiber-based installations.

conclusion

The ProLiant BL p-Class system is a comprehensive solution for meeting the needs of customers operating in a multi-tiered environment that demands space and power efficiency, on-demand computing, end-to end system redundancy, rapid hardware and software deployment, and scalable management. The ProLiant p-Class GbE Interconnect Kits are available as one of two interconnect options for the ProLiant BL p-Class system. Each GbE Interconnect Kit includes a pair of fully redundant, managed layer 2 Ethernet switches that consolidate up to thirty-two 10/100 Mbps network adapters to one to eight uplinks. These options are ideal for reducing Ethernet network cabling and the time required to deploy, manage, and service ProLiant BL p-Class systems.

feedback

Please direct comments regarding this communication to the ISS Technology Communications Group at this Internet address: TechCom@HP.com

03/2003

Document Number 16VZ-1102A-WWEN

appendix: features of the ProLiant BL p-Class GbE Interconnect Switch

switch performance

- non-blocking full wire speed architecture, store and forward mode layer 2 switching
- frame forwarding throughput of 1.488 million packets per second (per Gigabit port) and 0.1488 million packets per second (per 10/100 port) (64 byte packets)
- 4.0 Gbps bandwidth per switch (8.0 Gbps per server blade enclosure) (full duplex)
- 32MB main, 8MB flash, and 16MB packet buffer memory per switch (packet buffer memory shared between ports)

switch network features

- IEEE 802.3 10Base-T Ethernet, IEEE 802.3u 100Base-TX Ethernet, IEEE 802.3ab 1000Base-T Ethernet, and IEEE 802.3z 1000Base-SX
- IEEE 802.1D Spanning Tree Protocol, with ability to set STP in a fast forwarding mode on a per port basis (Cisco PortFast equivalent)
- IEEE 802.3ad link aggregation (excluding LACP) supporting up to six multilink trunk groups with eight ports per group (compatible with Cisco EtherChannel trunking
- IEEE 802.1p Quality of Service (QoS) packet prioritization with 4 classes of service mapped to 8 priority levels
- IEEE 802.1Q port based tagged VLANs and IEEE 802.3ac VLAN Ethernet extensions
- GARP VLAN Registration Protocol (GVRP) providing 802.1Q compliant VLAN pruning and dynamic VLAN creation
- IEEE 802.3x Flow Control with manual configuration capability
- full wire speed automatic MAC address learning, 8,191 MAC addresses per switch
- static MAC addresses manageable on a per port and per VLAN basis
- source and destination MAC address based prioritization of packets per VLAN with eight levels of priority
- MAC address aging time settable to any value from 1 to 1,000,000 seconds
- IGMP snooping v1 and v2, also configurable to a non-querier mode
- broadcast, multicast, and unknown packet storm control with a configurable threshold value
- 191 concurrent multicast groups (127 dynamically learned by IGMP, 64 static load balancing of both unicast and multicast traffic
- unicast and multicast static MAC address packet filtering
- auto-negotiation and auto-sensing with full-duplex support
- enable/disable ports with ability to manually force port speed and duplex mode on a per port basis
- auto-MDI/MDIX on all ports
- SNTP with support for a primary and secondary network time server and automatic daylight savings adjustment
- ping capability to test the connectivity on the Ethernet network
- ARP for IP to MAC address resolution

switch deployment and administration

- fully featured Web browser-based interface
- fully featured menu driven console interface
- command line interface with scripting capability
- Telnet-based access to the command line and menu-driven console interfaces
- serial line IP (SLIP)/RS-232 based access to the command line and menu-driven console interfaces
- SNMP-based scripting with best-case HP recommended example scripts
- SNMP v1 and RMON groups 1 (statistics), 2 (history), 3 (alarms), and 9 (events)
- Four configurable community strings for access via SNMP and SNMP trap managers
- MIB-II, Bridge MIB, Interface MIB, 802.1Q MIB, RMON1 MIB groups 1, 2, 3, and 9, Ethernet MIB, and HP enterprise switch MIBs
- Ability to manage both switches from a single Ethernet port
- ability to communicate to any and all server blade network adapters from any Ethernet uplink port
- manual, or automatic IP settings via BOOTP or DHCP server
- TFTP client to upgrade the switch firmware; to save the switch log file; and to save, restore, and update the switch configuration file
- Zmodem switch firmware transfer from the serial interface
- per port bandwidth control of ingress and egress traffic
- port mirroring with ability to mirror desired type of frames (egress, ingress, or both)
- ability to name ports on a per port basis
- human read/write XML-based configuration file
- power on self test (POST) at boot for hardware verification
- fully pre-configured for immediate plug-in operation in the server blade enclosure
- front panel system and per port speed and link activity LED annunciation panel per switch blade
- per port speed and link activity LEDs adjacent to all external Ethernet ports
- Web-based interface active virtual graphic of the front of each interconnect switch
- monitoring of port utilization, data packets received and transmitted, port error packets, packet size, trunk utilization, and more with both graphical and tabular displays

switch security

- password protected multi-level user accounts supported on all management interfaces
- configurable idle time-out period for web, Telnet, and console interfaces
- 255 port-based IEEE 802.1Q tagged VLANs per switch (510 per server blade enclosure)
- unicast and multicast static MAC address packet filtering
- restriction on the maximum number of MAC addresses learned on a per port basis
- ability to specify the IP-based management stations that are allowed to access each switch

switch availability

- front mounted hot-pluggable and hot-removable server blade, switch blade removed without need to cable/recable (p-Class)
- redundant switches per server blade enclosure
- redundant uplinks per switch
- redundant 10/100T inter-switch crosslinks for switch to switch communication and failover scenarios within the server blade enclosure
- server network adapters routed to different switches
- on-board redundant cooling per each
- redundant serial ports
- load balancing of unicast and multicast traffic
- IEEE 802.1D Spanning Tree Protocol
- IEEE 802.3ad multi-link load balancing (excluding LACP)
- Support for ProLiant network adapter teaming

switch port and power

p-Class C-GbE Kit

- four 10/100/1000T and four 10/100T external Ethernet ports on hot pluggable rear-mounted module
- four 10/100T external Ethernet ports on switch front panels
- all external Ethernet ports may be used for data, switch and iLO management, and/or PXE remote configuration
- 2 external front panel serial ports (one per switch)
- 32 internal 10/100T Ethernet ports to server blade network adapters (4 per server bay)
- 2 internal 10/100T Ethernet inter-switch crosslinks for switch to switch communication and failover scenarios
- I2C switch to management module communications

p-Class F-GbE Kit

- four 1000SX and four 10/100T external Ethernet ports on hot pluggable rear-mounted module
- four 10/100T external Ethernet ports on switch front panels
- all external Ethernet ports may be used for data, switch and iLO management, and/or PXE remote configuration
- 2 external front panel serial ports (one per switch)
- 32 internal 10/100T Ethernet ports to server blade network adapters (four per server bay)
- 2 internal 10/100T Ethernet inter-switch crosslinks for switch to switch communication and failover scenarios
- I2C switch to management module communications