


ProCurve Networking

ProCurve Switch 8212zl Series Reviewer's Guide V2.0a



Reviewer's Guide

Reflects Version K.13 updates



Table of Contents

Executive Summary.....	4
Introduction	4
ProCurve Switch 8212zl (J8715A)	4
Audience.....	4
Scope.....	4
Product Positioning.....	4
Overview	4
ProCurve Switch 8212zl – ProCurve Core/Distribution.....	5
ProCurve Switch 8212zl	6
Switch 8212zl Chassis Layout	7
Switch 8200zl Series Specific Modules and Components.....	7
8200zl Management Module (J9092A)	7
8200zl Fabric Modules (J9093A)	9
8200zl System Support Module (J9095A).....	9
Switch 8212zl Fan Tray (J9094A)	11
Common Switch zl Series Modules and Power Supplies	12
zl Series Interface Modules	12
zl Power Supplies	16
zl Power Supply Configuration Planning for PoE	17
Switch 8200zl Operational Characteristics	18
Management Module Failover	18
Updating SW code – minimizing downtime	18
ProVision ASIC Architecture	19
Inside the ProVision ASIC Architecture	20
Classification and Lookup	20
Policy Enforcement Engine.....	21
Network Switch Engine Programmability.....	21
Fabric Interface	21
ProVision ASIC CPU	22
Fabric ASIC	22
Management Subsystem	22
Advanced Capabilities of the Product Family	22
Versatile Intelligent Port (VIP).....	22
ProVision Hardware Resiliency.....	22
Overview of Features and Benefits	22
Feature Set Summary.....	25
Security Features	25
QoS Functions	26
Convergence	26
Layer 2 switching	26
Bridging Protocols	27
Routing Protocols	27
IPv6.....	27
Multicast Protocols	27
High Availability and Redundancy	27
Management	27
Standards and Protocols.....	28
Performance and Capacity	29
Capacity and Performance Features Comparison	29

Per-Port Buffer Sizes	29
Routing and Forwarding Tables	30
Optimizing the 10-GbE Port Configuration	30
Throughput and Latency Performance Data	32
10 Gigabit Performance Traffic Patterns	33
Throughput Test	34
Latency Measurements	34
Comparison of 8200zl/5400zl with 5300xl and 8100fl	34
Services and Support	36
Lifetime Software Updates (Best in the Industry)	36
Lifetime Warranty (Best in the Industry)	36
Free Telephone Support	36
Optional Support Services	36
Appendix A: Policy Enforcement Engine	37
Policy Enforcement Engine Benefits	37
Wirespeed Performance for ACLs	37
Appendix B: Power over Ethernet (PoE)	38
PoE device types	38
Power delivery options	38
PoE negotiation	39
Additional PoE Power – External Supplies	39
Support for Pre-802.3af Standard Powered Devices	39
Appendix C: PIM-Sparse Mode	39
Appendix D: LLDP-MED	40
Appendix E: Virus Throttle Security	42
Response options	42
Sensitivity	43
Connection-rate ACL	43
Appendix F: VRRP	43
XRRP support on 5300xl switch	44
Appendix G: OSPF Equal Cost Multipath	44
Appendix H: Troubleshooting	45
LED Status Indicators for 8200zl	45
Appendix I: Part Numbers and Field Replaceable Units	48

Executive Summary

ProCurve Networking has an extensive line of networking products built around the concept of the ProCurve Adaptive EDGE Architecture, which provides the security, mobility and convergence capabilities businesses demand, while giving IT administrators the ability to adapt to the changing needs of their organizations and control their infrastructure centrally.

This guide describes the ProCurve Switch 8212zl, which builds upon the ProCurve Switch 5412zl and the principles of the Adaptive EDGE architecture, bringing intelligence and performance to the network in a cost-effective way, and providing a platform for delivering the full Intelligent Edge capability using this newest of Interconnect Fabric Switches in the core of your network. The foundation for this switch is a purpose-built ProVision™ application-specific integrated circuit (ASIC) that allows the most demanding networking features, such as Quality of Service (QoS) and security, to be implemented in a scalable yet granular fashion. The Switch 8212zl offers a passive backplane and system architecture for high performance, system redundancy, 10-GbE capability, and programmable ASICs for simple feature upgrades through software, providing your network with excellent investment protection, flexibility and scalability.

The ProCurve Switch 8212zl is the industry's first core networking switch with a lifetime warranty and free software upgrades for whoever owns the switch. High-availability design, wirespeed performance, scalability, and future upgradeability make the Switch 8212zl a superior value proposition consistent with what ProCurve Networking has been delivering for more than 25 years.

Introduction

ProCurve Switch 8212zl (J8715A)

The ProCurve Switch 8212zl is a 9U switch (15.75 inch/40 cm), designed with high-availability features, dual redundant management modules, dual resilient fabric modules and redundant power supplies. The chassis can hold up to 12 ProCurve “zl” interface modules designed to support up to either 288 Gigabit or 48 10-Gigabit ports. The “zl” interface modules are the very same modules supported in the ProCurve 5400zl Series. The throughput of the 8212zl switch fabric is capable of up to 428 Mpps, with a backplane speed of 692 Gbps.

Audience

This guide is written primarily for technical evaluators and product reviewers of networking equipment and solutions.

Scope

This guide provides detailed information and specifications about the ProCurve Switch 8212zl product, with the assumption that details about standard protocols can be referenced externally by those familiar to general networking. Relatively new technologies will be covered in more detail than older, more broadly familiar technologies.

References to the ProCurve Switch 5400zl, 3500yl and 6200yl Series are used as they relate to the architecture and positioning of the Switch 8212zl Series.

For more information regarding the complete line of ProCurve Networking products, go to www.procurve.com.

Product Positioning

Overview

A widely used method for segmenting the areas in which switches are installed calls for three different classifications: access, distribution and core. Access switches provide aggregation of end nodes for connection to a distribution or core switch and usually are found in wiring closets. Distribution switches aggregate the links from access switches and, possibly, server farms. Distribution switches anchor the network in a building or, for small networks, across a campus. Core switches provide the focal point of the local network, aggregating the distribution switches, and providing connectivity for central site data centers, and often, the external network.

Regardless of a network's architecture – whether it's ProCurve's Adaptive EDGE architecture with its intelligence at the edge, or a traditional architecture that concentrates intelligence at the center of the network– there is a choice of wiring topologies to implement that architecture.

The intelligence and throughput of the ProCurve Switch 5400zl and 3500yl Series make them suitable for applications at or near the edge of networks of all sizes, or in the center of networks that are small-to-moderate in size. The Switch 8212zl offers the same ProVision ASIC intelligence and throughput with a redundant management module design and flexible switch fabric modules for demanding high-availability applications, whether at the core, distribution or edge of a network. The modular structure of both the ProCurve Switch 8212zl and 5400zl Series provides flexibility for tailoring the switch to match topology requirements. The ProCurve Switch 6200yl, an advanced Layer 3 stackable, is designed to be deployed as an aggregator of traffic from the edge to the core of the network.

The foundation for all of these switches is a purpose-built, programmable ProVision ASIC that allows the most demanding networking features, such as QoS and security, to be implemented in a scalable yet granular fashion. The 8212zl is a truly passive backplane design. The high-capacity switch fabric is designed for front-slot modularity on the Switch 8212zl, while the switch fabric is integrated on the backplane on the 5400zl/3500yl/6200yl products. The ProCurve Switch 8212zl and 5400zl/3500yl/6200yl Series have been designed as a product family, using the ProVision ASICs and software, and providing consistency and scalability across the family.

The ProVision ASICs are aimed at accomplishing several objectives:

- Providing a great engineering balance between feature capabilities, performance and price
- Bringing sophisticated control features to the edge of the network
- Having programmable features that allow future requirements to be accommodated in the ASICs

Key features of the ProCurve Switch 8212zl, 5400zl, 3500yl and 6200yl Series include:

- **Performance** — High-capacity switch fabric (from 115 Gbps to 692 Gbps backplane speed), bandwidth shaping and control, QoS, Layer 2 and Layer 3 jumbo frames
- **Security** — ACLs (per-port or identity-driven), virus throttle, switch CPU protection, detection of malicious attacks, DHCP protection, BPDU port protection, Dynamic ARP protection, Dynamic IP lockdown, STP root guard, IP and MAC lockdown/lockout, 802.1X, Web, and MAC user authentication, USB secure autorun, management access control (SSH, SSL, TACACS+, secure FTP)
- **Operational flexibility** — High-port density in 4U and 7U form factors, versatile intelligent ports (10/100/1000, Power over Ethernet (PoE)-enabled), power supply choices for optimum PoE
- **Resiliency** — Redundant power supplies, hot-swappable/hot-insertable modules, MSTP, switch meshing, VRRP, OSPF-ECMP
- **Layer 2 features** — GVRP, VLAN, Q-in-Q
- **IP Routing** — RIPv1, RIPv2, OSPF, PIM-SM, PIM-DM, static routes
- **IPv6** — IPv6 host, IPv4/IPv6 dual stack, MLD snooping
- **Convergence** — IP multicast snooping (data-driven IGMP), LLDP-MED, RADIUS VLAN, PoE
- **Diagnostic** — Remote intelligent mirroring, loopback interface, UDLD, sFlow support
- **Investment protection** — Upgradeable management engine and CPU, upgradeable to Premium License feature group, add-in modules and power supplies, versatile intelligent ports, programmable ASICs

ProCurve Switch 8212zl – ProCurve Core/Distribution

The ProCurve Switch 8212zl is one of the most advanced Layer 3/Layer 4 switches in the ProCurve Networking product line. The Switch 8212zl is a fully *passive backplane* 12-slot chassis and uses the same zl modules as the Switch 5400zl Series. With a variety of Gigabit interfaces, integrated PoE on all 10/100/1000Base-T ports, 10-GbE capability, the Switch 8212zl offers

excellent investment protection, flexibility and scalability, as well as ease of deployment, operation and maintenance.

The 8212zl is targeted as an enterprise-class, high-availability, distribution and medium-scale core switch, while the 5400zl and 3500yl switches are targeted as enterprise-class wiring closet switches — designed for low cost, with a choice of medium to high port density. Voice, video and data ready, the ProCurve Switch 8212zl, 5400zl and 3500yl/6200yl Series all offer extensive prioritization features that bring full convergence down to the desktop. Integrated PoE minimizes wiring requirements for Voice over Internet Protocol (VoIP) phones and wireless access points.

The 5400zl and 3500yl products ship with a standard offering of the Intelligent Edge feature group, and an optional Premium License feature group includes protocols such as PIM-SM, PIM-DM, OSPF-ECMP and VRRP. (RIP routing protocol and static routing features are included in the standard Intelligent Edge feature group).

The Switch 8212zl has OSPF-ECMP, VRRP, PIM-DM and PIM-SM features already enabled, making it suitable for distribution or medium-scale core switch applications, where routing and default-gateway use requires a higher degree of availability with the dual management, resilient fabric modules and redundant power supplies offered by the 8212zl. The passive backplane design and removable switch fabric modules make future upgrades to the Switch 8212zl easy.

ProCurve Switch 8212zl

The high-availability design of the ProCurve Switch 8212zl with its dual management and fabric modules, along with ProCurve's auto-synchronizing capability, positions this switch perfectly as a distribution or aggregator switch; it is even suitable for medium-scale core applications. It shares the same software and hardware capabilities with the 5400zl, 3500yl and 6200yl Series.

The Switch 8212zl enables the use of core features, such as dynamic routing protocols OSPF-ECMP and VRRP, as well as PIM-DM and PIM-SM for multicast routing support.

The 8212zl follows in the evolution of ProCurve-managed chassis switches, featuring PoE on every copper port, chassis flexibility, and a scalable architecture, with high availability designed in for critical network operations.

The 8212zl requires a single management module, which initializes, controls and monitors the various interface and fabric modules. The 8212zl has 12 slots available for population with interface modules (the management and fabric modules have their own dedicated slots and are not counted as part of the 12 slots). The 8212zl shares the same interface modules used with the ProCurve 5406zl and 5412zl switches. These interface modules are the 24-port 10/100/1000, 20-port 10/100/1000 + 4-port mini-GBIC, 24-port mini-GBIC, 4-port 10-GbE X2, and 4-port 10-GbE CX4. The 8212zl supports up to four internal power supplies (two are required for 12-slot operation – one for only six slots). The power supplies can be used for purposes of power redundancy or for additional PoE requirements.

The base configuration for the 8212zl (J8715A) as shipped from the factory includes:

- 12-slot 8212zl Chassis (J9091A), including a system Fan Tray (J9094A)
- One (1) 8200zl Management Module (J9092A), with an available slot for a second redundant management module
- Two (2) 8200zl Fabric Modules (J9093A)
- 8200zl System Support Module (J9095A)
- Dynamic routing protocols OSPF, VRRP, PIM-DM and PIM-SM are standard

Note that NO power supplies ship with the 8212zl base chassis: They must be ordered separately. There are two types available with different PoE power capacities (1500W or 875W). A minimum of two power supplies are required for proper operation of all 12 slots in the 8212zl. Up to four power supplies can be installed in the 8212zl, providing N+N power redundancy, which is useful in networks that provide two separate AC power feeds. In the event of a failure down to one single power supply, the 8212zl will continue to operate, but only the top six slots (labeled A through F) will remain powered on, so critical network connections should be prioritized toward the top of the switch.

Switch 8212zl Chassis Layout

The Switch 8212zl is a rack-mountable 9U height chassis. The interface modules are inserted in the front slots, labeled A through L. Management and fabric modules are labeled MM1/MM2 and FM1/FM2 respectively.

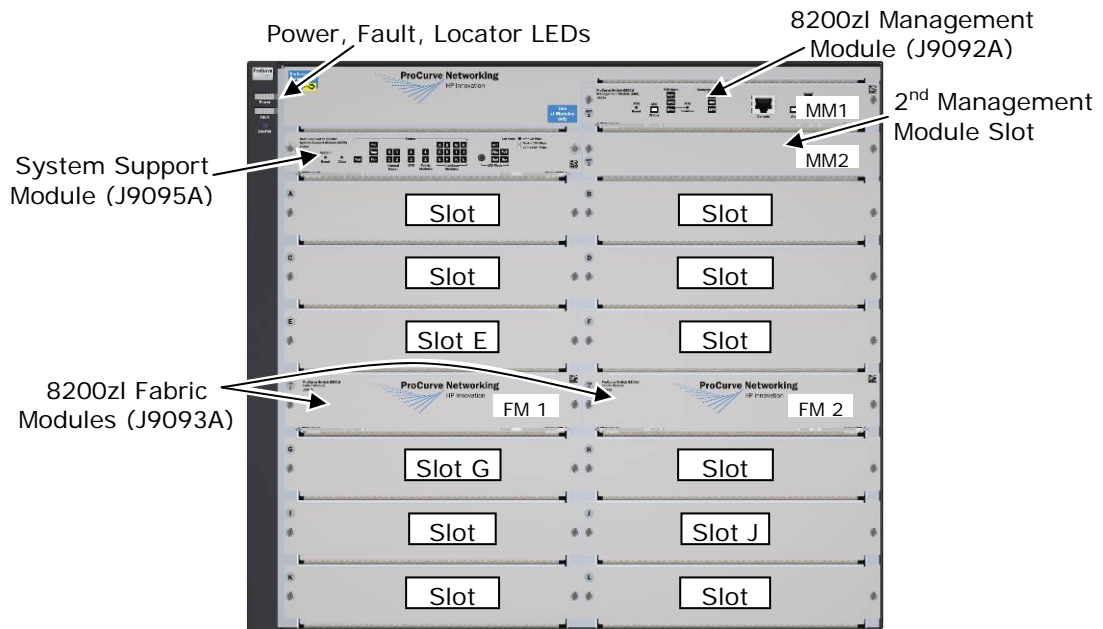



Figure 1: ProCurve Switch 8212zl Base System (J8715A) slot labeling (MM = Management Module; FM = Fabric Module; all others = Interface Module)

The 8200zl Management Modules are hot-swappable, i.e., the switch does not have to be powered off to remove either management module. The dual management module (MM) design allows either module to manage the system, in an active/standby model. By default, MM1 will assume the “active” role, but in general, whichever MM was the last one to be “active,” the other MM will assume the “standby” role. All operations are performed through the “active” MM (either MM1 or MM2). If the active MM were to fail, the standby MM takes control of the switch and continues operation. However, if only one management module is running and then is removed, all ports lose communication and the system will be powered down.

Switch 8200zl Series Specific Modules and Components

The Switch 8200zl Series shares interface modules and power supplies with the Switch 5400zl Series. This next section will cover details of modules and components specific to the 8200zl Series.

8200zl Management Module (J9092A)



The ProCurve Switch 8200zl Series requires at least a single management module that oversees (or “supervises”) the operation of the interface modules and fabric modules. The management module is responsible for network control processing (e.g., OSPF updates or ARP requests), while each interface module in conjunction with the fabric modules handles the traffic switching in ASIC hardware.

The use of dual management modules by default synchronizes configuration information and code images automatically for the user. There is no need for a “synchronize” command. When configuration changes are written to flash memory (“write memory”) or SW updates are performed on the Active Module (a tftp copy of newer SW into Flash), they are automatically copied over to the Standby Management Module.




Figure 2: Close-up of the Management Module LEDs in a typical operational state

The LEDs on the 8200zl Management Module are grouped into two columns:

- One set to indicate the management module’s “state” (Active, Standby or Down)
- One set to indicate the status of components (CompactFlash and DIMM system memory) and communication status with the System Support Module (SSM)

The “MM Status” LED indicates general health of the management module, indicated by a green color after the module has passed power on self-test.

The “MM Reset” is a recessed button used to manually reset the management module.

The management modules are designed to be hot-swappable, and can be removed from the chassis without damage. The synchronization of files (configuration, code images, state and default condition directives) may indeed be occurring, so to minimize the possibility of corruption between MMs, when manually removing the module, use the “MM Shutdown” button. The “MM Shutdown” is a recessed button used to gracefully shutdown the management module, completing any synchronization of files and stating information to the second management module. When the “Down” LED is lit, the management module can be removed.

The LED indicators are covered in more detail later in this document in Appendix H: Troubleshooting.

Processor

The CPU processor is a Freescale PowerPC 8540 operating at 667 MHz.

Memory

SDRAM

Synchronous Dynamic RAM is used for the storage of uncompressed executable code and data structures. The SDRAM consists of a 256 MB DDR-1 DIMM in the base module, expandable up to 1GB. The DDR-1 interface is 64-bits running at 166 MHz bus speed (333 MHz data rate).

FLASH

The flash consists of a 128 MB CompactFlash expandable up to 1GB, and a 4 MB mirror-bit flash. The mirror-bit flash is used for initial boot code. The CompactFlash is used for nonvolatile configuration storage, and compressed code storage. The CompactFlash is socketed for future upgrade capability.

The CompactFlash may be programmed in a bulk fashion or one sector at a time. Since all application code is executed out of SDRAM, the CompactFlash may be programmed while the switch is operational; that is, the new code can be downloaded onto the CompactFlash during system operation. The CompactFlash is sized such that a backup copy of an older revision of application code also may be stored. The system also allows up to three copies of configuration files to be held, associating them to a particular flash image (primary => Config1, secondary => Config2, Active Running session => Config3).


Console Port

The management module incorporates one RS-232 serial port for local management and configuration. This port uses an RJ-45 connector mounted on the front panel. To connect to the console, an RJ45-to-DB9 cable is provided with each switch and is similar to the “rollover” cable used on Cisco products.

Auxiliary Port

The management module includes a USB auxiliary port used for offline data transfer of files without the need for a network to be set up. System code can be copied to a USB memory stick, and the system can copy this image from USB, just as it would be over the network.

8200zl Fabric Modules (J9093A)



The switching fabric is provided by two fabric modules, each housing a single ProVision Fabric ASIC. Much like the design of the award-winning ProCurve 5412zl switch, the Fabric ASICs are load-sharing, and resilient to each other. With both fabric ASICs operational, all 12 slots of the 8212zl provide wirespeed switching capacity for up to 288 gigabit ports, or up to 28.8 Gbps per slot when used with 10 Gig ports.

In the unlikely case of a failing ASIC on either of the fabric modules, the 8212zl resilient design continues to operate with all 12 slots interconnected by the remaining fabric module – albeit at half-capacity (14.4 Gbps capacity per slot) – until a repair cycle can be performed. Replacing a failed module is simply a matter of removing two screws and replacing the module. Sub-second recovery of the switching fabric is performed automatically.


8200zl System Support Module (J9095A)

The Switch 8212zl is a fully repairable core switch. All active components for the system are mounted on removable, quickly replaceable modules. Where the management and fabric components of the system are redundant and resilient to failure conditions, other components that normally would be designed into the backplane are instead designed into the System Support Module (SSM).

The SSM contains:

- System clock
- Multiplexer circuitry for management module to interface module communication
- Fan controller circuitry

On a 5400zl Series chassis, these components are designed into the chassis backplane.



Rather than duplicate status information and control buttons on each of the dual management modules, the System Support Module creates a common location for:

- System “Reset” and “Clear” buttons,
- Status LEDs for internal power supplies, External Power Supply (EPS), fabric and interface modules

- LED Mode indicators (which dictate the function of the Mode LED on individual ports) to indicate activity (Act), full duplex operation (FDx), speed (Spd), delivery of PoE power, and a configurable User Mode (USR)¹

All components contained on the SSM are very low failure rate functions (as evidenced from the very low failure rate of ProCurve 5400zl Series chassis and predecessor chassis products). But in the event of a failure of any of these components, this module can be replaced without removing any network cables or interface modules as other competitor designs require when there is need to service some components mounted to the backplane.

Since the SSM is a required module for the operation of the 8212zl chassis, it is not a hot-swappable module. Although no damage will occur to the chassis if the SSM is removed during operation, its removal will indeed power down the system (similar to pulling a sole management module with no standby MM in place). To prevent the accidental removal of the module, the SSM is mounted to the chassis with Tamper-Resistant Screws (TRS). TRS also contribute to the low failure rate probability for the SSM – static discharge through user handling, or misalignment of pins upon reinsertion, just to name a few contributors to module failure.




Figure 3: Tamper-Resistant Torx - 20 screws and wrench (included with Spare SSM)

As with many critical components, a best practice would be to have a spare SSM on-site. ProCurve includes a TRS wrench with every separately purchased J9095A module, should a failed SSM ever need replacing. All other modules can be removed by using a blade screwdriver or T10 Torx driver.




Figure 4: Close-up of System Support Module (SSM) in a typical state

The above figure shows that PoE, temperature, and fan status are good (solid green) and that there are two power supplies installed, both fabric modules are in a good state, as are modules installed in Slots A through D and also, I through L. Modules E through H are empty. The Mode LEDs on ports indicate any network activity on any port that is linked.

The LED indicators are covered in more detail later in this document in Appendix H: Troubleshooting.

¹ User Mode (USR) will be enabled in a future software release.

Switch 8212zl Fan Tray (J9094A)

The fan tray assembly contains the cooling fans for the interior of the 8212zl chassis; the power supplies have their own internal cooling fans. The 8212zl fan tray consists of six variable speed fans – four to cool the modules and two to cool the SSM and management modules. The fan speed is based on the sensed ambient temperature of the chassis.

The fan tray is installed/replaced from the rear of the chassis, and mounted on the left side of the chassis (from a front-view perspective). The fans draw air through ventilation holes to the left, through the system, and to ventilation holes to the right (side-to-side airflow).




Figure 5: Airflow direction and rear view of 8212zl with four redundant power supplies installed

The fan tray is hot-swappable in the 8212zl; it can be removed and replaced without removing power from the switch. However, the new fan tray should be installed immediately after removing the old fan tray to avoid overheating and automatic switch shutdown within three minutes.

In the event of an individual fan failure, an SNMP trap and event log entry is generated. A system typically can operate for quite a long time with a single fan failure (out of the six), while the other fans step up in speed to compensate for the loss of airflow.

The rear panel of the fan tray replicates the Power/Fault/Locator LEDs found on the front of the 8212zl switch. This is useful when attempting to locate an 8212zl switch while walking behind a row of equipment cabinets.




Figure 6: Close-up view of rear-mounted Power, Fault and Locator indicator LEDs

The Locator function is enabled through the following CLI command:

```
ProCurve Switch 8200zl# chassislocate ?
  blink                Blink the chassis locate led (default 30 minutes).
  off                  Turn the chassis locate led off.
  on                   Turn the chassis locate led on (default 30 minutes).
ProCurve Switch 8200zl# chassislocate blink ?
<1-1440>              Number of minutes duration (default 30).
<cr>
```

By indicating a number N after either the “blink” or “on” parameter, the locator LED will extinguish automatically after N minutes. Without specifying a value, the default is 30 minutes.

Common Switch zl Series Modules and Power Supplies

The ProCurve Switch 8212zl shares interface modules with the ProCurve Switch 5400zl Series. The ProCurve Switch zl Series supports a variety of popular interface modules (line cards) providing customers with the ability to change or scale their LAN links and adapt as the needs of the business environment change over time.

zl Series Interface Modules

A variety of interface modules are available for ProCurve Switch zl Series:

- ProCurve Switch zl 24-port 10/100/1000 PoE Module (J8702A)
- ProCurve Switch zl 20-port 10/100/1000 + 4-port Mini-GBIC Module (J8705A)
- ProCurve Switch zl 24-port Mini-GBIC Module (J8706A)
- ProCurve Switch zl 4-port 10-GbE X2 Module (J8707A)
- ProCurve Switch zl 4-port 10-GbE CX4 Module (J8708A)

Each of the copper-based interface modules provides integrated PoE capability.

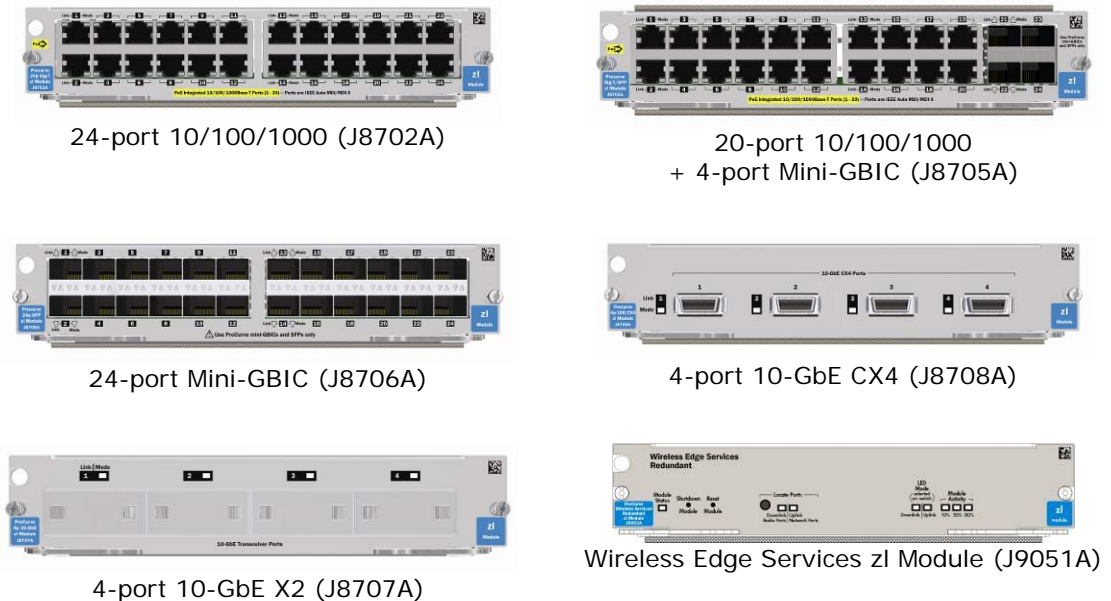



Figure 7: ProCurve Switch 5400zl Series Interface Modules

The following is a description of the interface modules supported by the ProCurve Switch 8212zl and 5400zl Series.

ProCurve Switch zI 24p 10/100/1000 PoE Module (J8702A)



Description

This interface module has 24 10/100/1000Base-T ports that provide gigabit-over-copper connectivity for wiring closets, enabling high-density gigabit connectivity to the desktop over Category 5 copper cabling. Each port is capable of providing IEEE 802.3af compliant PoE to power IP phones, wireless access points and other devices. Pre-standard powered devices also can be supported.


Ports

IEEE Auto-MDI: yes

Duplex: half or full

Connectors: RJ-45

ProCurve Switch zI 20-port 10/100/1000 + 4-port Mini-GBIC Module (J8705A)



Description

This interface module has 20 10/100/1000Base-T ports that provide gigabit-over-copper for wiring closets, enabling high-density gigabit connectivity to the desktop over Category 5 copper cabling. Each port is capable of providing IEEE 802.3af compliant PoE to power IP Phones, wireless access points and other devices.

In addition, this module provides four mini-GBIC ports for uplinks and intra-building connections. They can be trunked (aggregated links or “channels”) to provide up to 4 Gigabits of connectivity. Two such modules can have their ports trunked across both modules to provide module redundancy, as well.

Ports

4 open mini-GBIC slots

IEEE Auto-MDI: yes

Duplex: half or full


Connectors: RJ-45

Mini-GBICs supported (ordered separately) ²

² The Switch 8212zI supports mGBIC products from ProCurve with a “B” or later (“C” or “D”) designator as the last letter on the part number (e.g., J4858B or J4858C). mGBICs that are migrated from an older switch are not supported for operations in the Switch 8212zI .

- J4858B ProCurve Gigabit-SX-LC Mini-GBIC
- J4859B ProCurve Gigabit-LX-LC Mini-GBIC
- J4860B ProCurve Gigabit-LH-LC Mini-GBIC
- J8177B ProCurve Gigabit 1000Base-T Mini-GBIC

ProCurve Switch 5400zl 24-port Mini-GBIC Module (J8706A)



Description

This interface module has 24 mini-GBIC ports and is appropriate for use as an aggregator in a distribution environment. This module supports the same mini-GBICs as the ProCurve Switch zl 20-port 10/100/1000 + 4-port Mini-GBIC module previously described. The mini-GBICs are ordered separately.

Ports

24 open mini-GBIC slots

ProCurve Switch 5400zl 4-port 10-GbE X2 Module (J8707A)



Description

This interface module has four 10-GbE X2 Transceiver Ports that support any combination of SR, LR, ER, or CX4 transceiver types. This provides maximum flexibility for connecting 10-GbE high-speed downlinks to any other switch supporting that connection type. The wide variety of distances supported makes this module an ideal choice for intra-building connections. Ports can be trunked to provide higher throughput. Two such modules can have their ports trunked across both modules to provide module redundancy, as well.


Ports

4 open X2 transceiver slots

Transceivers supported (ordered separately)

- J8436A ProCurve 10-GbE X2-SC SR Optic
- J8437A ProCurve 10-GbE X2-SC LR Optic
- J8438A ProCurve 10-GbE X2-SC ER Optic
- J8440A ProCurve 10-GbE X2-CX4 Transceiver

ProCurve Switch 5400zl 4-port 10-GbE CX4 Module (J8708A)



Description

This interface module has 4 10-GbE CX4 ports.

Ports

4 10-Gigabit Ethernet ports (IEEE 802.3ak Type 10GBase-CX4)

Connectors: CX4

Transceivers supported (ordered separately)

- J8439A ProCurve 10-GbE CX4 Media Converter

Maximum Distance

15 m using CX4 cable

300 m using optical media converters and multimode fiber cable

Notes

Use CX4 10-GbE cable (0.5 m–15 m) or ProCurve 10-GbE CX4 Media Converter (J8439A).

Wireless Edge Services zl Module (J9051A)



Description

Working in conjunction with ProCurve Radio Ports, the ProCurve Wireless Edge Services zl Module provides centralized wireless LAN configuration and management of advanced wireless services, enabling a resilient, highly secure, mobile multi-service network. Each “WES” Module can control up to 156 ProCurve Radio Ports (light access points) to provide Layer 2/3 seamless roaming and a secure mobility environment.

Wireless sFlow support to provide leveraged network management for both wired and wireless connections.

Up to four modules can be installed in a single chassis, and a redundant WESM (J9052A) version automatically adopts radio ports if the primary module is unavailable or should fail.

ProCurve Radio Ports supported (ordered separately)

- J9004A ProCurve Radio Port 210 (single 802.11g radio)
Integrated antenna
- J9006A ProCurve Radio Port 230 (dual radio a+b/g)
Integrated antenna

- J9005A ProCurve Radio Port 220 (dual radio a+b/g)
Penum rated, external antennas required

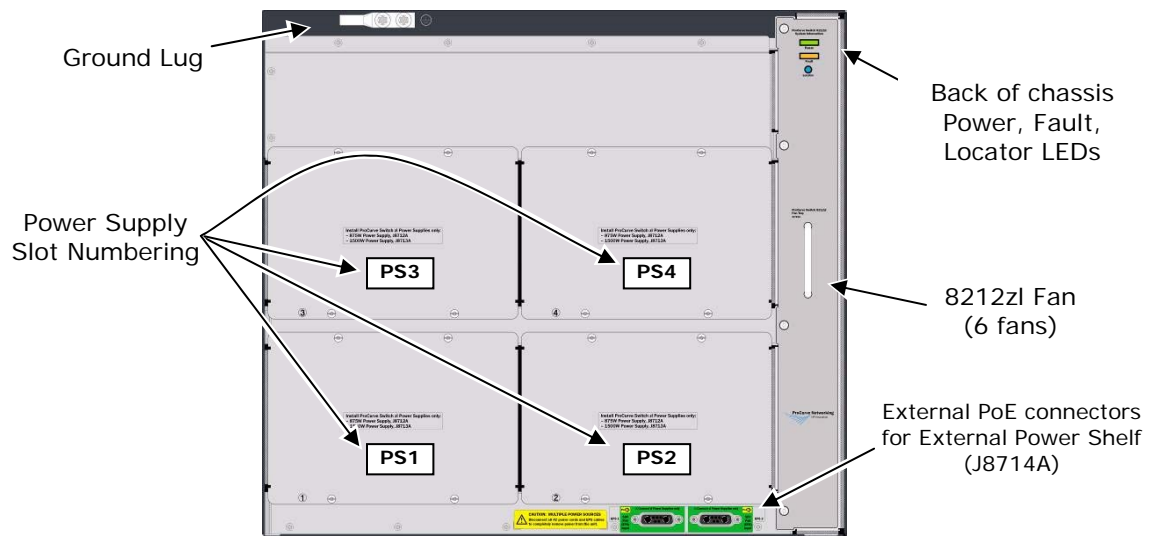


Figure 8: Rear view of 8212zl switch (power supply slots)

zl Power Supplies

The internal power supplies are inserted in the back slots. These slots are labeled PS1 through PS4. A bank of status LEDs on the front of the chassis displays whether power supplies are inserted and providing power, or indicate which of the four power supply slots may have a fault or issue (e.g., supply is inserted into the chassis, but not delivering power or is unplugged). A power supply is hot-swappable, provided at least one other power supply is operational.

The 8212zl requires at least two power supplies for all 12-slots to operate. In the case where the system is reduced to one operational power supply, rather than power down the complete system, only the upper six slots (slots A through F) will receive power. A best practice would be to ensure that the management VLAN access to the switch is connected through one of the top six slots. Reinstatement of a second (or more) power supplies requires a reboot/restart of the 8212zl switch to bring the lower six slots (slots G through L) back online.

There are two different power supplies available for the ProCurve Switch 8212zl Series. These are the same power supplies used for the 5400zl Series, so sparing is minimized for these ProCurve products. Both types of power supplies provide:

- System power (the power needed to run the switch itself), and
- PoE power (the power sent down the Ethernet cable to power the device at the other end).

The only difference between the two supplies is the amount of PoE power available from the supply.

There are two internal power supplies available:

- J8712A 875W (110V/220V)
- J8713A 1500W (220V)

The ProCurve Switch 8212zl and 5400zl Series, as well as the 3500yl Series, can power any device that adheres to the IEEE 802.1af PoE standard. In addition, devices using pre-standard Cisco PoE power also can be powered. The switches will automatically detect what type of and how much power is needed when a compatible device is plugged into the port.

It is recommended that the same power supply model be used for all power supplies installed in a given 8212zl chassis (mixing types in the same chassis, although not detrimental to the operation of the system, creates redundancy cases that may not be what was intended for the power supply and PoE redundancy strategy).



Power Supply Types	System Power	PoE Power	
875W Power Supply J8712A (110/220 VAC)	600W	273W	875W or 1500W Power Supply 
1500W Power Supply J8713A (220 VAC only)	600W	900W	
External Switch zl Power Shelf J8714A	0 (only provides PoE power)	Up to 1800W	Install 1 or 2 of J8712A or J8713A 

Figure 9: Power Supply Choices

The internal power supplies provide system power for all internal components (+12V output within +/- 5 percent tolerance, regulated) and power for PoE ports (-48V output within +/- 5 percent tolerance, regulated), which meets isolation and noise requirements of the IEEE 802.3af specification.

Both internal supplies have over-current, over-temperature and over-voltage protection, and integrated fans. Hot-swapping is allowed, taking into account that disconnecting the power supplies may interrupt PoE operation.

The External Switch zl Power Shelf (J8714A) uses the same power supplies as the 8212zl or 5400zl products. This power shelf connects to an 8212zl/5400zl chassis with rear-mounted cables, to provide additional PoE power when the amount necessary to support additional PoE ports is more than the internal power supplies can provide. The next section provides direction for how to decide which power supply is most appropriate.

zl Power Supply Configuration Planning for PoE

The following criteria can be used for decision-making when choosing which or how many power supplies to configure for an 8212zl or 5400zl chassis switch:

- How much power will be required for each PoE port, full (15.4 watts) or phone (8 watts) power?
- How many network devices will require power (how many PoE ports are needed)?
- Is redundant power required (for the switch power and/or PoE power)?

Note that the power supplies provide a “pool” of power for all interface modules to draw from, that is, PoE power is not limited on a per-interface-module basis.

Power Supply Configuration	Total PoE Power (Watts)	Number of PoE ports at 15.4 W consumption		Number of PoE ports at 8 W consumption	
		No Redundancy	With N+N Redundancy	No Redundancy	With N+N Redundancy
Two 875W (J8712A)	546 W (2 X 273)	35	17	68	34
Four 875W (J8712A)	1092 W (4 x 273)	70	35	136	68
Two 1500W (J8713A)	1800 W (2 x 900)	116	58	225	112

Power Supply Configuration	Total PoE Power (Watts)	Number of PoE ports at 15.4 W consumption		Number of PoE ports at 8 W consumption	
		No Redundancy	With N+N Redundancy	No Redundancy	With N+N Redundancy
Four 1500W (J8713A)	3600 W (4 x 900)	233	116	288 *	225

Table 1: Examples of Maximum Number of PoE Ports for 8212zl/5412zl

* 8212zl/5412zl fully loaded with 12 24-port modules provides up to 288 ports

The table above shows examples of the maximum number of PoE ports that can be supported by the 8212zl/5412zl switch. One column shows the maximum number of PoE ports at full power (Class 0 – 15.4 watts) and the next column shows the maximum number of PoE ports at phone power (8 watts) for each configuration. These categories are further subdivided based on whether the internal power supplies (1 of 2 or 2 of 4) are to provide power redundancy in the event a power supply fails.

Note that four 875W internal power supplies cannot supply sufficient PoE power to an 8212zl/5412zl switch fully populated with 24-port interface modules (288 ports). Four 1500W power supplies can be used to supply full PoE power (15.4 watts) to 233 ports. Providing full PoE power to the maximum number of ports in an 8212zl chassis (288 ports) requires the use of a Switch zl Power Shelf, with at least one (1) 1500W power supply.

For additional information about power supply configurations, see the *ProCurve Switch 8212zl/5400zl/3500yl Series Ordering Guide* that can be found in the ProCurve Networking Reference Library at <http://www.hp.com/rnd/library>.

Switch 8200zl Operational Characteristics

The Command Line Interface and Web GUI are consistent with other Switch zl and yl Series with a few notable differences.

The same K.13.xx code used for the Switch 5400zl/3500yl/6200yl is the same code that runs on the Switch 8200zl. At runtime, the system image detects which platform it is running on, and enables the commands that are unique to the Switch 8212zl.

Management Module Failover

By default, the system is enabled for automatic failover in the case of a management module failure; software crash or hung state; or hardware failure. At initial release of the Switch 8212zl, the time for a “standby” management module to take over should be less than 45 seconds.

This initial failover scenario will be equivalent to a “warm restart,” but the standby module will be in a warm state, partially booted and waiting to take over command of the system when the need arises. The standby MM begins reading the configuration that was synchronized while it was waiting in the wings. At this time, this process requires resetting the interface modules and restarting them with new images. Layer 2 connectivity will be interrupted during this process, and Layer 3 reconvergence will begin approximately 45 seconds after the failover. No state tables are preserved with this initial release. ProCurve will be improving this failover time with major updates to the K-code over the coming months.

Updating SW code – minimizing downtime

With dual management modules, it’s now possible to upgrade code and “warm boot” the standby with the new code.

The two MMs, by default, synchronize their environments: Flash Images (both Primary and Secondary images), configuration files, start-up directives. Since all activity is operated on the “active” MM, simply update the code there and wait for the synchronization to complete.

For example, copy the “newer” code to the primary flash area with the following command:

```
8200zl#copy tftp flash 10.10.10.100 k_12_18.swi primary
```

When the command is complete, the standby card will begin to synchronize and also copy this K.12.18 code to its primary flash area.

The standby MM then would be restarted with the following command:

```
8200z1#boot standby
```

Assuming no directives were changed from their factory default settings, the standby module now will reboot from the Primary Flash area. Since the standby MM is being booted, no interruption to the switch operations will occur (i.e., the “active” MM is still running the “old” version of code).

Once the standby module is fully running in standby mode (and now running the K.12.18 code), the failover can be caused by the following command:

```
8200z1#redundancy switchover management
```

If the desired outcome is for the system to fail back to the earlier code (probably still in the Secondary Flash only the image in the Primary Flash area was overwritten), change the reboot directive to boot from the Secondary Flash area:

```
8200z1#boot set-default flash secondary  
8200z1#boot standby
```

This will now cause the “standby” MM to reload, but from the Secondary Flash area. In the unfortunate occurrence of an unattended crash or other failure, the standby MM will now take control of the system, but this time, running the previous code image still stored in the Secondary Flash area.

ProVision ASIC Architecture

The ProVision ASIC architecture is the latest generation ProCurve ASIC technology, and is used in the ProCurve Switch 8212zl, 5400zl, 3500yl, and 6200yl Series. The ProVision ASIC architecture consists of multiple network chips interconnected by an active crossbar consisting of the fabric chip. A network chip is implemented on each of the various interface modules (also known as “line cards”). The 8212zl’s fabric ASICs are contained on two fabric modules, offering a resilient design. Only one fabric module is required for full operation of all 12 slots; two fabric modules offer full wirespeed capacity for up to 12 slots. On the 5400zl Series, the fabric ASICs are integrated on the switch backplane, which helps reduce the switch’s overall size (7U vs. 9U).

Each network chip represents a node in the system with “High-Speed Links” (HSLs) connecting to the interconnect fabric. Each HSL provides approximately 14.4 Gbps of data bandwidth (28.8 Gbps per module slot). In addition, a management module with a dedicated CPU provides communications control between the network chips and fabric chips. The 8212zl dual-management design offers an “Active/Standby” design so that a failure of an Active Management module will be backed up by a Standby Management module, minimizing downtime of the total switch.

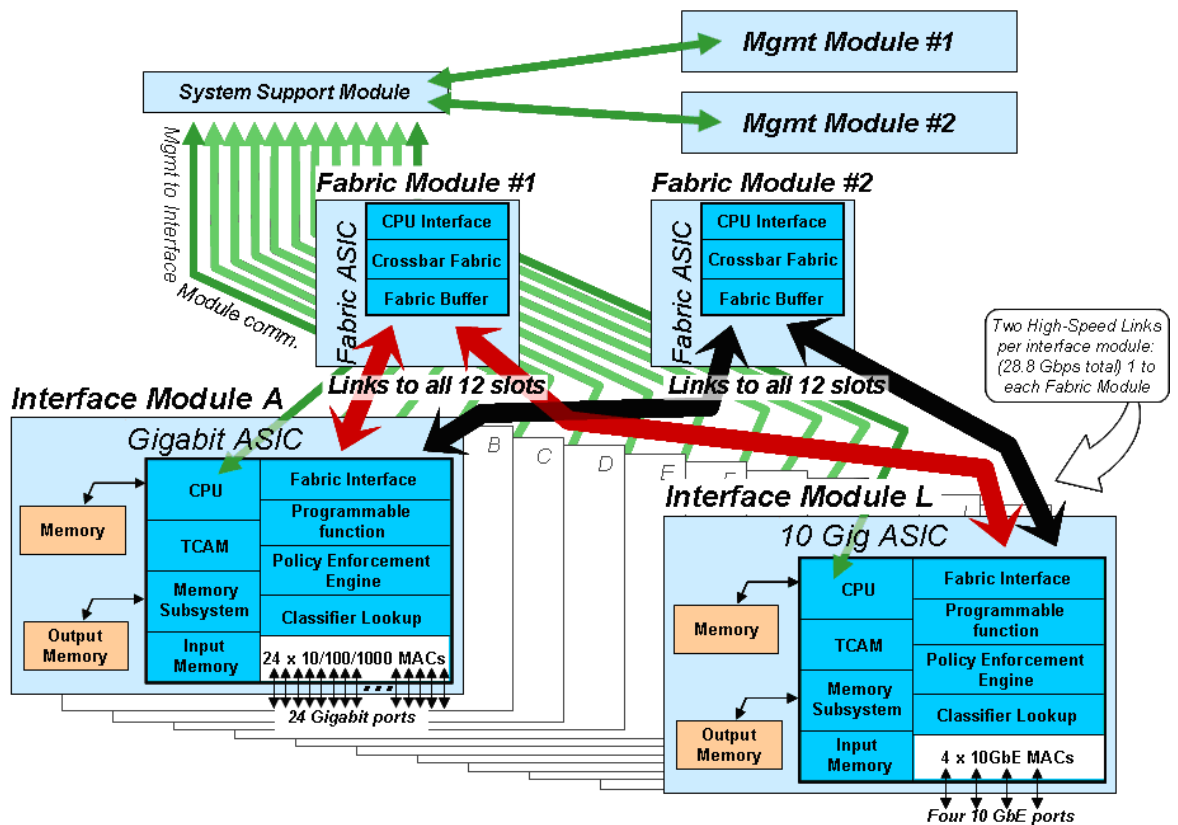


Figure 10: ProVision ASIC Architecture for the Switch 8212zl

The diagram above illustrates an example of the logical interconnection of the ProVision ASICs on the 12-slot Switch 8212zl. All of the key elements are interconnected through the passive backplane. The red and black heavy lines represent the data path on the passive backplane. A physically separate power plane distributes power to all modules to avoid possible issues when power routing is near HSLs between interface modules. Compared to the ProCurve Switch 5412zl, where the Fabric ASIC and components of the SSM are all contained on the active backplane, the 8212zl has a true passive backplane, with removable, customer-repairable modular components.

Inside the ProVision ASIC Architecture

Each interface module contains a full, ASIC-based, Layer 3 routing switch engine as well as Layer 4 filtering and metering. These new ProVision ASICs are ProCurve's fourth-generation switching ASICs. This network switch engine, in the ProVision ASICs, provides all the packet processing: Layer 2 and Layer 3 lookups, filtering and forwarding decisions, VLAN, trunking³ and priority queuing determinations. The ProVision ASIC on each line card contains its own CPU. These features of the ProVision ASIC are common for all products used in the ProCurve Switch 8212zl, 5400zl, 3500yl and 6200yl Series.

Classification and Lookup

When a packet first comes in, the classifier section determines the packet characteristics, its addresses, VLAN affiliation and any priority specification. The packet is stored in input memory; lookups into the table memory are conducted to determine routing information; and a ProVision ASIC-specific packet header is created for the packet with this information. Then, this header is forwarded to the Policy Enforcement Engine.

³ ProCurve's term "trunking" is the aggregation of multiple physical links into one logical link. Other vendors may refer to this as "channels" or "LAG" (Link Aggregated Groups).

Policy Enforcement Engine

The ProVision ASICs on each interface module contain the Policy Enforcement Engine. This engine provides fast packet classification to be applied to ACLs, QoS, rate-limiting and some other features through an onboard TCAM (Ternary Content Addressable Memory). Some of the variables that can be used include source and destination IP addresses (can follow specific users), TCP/UDP port numbers, and ranges (apply ACLs to an application that uses fixed port numbers or ranges). Over 14 different variables can be used to specify the packets to which ACL and QoS rules, rate-limiting counters, and others are to be applied.

Partially implemented in the initial software release, the Policy Enforcement Engine will provide a common front end for the user interface to ACLs, QoS, rate-limiting and some other services. In subsequent software releases for the switches, more features can take advantage of the Policy Enforcement Engine to provide a powerful, flexible method for controlling the network environment. For example, traffic from a specific application can be raised in priority for some users, blocked for other users, and limited in bandwidth for yet other users. After the policy enforcement engine, the header is then forwarded to the programmable section of the network switch engine.

Network Switch Engine Programmability

Each ProVision ASIC switch engine contains multiple programmable units, making them truly Network Processor Units (NPU). One of the functions of the NPU is to analyze the header of each packet as it comes into the switch. The packet's addresses can be read with the switch making forwarding decisions based on this analysis. For example, if a packet's 802.1Q tag needs to be changed to re-map the packet priority, the ProVision ASIC needs to look at each packet to see if any particular one needs to be changed. This packet-by-packet processing has to occur very quickly to maintain overall wirespeed performance – a capability of the ProVision ASICs.

To broaden the flexibility of the ProVision ASICs, a programmable function is included for its packet processing. This NPU function allows the ProCurve designers the opportunity to make future changes or additions in the packet processing features of the ASIC by downloading new software to it. Thus, new features needing high-performance ASIC processing can be accommodated, extending the useful life of the switch without the need to upgrade or replace the hardware. In the first release of the ProCurve Switch 8212zl, 5400zl, 3500yl and 6200yl Series, the NPU function within the ProVision ASICs is totally unused, awaiting future upgrades.

ProCurve's first venture into switching ASIC designs began in 1995, with the introduction of the Switch 2000. The concept of adding the programmable functionality of the NPU within a switching ASIC was designed and implemented in the popular ProCurve Switch 4000M switch family introduced in 1998. ProCurve's 5300xl programmable capability was a third-generation design based on the original ProCurve Switch 4000M implementation. The programmable capability was used to give both the ProCurve Switch 4000M and Switch 5300xl new ASIC-related features well after initial release of those products. Customers with existing units could benefit from the new features through free software downloads. The customer's investment in the ProCurve Switch 4000M and 5300xl is preserved by providing new functionality not otherwise possible without the ASIC NPU programmability.

Being based on the ProCurve Switch 4000M and 5300xl implementations, the NPU capabilities of the ProVision ASICs used in the ProCurve Switch 8212zl, 5400zl, 3500yl and 6200yl Series are a fourth-generation design, following the designs of the Switch 5300xl, Switch 4000M and the original Switch 2000.

Fabric Interface

After the packet header leaves the programmable section, the header is forwarded to the Fabric Interface. The Fabric Interface makes final adjustments to the header based on priority information, multicast grouping, and so on, and then uses this header to modify the actual packet header as necessary.

The Fabric Interface then negotiates with the destination ProVision ASICs for outbound packet buffer space. If congestion on the outbound port is present, WRED (Weighted Random Early Detection) also can be applied at this point as a congestion avoidance mechanism. Finally, the ProVision ASICs Fabric Interface forwards the entire packet through the Fabric-ASIC to an awaiting output buffer on the ProVision ASICs that control the outbound port for the packet. Packet transfer from the ProVision ASICs to the Fabric ASIC is accomplished using the 28.8 Gbps full duplex backplane connection, also managed by the Fabric Interface.

ProVision ASIC CPU

Each ProVision ASIC contains its own CPU for learning of Layer 2 nodes, packet sampling for the XRMON/sFlow function, handling local MIB counters, and running other module-related operations. Overall, the local CPU offloads the master CPU by providing a distributed approach to general housekeeping tasks associated with every packet. MIB variables, which need to be updated with each packet, can be done locally. The Layer 2 forwarding table is kept fresh through the use of this CPU. Other per-port protocols, such as Spanning Tree and LACP, also are run on this CPU. The local CPU, being a full-function microprocessor, allows functionality updates through future software releases.

Fabric ASIC

The Fabric ASICs — one ASIC located on each resilient 8200zl Fabric Module — provides the crossbar fabric for interconnecting the modules together. The use of a crossbar allows wirespeed connections simultaneously from any module to any other module. As mentioned in the ProVision ASICs section, the connection between the Fabric ASIC and each interface module's ProVision ASIC is through a 28.8 Gbps full duplex link.

Each resilient fabric module is connected to all 12 slots. In the event of a single ASIC failure (i.e., one of the two fabric modules fails), the remaining fabric module continues to provide connectivity to all 12 slots, albeit at half-capacity, or 14.4 Gbps full duplex. For networks running at 50 percent utilization or less, there will be no discernable difference in throughput.

Management Subsystem

The management subsystem is responsible for overall switch management. It consists of a CPU, flash memory to hold program code, processor memory for code execution, status LEDs (for management module components) and push buttons, a console interface, and other system support circuitry to interface and control each interface module. In the case of the Switch 8212zl, the management subsystem is on a module that is removable and upgradeable. Each 8200zl chassis requires one management module to function. Redundancy is provided by installing a second management module.

Advanced Capabilities of the Product Family

The ProCurve Switch 8200zl, 5400zl, 3500yl and 6200yl Series family includes a number of advanced capabilities that offer a highly reliable, robust chassis data environment that leads to increased network uptime, keeping overall network costs down.

Versatile Intelligent Port (VIP)

All gigabit copper ports are PoE-enabled and tie into the overall capabilities of the switch, such as the Policy Enforcement Engine. By having the VIP, any port can be connected to any place in the network. If a PC needs gigabit connectivity, it is there. If a wireless access point needs PoE, it is there. And if for some reason these ports need to be swapped on the switch, it can be done simply by moving the cables at the switch or patch panel, without concern for whether the module is PoE-enabled or gigabit-capable.

The VIP provides a high degree of flexibility, as well as future-proofing, and makes troubleshooting easier.

ProVision Hardware Resiliency

Many functions required in a switch have been implemented in the single ASIC on the module. What takes a number of chips in other vendor products is achieved in a single ProVision ASIC. This keeps the part count down, raising overall reliability of the module a significant degree.

Another engineering aspect in the ASIC is hardware error detection, with correction in software for the memory used by the switch. This includes the memory used for forwarding the network traffic, such as the routing and forwarding tables, the Policy Enforcement Engine information, multicast tables, and other data structures. Traffic sent across the backplane uses a protocol to check that there is space available at the destination module, so fabric data is not lost.

Overview of Features and Benefits

The ProCurve Switch 8200zl, 5400zl, 3500yl, and 6200yl Series use the same software image base (letter designated K-code, e.g., K.12.xx). The Switch 8200zl and 6200yl Series support additional Layer 3 services: PIM-SM, PIM-DM, OSPF-ECMP, and VRRP. The primary differences

among these switch families are hardware-related and include such aspects as port density and the number of power supplies and fans.

The following summary of features and benefits applies to the ProCurve Switch 8200zl, 5400zl, 3500yl, and 6200yl Series. Any differences existing among the switches are noted.

Feature	Benefit
Performance	
ProVision ASIC technology	Powered by the ProVision ASICs, the switch families offer state-of-the-art, high-capacity, switch fabric performance: 692 Gbps for the 8212zl/5412zl; 346 Gbps for the 5406zl; 173 Gbps for the 3500yl-48G-PWR; and 115 Gbps for the 3500yl-24G-PWR and 6200yl-24G-mGBIC.
Selectable queue configurations	Increase performance by selecting the number of hardware forwarding queues and associated memory buffer that best meet the requirements of network applications.
Security-Enabled:	
Virus throttle	Connection-rate filtering thwarts viruses from spreading by blocking routing from certain hosts exhibiting abnormal traffic behavior.
ICMP rate-limiting	Throttles denial-of-service (DoS) attacks or other malicious behavior that uses high-volume ICMP traffic.
Switch CPU protection	Provides automatic protection against malicious network traffic trying to shut down the switch.
Detection of malicious attacks	Monitors 10 types of network traffic and sends warnings when anomalies potentially caused by malicious attacks are detected.
DHCP protection	Blocks DHCP response packets from being forwarded if received from an unauthorized port.
BPDU port protection	Blocks Bridge Protocol Data Unit (BPDU) on ports that should not be receiving BPDUs, preventing forged BPDU attack.
Dynamic ARP protection	Blocks ARP broadcast from unauthorized hosts, preventing eavesdropping or data theft of network data.
Dynamic IP lockdown	Works with DHCP protection to block traffic from unauthorized host, preventing IP source address spoofing.
Filtering capabilities	Include fast, flexible Access Control Lists (ACLs), up to 3,000 per module, source port, multicast MAC address, and other protocol-based filtering capabilities.
Identity Driven Manager (IDM)	ProCurve IDM to dynamically apply security, access and performance settings to infrastructure devices based on approved user, location and time.
Port security	Port security, MAC lockdown and MAC lockout protection for restricting access to the network through a switch port.
User Authentication	Client-based access control using 802.1X, Web-based, MAC-based authentication, RADIUS, and TACACS+. At initial release, the combinations of authentication methods allowed simultaneously on a port are 802.1X/Web and 802.1X/MAC.
Management security	SSH v1 and v2, SSLv3 and TLSv1, TACACS+, and Secure FTP for management access. Management VLAN—Limit CLI/GUI/telnet/SNMP access to the switch to a particular VLAN. Manager authorized list limits CLI/Web/Telnet/SNMP management access to the switch to particular authorized hosts.
Redundancy and High Availability	
Redundancy Protocols	Protocols providing high availability include 802.1Q Multiple Spanning Tree Protocol, Switch Meshing, and Virtual Router Redundancy Protocol (Premium License).

Feature	Benefit
Power Supplies	Optional redundant power supplies available for Switch 8200zl/5400zl Series.
Management Module	Removable/replaceable management module for Switch 5400zl Series.
Dual Mgmt Module	Active/Standby management module for Switch 8200zl Series.
Operational Intelligence	
Policy Enforcement Engine	Policy Enforcement Engine is user-configured to select packets that are then forwarded or dropped (based on ACLs, QoS, and rate-limiting). The engine is fast, and can look for multiple variables, such as an IP address and port number, in a single pass through a packet. Provides a common user experience regardless of switch to which the user is connected.
Operational Flexibility	
Versatile Intelligent Port (VIP)	VIP supports flexible client port connections as every copper port is: 10/100/1000 capable, PoE-enabled, has all Layer 3 services applied, and has eight hardware queues. All copper ports are created equal, so any 10/100/1000 copper port can be connected to any 10/100/1000 copper client. The Switch 8212zl, 5400zl and 3500yl Series support copper ports, while the Switch 6200yl Series only supports mGBIC fiber or mGBIC RJ45 (mGBIC connections do not allow PoE).
PoE Power	Power supply choice for optimum PoE power on 8212zl/5400zl.
Port Density	Choice of 24 to 288 ports in single, consistent product family.
L3 Services at L2	ACLs, QoS, and other features can use Layer 3 information, such as IP addresses, without requiring the switch to have a routing feature set.
Intelligent Edge/ Premium License	Flexible approach to licensing. The 5400zl and 3500yl switches can run the Intelligent Edge feature group initially, and then be upgraded later to run the Premium License feature group, if advanced Layer 3 features are needed. The Layer 3 features include VRRP, PIM-SM, PIM-DM, and OSPF-ECMP. A Premium License can be transferred to another switch, as long as the license stays in the same hardware family. The Switch 8200zl and 6200yl Series come standard with these features.
Convergence Ready	
Bandwidth Shaping / Control	Guaranteed minimums can be applied to traffic that must always get through the switch. Enforced maximums can be used to limit problem clients' bandwidth to no more than a set amount.
QoS	QoS based on TCP/UDP ports and other variables allows precise control of packet priority. Eight priority levels mapped to eight hardware queues can be used to set outgoing IP priority.
L2/L3 Jumbo Frames	Layer 2/Layer 3 jumbo frames provide scalability in throughput.
Supports Standard Protocols	
IP Routing	Supports RIPv1, RIPv2, OSPF, OSPF-ECMP and static routes. OSPF requires the Premium License for the 5400zl and 3500yl. Standard on the Switch 8200zl and 6200yl Series.
IPv6	Supports deployment on IPv6 network with IPv6 host, IPv4/IPv6 dual stack, and MLD snooping
VLANs	Supports 802.1Q, 2048 concurrent VLANs (4094 in later release), Group VLAN Registration Protocol (GVRP).
Q-in-Q	Increases the scalability of Ethernet network by providing a hierarchical structure; connects multiple locations on high-speed campus or metro Ethernet network.
IP Multicast	Supports PIM-SM, PIM-DM, and IGMP data-driven snooping.

Feature	Benefit
Low Cost of Ownership	
Future-proofing	<p>Versatile Intelligent Port (VIP) implies all gigabit copper ports on the 8212zl, 5400zl and 3500yl Series are PoE-enabled.</p> <p>The 8212zl and 5400zl Series chassis management module is upgradeable in order to meet future needs for processing power.</p> <p>Programmable ASICs allow some future requirements to be implemented without replacing the switch.</p> <p>Intelligent Edge features can be upgraded to Premium License features on the 5400zl and 3500yl Series, allowing an edge switch to be redeployed as a distribution switch.</p>
Lifetime warranty	<p>Next-business-day advance replacement for as long as you own the product (available in most countries).</p> <p>Applies to all components of the 8212zl, 5400zl, 3500yl and 6200yl products including the power supplies.</p>
Lifetime SW Updates	No additional charge for software feature enhancements to the Intelligent Edge and Premium licenses.

Feature Set Summary

Additional details about the ProCurve Switch 8212zl, 5400zl, 3500yl, and 6200yl Series features include the following:

Security Features

- Virus Throttle: Connection-rate filtering thwarts virus spreading by blocking routing from certain hosts exhibiting abnormal traffic behavior.
- ICMP throttling: Defeats ICMP DoS attacks by enabling any switch port to automatically throttle ICMP traffic.
- Filtering Capabilities: Includes fast, flexible Access Control Lists (ACLs), up to 3,000 per module (in later release, more precise detailed control via the fast Policy Enforcement Engine), source port, multicast MAC address, and other protocol-based filtering capabilities.
- Detection of malicious attacks: Monitors 10 types of network traffic and sends a warning if an anomaly occurs signaling the detection of a potential malicious attack.
- USB secure autorun: Uses USB flash drive to deploy, troubleshoot, or update switches; work with secure credential to prevent tampering.
- STP root guard: Protects STP root bridge from malicious attack or configuration mistakes.
- DHCP protection: Blocks DHCP packets from unauthorized DHCP servers, preventing DoS attack.
- BPDU port protection: BPDU on ports that do not require BPDU, preventing forged BPDU attack.
- Dynamic ARP protection: Blocks ARP broadcast from unauthorized hosts, preventing eavesdropping or data thief of network data.
- Dynamic IP lockdown: Works with DHCP protection to block traffic from unauthorized host, preventing IP source address spoofing.
- Multiple user authentication methods:
 - IEEE 802.1X: Industry-standard way of user authentication using an 802.1X supplicant on the client in conjunction with a RADIUS server.
 - Web-based authentication: Similar to 802.1X, provides a browser-based environment to authenticate clients that do not support the 802.1X supplicant.
 - MAC-based authentication: Client is authenticated with the RADIUS server based on the MAC address of the client; useful for clients who have minimal or no user interface.
- Authentication flexibility:
 - Multiple 802.1X users per port: Provides authentication of multiple 802.1X users per port; prevents user "piggybacking" on another user's 802.1X authentication.

- Concurrent 802.1X and Web or MAC authentication schemes per port: At initial release, the combinations of authentication methods allowed simultaneously on a port are 802.1X/Web and 802.1X/MAC.
- Access control lists (ACLs): Provide filtering based on the IP field, source/destination IP address/subnet, and source/destination TCP/UDP port number on a per-VLAN or per-port basis.
- Identity-driven ACL: Enables implementation of a highly granular and flexible access security policy specific to each authenticated network user.
- Port security: Prevents unauthorized access using MAC address lockdown.
- MAC address lockout: Prevents configured particular MAC addresses from connecting to the network.
- Source-port filtering: Allows only specified ports to communicate with each other.
- Security banner: Displays customized security policy when users log into the switch.
- Management access:
 - All access methods — CLI, GUI or MIB — are encrypted securely through SSHv2, SSL, and/or SNMPv3
 - RADIUS and TACACS+: Can require either RADIUS or TACACS+ authentication for secure switch CLI logon.
 - Secure FTP: Allows secure file transfer to/from the switch and protects against unwanted file downloads or unauthorized copying of switch configuration file.

QoS Functions

Layer 4-prioritization: Enables prioritization based on TCP/UDP ports.

Traffic prioritization: Allows real-time traffic classification into eight priority levels mapped to eight queues.

Bandwidth shaping using:

- Rate-limiting: Per-port, ingress-based, enforced bandwidth maximums
- Guaranteed minimums: Per-port, per-queue, egress-based, guaranteed bandwidth minimums

Class of Service (CoS): Sets 802.1p priority tag based on IP address, IP Type of Service (ToS), Layer 3 protocol, TCP/UDP port number, source port, and DiffServ.

Convergence

- IP multicast routing: Includes PIM Sparse and Dense modes to route IP multicast traffic.
- IP multicast data-driven IGMP: Automatically prevents flooding of IP multicast traffic.
- LLDP-MED (Media Endpoint Discovery): A standard extension of LLDP that stores values for parameters, such as QoS and VLAN, to automatically configure network devices such as IP phones.
- RADIUS VLAN for voice: Uses standard RADIUS attribute and LLDP-MED to automatically configure VLAN for IP phones.
- PoE allocations: Supports multiple methods (automatic, 802.3af class, LLDP-MED or user-specified) to allocate PoE power for optimal energy saving.
- iSCSI support: Enables the deployment of Ethernet storage area network solutions using the iSCSI standard.

Layer 2 switching

- ProCurve switch meshing: Dynamically loads balances across multiple active redundant links to increase available aggregate bandwidth (routing must be disabled).
- VLAN support and tagging: Support for complete 802.1Q standard and 2,048 VLANs simultaneously.
- 802.1v protocol VLANs: Isolate select non-IPv4 protocols automatically into their own VLANs.
- GVRP: Group VLAN Registration Protocol allows automatic learning and dynamic assignment of VLANs.

- Q-in-Q: Increases the scalability of Ethernet network by providing a hierarchical structure; connects multiple LANs on high-speed campus or metro Ethernet network.

Bridging Protocols

- MSTP: Provides high link availability in multiple VLAN environments by allowing multiple spanning trees; encompasses 802.1D STP and 802.1w RSTP.

Routing Protocols

- RIP (v1, v2, and v1 compatible v2)
- OSPF with host-based ECMP (Equal Cost Multipath) and NSSA
- Static IP routes: To manually add routes directly to the routing table
- 10,000 network address routes, 65,536 Layer 3 host address routes
- UDP helper function: UDP broadcasts can be directed across router interfaces to specific IP unicast or subnet broadcast addresses, and prevent server spoofing for UDP services such as DHCP
- Loopback interface address: Define an interface in RIP and OSPF that can always be reachable, improving diagnostic capability
- IPv4 routing

IPv6

- IPv6 host: Switches can be deployed and managed at the edge of IPv6 LAN
- IPv4/IPv6 Dual Stack: Provides transition mechanism for IPv4 and IPv6
- MLD Snooping: Prevents multicast traffic from flooding the network

Multicast Protocols

- IGMP data-driven
- PIM-SM, PIM-DM (Permanently enabled on the 8212zl and 6200yl; requires Premium License for the 5400zl/3500yl.)

High Availability and Redundancy

- VRRP: Virtual Router Redundancy Protocol (permanently enabled on the 8212zl and 6200yl; requires Premium License for the 5400zl/3500yl) allows groups of routers to dynamically back each other up to create highly available routed environments.
- 802.3ad LACP: Link Aggregation Control Protocol and ProCurve trunking support up to 36 trunks, each with up to eight links (ports) per trunk.
- Port trunks, or link aggregation groups, can operate across multiple modules to increase redundancy.
- Supports various redundant power supply configurations: The 8212zl and 5412zl provides slots for four internal supplies. This can provide N+N redundancy similar to the 5406zl. The internal supplies can be used for N+N backup; two backing up two for both PoE and system power needs.
- Hot-swappable/hot-insertable interface modules, so as not to interrupt network operation.
- Management module is removable/upgradeable: The switch does not have to be powered off to remove a management module from an 8212zl/5400zl Series. However, when the system is operating with only one management module and is removed, all ports will lose communication and the system will be powered down.
- Other Layer 2 and Layer 3 redundant protocols include MSTP, ProCurve switch meshing, and OSPF-ECMP (requires Premium License for the 5400zl/3500yl).

Management

- Remote intelligent mirroring: Mirrors selected ingress/egress traffic based on ACL, port, MAC address, or VLAN to a local or remote 8200/6200/5400/3500 switch anywhere on the network.
- RMON, XRMON, and sFlow: Provide advanced monitoring and reporting capabilities for statistics, history, alarms and events.

- Command authorization: Leverages RADIUS to link a custom list of CLI commands to individual network administrator's login; also provides audit trail.
- LLDP: Link Layer Discovery Protocol (LLDP) allows automated device discovery for easy mapping by network management applications.
- Friendly port names: Allow assignment of descriptive names to ports.
- Dual flash images: Provide independent primary and secondary OS files for backup while upgrading.
- Multiple configuration files: Allows multiple configuration files to be stored to flash image.
- Unidirectional Link Detection (UDLD): Monitors a link between two switches and blocks the ports on both ends of the link if only one of a fiber pair goes down at any point between the two devices.
- USB support: Allows users to copy files to/from the switch to a USB device.
- Configuration, monitoring and management tasks can be performed using the CLI (through direct console or Telnet) connection, Web browser, Menu interface, and ProCurve Manager.

Standards and Protocols

RFC 768 UDP	RFC 783 TFTP
RFC 792 ICMP	RFC 793 TCP
RFC 826 ARP	RFC 854 Telnet
RFC 951 BootP	RFC 1058 RIP
RFC 1213 MIB II	RFC 1492 TACACS+
RFC 1493 Bridge MIB	RFC 1724 RIPv2 MIB
RFC 1542 BootP	RFC 1850 OSPF MIB
RFC 2021 RMON v2 Probe Configuration	RFC 2030 Simple Network Time Protocol
RFC 2096 IP Forwarding Table MIB	RFC 2138 RADIUS
RFC 2328 OSPFv2 ¹	RFC 2453 RIPv2
RFC 2474 DiffServ Precedence	RFC 2597 DiffServ Expedited Forwarding
RFC 2598 DiffServ Assured Forwarding	RFC 2613 SMON
RFC 2618 RADIUS Client MIB	RFC 2620 RADIUS Accounting MIB
RFC 2665 Ethernet MIB	RFC 2668 802.3 MAU MIB
RFC 2674 802.1p and IEEE 802.1Q Bridge MIB	RFC 2737 Entity MIB
RFC 2787 VRRP MIB	RFC 2819 RMON MIB ²
RFC 2863 Evolution of Interface	RFC 2866 RADIUS Accounting
RFC 3046 DHCP Relay Agent Information Option	RFC 3376 IGMPv1/v2/v3
RFC 3768 VRRP ¹	IEEE 802.1AB Link Layer Discovery Protocol
IEEE 802.1D Spanning Tree	IEEE 802.1p Priority
IEEE 802.1Q GVRP	IEEE 802.1Q VLAN tagging
IEEE 802.1s Multiple Spanning Tree	IEEE 802.1w Rapid Convergence Spanning Tree
IEEE 802.1X Network Login	IEEE 802.3ad Link Aggregation Control Protocol
IEEE 802.3af Power over Ethernet (PoE)	IEEE 802.3x Flow Control
ANSI/TIA-1057 LLDP Media Endpoint Discovery	DHCP Relay
HTML and Telnet management	IGMPv3
PIM Dense Mode ¹	PIM Sparse Mode ¹
Secure Sockets Layer	sFlow
SNMPv1/v2c/v3	SSHv1/SSHv2 Secure Shell
XRMON	RFC 2925 Ping MIB

Notes:

1. Requires Premium License for 5400zl and 3500yl. Standard on 8212zl and 6200yl
2. RFC 2819 RMON MIB four groups: 1 (statistics), 2 (history), 3 (alarm) and 9 (events)

Performance and Capacity

Capacity and Performance Features Comparison

The following table provides a comparison of several categories of capacity and performance levels for the ProCurve Switch 8212zl, 5400zl, 3500yl and 6200yl Series. These include the routing and switching capacity (Gbps), switch fabric speed (Gbps), maximum number of 1 Gbps that can operate concurrently at wirespeed, maximum number of 10 Gbps ports that can operate concurrently at wirespeed, and the size of the routing table (entries).

Capacity / Performance Feature	ProCurve Switch 8212zl/ 5412zl	ProCurve Switch 5406zl	ProCurve Switch 3500yl 24G-PWR	ProCurve Switch 3500yl 48G-PWR	ProCurve Switch 6200yl-24G-mGBIC
Routing / Switching capacity (Gbps)	576	288	101	148	101
Switch Fabric Speed (Gbps)	692	346	115	173	115
Max 1 Gigabit ports at wirespeed	288	144	24	48	24
Max 10 Gigabit ports at wirespeed	24	12	2	2	2
Routing table size (entries)	10,000	10,000	10,000	10,000	10,000

Table 2: Capacity and Performance Features Comparison of 5400zl, 3500yl and 6200yl Series

Per-Port Buffer Sizes

The ProCurve Switch 8212zl, 5400zl, 3500yl and 6200yl Series have eight hardware queues to which prioritized traffic can be assigned. Software can designate a certain amount of output memory buffer for each queue to minimize impact from events, such as network traffic bursts that exceed line rate. Users can choose to configure the number of queues — two, four or eight — so that they can optimize the amount of memory buffer per queue for their applications. For example, if a user has only two priorities on the network, that user should configure for two queues so all buffer memory is assigned to only two queues. The table below identifies the buffer sizes allocated on a per-port basis to the normal and other priority queues in each configuration.

Port Speed	Size of Normal Priority Queue (KB)	Size of Other Priority Queues (KB)
1 Gbps	194 KB (30%)	70 KB (10%)
10 Gbps	225 KB (30%)	77 KB (10%)

Table 3a: Per-port Priority Queue Buffer Sizes for 8 queues

Port Speed	Size of Normal Priority Queue (KB)	Size of Other Priority Queues (KB)
1 Gbps	480 KB (70%)	68 KB (10%)
10 Gbps	538 KB (70%)	76 KB (10%)

Table 3b: Per-port Priority Queue Buffer Sizes for 4 queues

Port Speed	Size of Normal Priority Queue (KB)	Size of Other Priority Queues (KB)
1 Gbps	618 KB (90%)	68 KB (10%)
10 Gbps	691 KB (90%)	76 KB (10%)

Table 3c: Per-port Priority Queue Buffer Sizes for 2 queues

Routing and Forwarding Tables

The ProCurve Switch 8200zl, 5400zl, 3500yl, and 6200yl Series have several routing and forwarding table features that enhance routing and switching performance.

Every interface module has its own “best match prefix” routing table that contains IP routes and is used for determining how to route the vast majority of incoming packets. Using the best match prefix routing table is extremely fast and enables wirespeed routing to be achieved.

If the best match prefix routing table does not contain an entry that can be used to determine the route of a received packet, then the main routing table is used. The main routing table can contain up to 10,000 routing table entries.

For Layer 2 forwarding, each switch has a forwarding or MAC table that can contain up to 65,536 entries (64K).

Optimizing the 10-GbE Port Configuration

The 10-GbE modules (J8707A, J8708A, and J8694A) used in the ProCurve Switch 8212zl are designed to deliver full 10 Gbps wirespeed to each port, where either one or two ports are in a linked state with another device. When three or four 10-GbE ports are in a linked state, the module supports an aggregate bandwidth of 28.8 Gbps across the linked ports.

As illustrated in the figure below, internally, there are two 14.4-Gbps channels between each four-port 10-GbE module and the switch fabric.

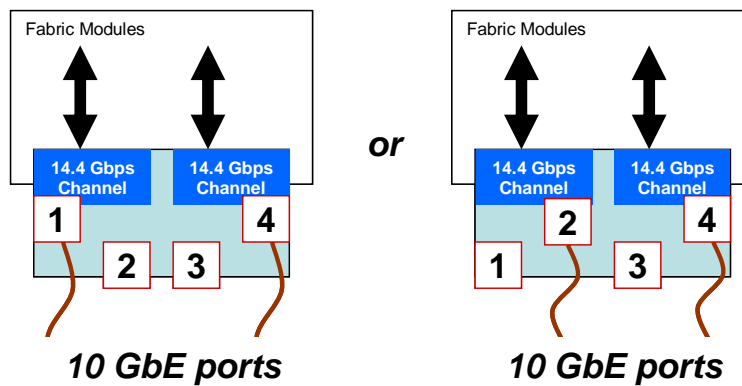


Figure 11: 10-GbE module architecture showing two (2) ports grouped to 14.4 Gbps Channel

- When any two 10-GbE ports on a given module are in a linked state, each port automatically operates on its own channel, which guarantees 10 Gbps of bandwidth for each port. The two ports are *dynamically* mapped to an available high-speed channel and support full 10 Gbps wirespeed operation.




Figure 12: 10-GbE module architecture showing four (4) ports grouped to 14.4 Gbps Channel

- When more than two ports of a 10-GbE module are in a linked state, ports 1 and 4 are statically mapped to share one 14.4 Gbps channel, while ports 2 and 3 are statically mapped to share the other 14.4 Gbps channel.
- If only one port in a given channel is in a linked state, then that port operates at wirespeed and the other port uses no bandwidth.
- If both ports in a given channel are in a linked state, then the 14.4 Gbps of bandwidth is balanced fairly between the two ports.

For example, in an application where three ports on the module are needed and it is important to ensure that port 1 always has a full 10 Gbps available, the connection choices shown in the next figure would be used.

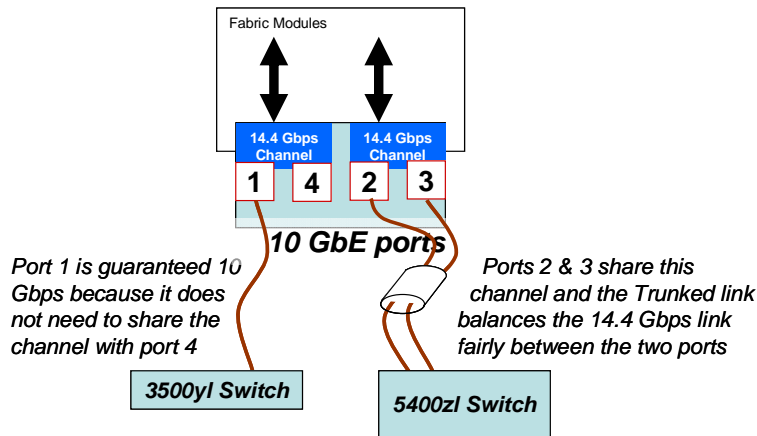


Figure 13: Approach for guaranteeing 10 Gbps on a specific port

Connection choices are also important where equally balanced bandwidth is needed, such as in a Virtual Router Redundancy Protocol (VRRP) application. This scenario is illustrated in the next figure.

- Each trunk is guaranteed 20 Gbps bandwidth because, in the VRRP scenario, one router is the Owner and the other is the Backup
- As a result, only one port in each 14.4 Gbps channel is in use since the trunks are using different channels




Figure 14: Guaranteeing Equally Balanced Bandwidth in a VRRP Environment

Throughput and Latency Performance Data

The following three tables describe the performance capabilities of the 10 Gbps, 1 Gbps and 100 Mbps interface module ports of the ProCurve Switch 8200zl, 5400zl, 3500yl and 6200yl Series.

Over a range of packet sizes from 64 bytes to 1518 bytes, the performance levels achieved are described in terms of the following metrics:

- Throughput level percentage achieved when performing Layer 2 switching and Layer 3 routing; 100 percent in all cases.
- Throughput in the form of the number of packets per second when performing Layer 2 switching and Layer 3 routing.
- Average per-packet latency in microseconds when performing Layer 2 switching and Layer 3 routing. These values shown are listed for both First In First Out (FIFO) and Last In First Out (LIFO) processing. The LIFO values represent the packet transmit time, whereas the FIFO values include the switch decision time and the packet transmit time.
- Percentage of packets dropped in a full mesh configuration; 0 percent in all cases.

The following measurements were performed by ProCurve using test equipment manufactured by Ixia Communications (www.ixiacom.com). In these tests, the maximum number of supported ports on the ProCurve Switch 5412zl was used. Proportional performance results at 100 percent of throughput are achieved on the ProCurve Switch 8200zl, 5406zl, 3500yl-24G-PWR, 3500yl-48G-PWR, and 6200yl-24G-mGBIC for the maximum number of ports supported by each of those models. The 100 Mbps table is not applicable to the ProCurve Switch 6200yl-24G-mGBIC, since only 1 Gigabit and 10 Gigabit ports are supported.

10 Gbps	Throughput%	Packets per Second	Full Mesh % Drops	Latency (FIFO)		Latency (LIFO)	
Packet Size	L2/L3	L2/L3		L2 (μsecs)	L3 (μsecs)	L2 (μsecs)	L3 (μsecs)
64 ¹	100	14880952	0	1.95	2.05	1.9	2.0
128	100	8445946	0	2.01	2.11	1.9	2.0
256	100	4528986	0	2.31	2.31	2.1	2.1
512	100	2349624	0	2.62	2.72	2.2	2.3
1024	100	1197318	0	3.73	3.73	2.9	2.9
1280	100	961538	0	4.13	4.23	3.1	3.2
1518	100	811688	0	4.52	4.52	3.3	3.3

1 Gbps	Throughput%	Packets per Second	Full Mesh % Drops	Latency (FIFO)		Latency (LIFO)	
Packet Size	L2/L3	L2/L3		L2 (μsecs)	L3 (μsecs)	L2/L3	L2/L3
64	100	1488095	0	3.28	3.19	2.7	2.7
128	100	844595	0	3.79	3.79	2.7	2.7
256	100	452899	0	4.91	4.91	2.8	2.8
512	100	234962	0	7.16	7.16	3.0	3.0
1024	100	119732	0	11.76	11.76	3.5	3.5
1280	100	96154	0	14.00	14.2	3.7	3.9
1518	100	81274	0	16.20	16.20	4.0	4.0

100 Mbps	Throughput%	Packets per Second	Full Mesh % Drops	Latency (FIFO)		Latency (LIFO)	
Packet Size	L2/L3	L2/L3		L2 (μsecs)	L3 (μsecs)	L2/L3	L2/L3
64	100	148810	0	3.28	3.19	2.7	2.7
128	100	84459	0	3.79	3.79	2.7	2.7
256	100	45290	0	4.91	4.91	2.8	2.8
512	100	23496	0	7.16	7.16	3.0	3.0
1024	100	11973	0	11.76	11.76	3.5	3.5
1280	100	9615	0	14.00	14.2	3.7	3.9
1518	100	8127	0	16.20	16.20	4.0	4.0

Table 4: Throughput Performance

Note:

1. See the explanation about 10 Gigabit Performance Traffic Patterns in the next section.

10 Gigabit Performance Traffic Patterns

In the prior table, the performance levels for **10 Gigabit** ports assume the underlying traffic patterns reflect either one of the following minimum conditions to achieve wirespeed throughput.

- A single source traffic stream with an average packet size of 88 bytes or larger.

- Two or more source traffic streams of any packet size down to the minimum value of 64 bytes.

In the unlikely case where the average packet size is consistently smaller, the throughput will be less than wirespeed. For example, consider a worst-case scenario, where the average packet size is 64 bytes. This would result in a throughput of approximately 70 percent to 80 percent of the rated wirespeed capacity. ProCurve considers such minimum-sized, packet traffic scenarios being realized over an extended period of time to be extremely atypical and unlikely to be experienced by customers in the field.

Note that the limits described above do not apply to Gigabit ports.

Throughput Test

A fully meshed performance test sends packets from each port to every other port during the test. This type of test exercises both the modules and the backplane. These tests show the ProCurve Switch 8200zl, 5400zl, 3500yl and 6200yl Series to be wirespeed on all ports simultaneously.

Latency Measurements

Latency is commonly measured as the amount of time it takes for a byte inside a packet to enter and then leave the switch. Latency statistics typically are documented as including both the processing time of the switch as it makes its forwarding decision, and the time for the packet itself to enter and leave the switch. In the prior tables, this definition of latency corresponds to the FIFO latency statistics. The LIFO latency statistics that also are listed in the tables represent only the packet transmission time.

Almost all switches currently on the market are store and forward, so the entire packet is received into the switch before the switch begins to transmit the packet out the egress port. Including the packet receive time in the FIFO latency statistics is appropriate, since this extra time is a contributing component of the overall transit time of the packet as it moves through the network.

The latency figures for the ProCurve Switch 8200zl, 5400zl, 3500yl/6200yl Series are consistently low. Latencies this low will not be a factor in general network operation, even with streaming video or VoIP applications. The LIFO latency values are fairly consistent across all packet sizes due to the fact that ingress and egress packet processors operate on the header of the frame (not the whole frame), while the full frame is buffered in and out of packet buffer memory. Memory transfers are scheduled to fit a full 1518-byte frame, so frames are transferred in and out of memory in approximately the same amount of time, regardless of packet size. While the frame headers are being looked up, and actions required for the frame on egress are being coordinated among interface modules, the frame is transferred through the switching fabric module.

Comparison of 8200zl/5400zl with 5300xl and 8100fl

The table below provides a general comparison of the ProCurve Switch 8200zl, 5400zl and 3500yl Series with the ProCurve Switch 5300xl and 8100fl Series products. The features compared include the number of 1 Gigabit and 10 Gigabit ports supported, and routing table sizes.

The 8100fl Series offers high-performance, high-availability, cost-effective connectivity for Intelligent Edge devices while delivering a flexible, scalable, high port density 1 Gigabit and 10 Gigabit Ethernet core networking solution. The 8100fl Series offers two chassis configurations: the ProCurve Switch 8108fl and ProCurve Switch 8116fl. The ProCurve Switch 8108fl is an eight-slot chassis-based routing switch with support for up to 16 10GbE ports (2:1 oversubscribed) or 80 100/Gigabit Ethernet ports (wirespeed). The ProCurve Switch 8116fl is a 16-slot chassis-based routing switch with support for up to 32 10GbE (2:1 oversubscribed) or 160 100/Gigabit Ethernet ports (wirespeed).

The ProCurve Switch 5300xl Series offers scalable Layer 2, 3 and 4 switching in compact four- and eight-slot modular form factors. The ProCurve Switch 5304xl is a four-slot chassis-based routing switch with support for up to 64 10/100/1000 Ethernet ports. The ProCurve Switch 5308xl is an eight-slot chassis-based routing switch with support for up to 128 10/100/1000 Ethernet ports.

Feature	ProCurve Switch 8200zl	ProCurve Switch 5400zl	ProCurve Switch 3500yl	ProCurve Switch 5300xl	ProCurve Switch 8100fl
Routing / Switching capacity (Gbps)	576	576	148	76	320
Chassis Height	9U(8212zl)	4U (5406zl) 7U (5412zl)	1U	3U (5304xl) 5U (5308xl)	13U (8108fl) 20U (8116fl)
Maximum Number of 1 GbE Ports	288	288	48	128	160
Maximum Number of 10 GbE Ports	48	48	4	0	16
Power over Ethernet (PoE)	Yes	Yes	Yes	Yes	No
Jumbo Frames	Yes	Yes	Yes	No	Yes
Number of Priority Queues	8	8	8	4	5
Number of ACLs (entries)	3,000	3,000	3,000	1000	1000
Number of VLANs	2,048	2,048	2,048	256	2,048
Forwarding Table Size (entries)	64K	64K	64K	16K	100,000
Routing Table Size (entries)	10,000	10,000	10,000	10,000	100,000
RIP v1/v2	Yes	Yes	Yes	Yes	Yes
OSPF / OSPF-ECMP	Yes	Yes ¹	Yes ¹	Yes	Yes
PIM-SM	Yes	Yes ¹	Yes ¹	No	No
PIM-DM	Yes	Yes ¹	Yes ¹	Yes	No
Spanning Tree	MSTP	MSTP	MSTP	RSTP, MSTP	MSTP
LACP	Yes	Yes	Yes	Yes	Yes
Virtual Router Redundancy	VRRP	VRRP ¹	VRRP ¹	XRRP	VRRP
Redundant Power	Yes	Yes	Yes (external unit)	Yes	Yes
Redundant Switch Fabric	Yes	No	No	No	Yes
Redundant Management module	Yes	No	No	No	Yes

Table 9: Comparison of 8200zl, 5400zl, 3500yl, 5300xl and 8100fl Series

Note:

1. Requires Premium License to enable this feature.

The architecture used for ACL storage and processing is different among the ProCurve switches listed in the table above. The ProCurve Switch 8200zl, 5400zl and 3500yl/6200yl Series have the additional performance advantage gained through the use of the TCAM (see the appendix for more details). The ProCurve Switch 8200zl, 5400zl and 3500yl/6200yl Series also have an additional performance advantage in the processing of the routing table through the use of a "best match prefix" table (described previously in this section). The ProCurve Switch 8200zl, 5400zl and 3500yl/6200yl Series support the use of PIM-SM and PIM-DM, but not concurrently.

Services and Support

Lifetime Software Updates (Best in the Industry)

As with other ProCurve products, the ProCurve Switch 8200zl, 5400zl, 3500yl, and 6200yl Series come with lifetime software updates. For as long as software updates are available for these switches, customers can download updates from the ProCurve web site, free of charge.

Lifetime Warranty (Best in the Industry)

Warranties, and the ease of obtaining warranty service for the customer, is a product benefit that is easily overlooked in a technical evaluation, but ranks high as a concern of customers as they get ready to actually make a purchase decision. The ProCurve Switch 8200zl, 5400zl, 3500yl, and 6200yl Series have a limited lifetime warranty, for as long as you own the product. If any part of the switch fails due to a defect in material or workmanship, including the power supply or fans, it will be replaced. In most parts of the world, the replacement unit is sent with next-business-day delivery in advance of the failing unit being returned to ProCurve.

Advance replacement gets the unit to the customers as fast as possible and minimizes down time by allowing the impaired unit to continue to be used if possible until the replacement unit arrives. This also allows for easy scheduling for when the actual unit swap occurs on the network. An optional upgrade to on-site replacement is also available in most parts of the world. For more details about warranty coverage, refer to the warranty statement that is shipped with the product.

The ProCurve Switch 8200zl, 5400zl, 3500yl, and 6200yl Series warranties are industry-leading.

Free Telephone Support

ProCurve provides free pre-sales and post-sales telephone support during normal business hours to end users and ProCurve resellers through the ProCurve Customer Care Centers located worldwide.

Optional Support Services

In addition to free support services, such as the warranty and telephone support, ProCurve offers an extensive range of fee-based support services to meet more specialized needs. The following optional services are available for the ProCurve Switch 8200zl, 5400zl, 3500yl, and 6200yl Series:

- On-site next business day
- On-site in four hours, same business day
- On-site in four hours 24x7
- Six-hour call-to-repair
- 24x7 telephone support

Hewlett-Packard also can provide more broad-based services such as site surveys, installation services, and actual management of the network, depending on customer needs.

More information can be found at <http://www.hp.com/rnd/services/index.htm> or by contacting a local ProCurve sales office.

For more information on the ProCurve Networking Adaptive EDGE Architecture, see the ProCurve Networking website at <http://www.hp.com/rnd/architecture/index.htm>.

Appendix A: Policy Enforcement Engine

The ProVision ASIC architecture used in the ProCurve Switch 8200zl, 5400zl, 3500yl and 6200yl Series brings a number of advanced capabilities to the network that offer a highly reliable, robust environment that leads to increased network uptime, keeping overall network costs down. One major feature is the ProVision *Policy Enforcement Engine*, which is implemented in the ProVision ASIC of each interface module.

Policy Enforcement Engine Benefits

The Policy Enforcement Engine has several benefits:

- **Granular policy enforcement**
The initial software release on these products takes advantage of a subset of the full Policy Enforcement Engine capabilities, which will provide a common front end for the user interface to ACLs, QoS, Rate-Limiting, and Guaranteed Minimum Bandwidth controls. Fully implemented in later software releases, the Policy Enforcement Engine provides a powerful, flexible method for controlling the network environment. For example, traffic from a specific application (TCP/UDP port) can be raised in priority (QoS) for some users (IP address), blocked (ACL) for some other users, and limited in bandwidth (Rate-Limiting) for yet other users.

The Policy Enforcement Engine provides fast packet classification to be applied to ACLs and QoS rules and Rate-Limiting and Guaranteed Minimum Bandwidth counters. Parameters that can be used include source and destination IP addresses, which can follow specific users, and TCP/UDP port numbers and ranges, which are useful for applications that use fixed port numbers. Over 14 different variables can be used to specify the packets to which ACL, QoS, Rate-Limiting, and Guaranteed Minimum Bandwidth controls are to be applied.

- **Hardware-based performance**
As mentioned above, the Policy Enforcement Engine is a part of the ProVision ASIC. The packet selection is done by hardware at wirespeed except in some very involved rules situations. Therefore, very sophisticated control can be implemented without adversely affecting performance of the network.
- **Works with Identity Driven Manager (IDM)**
IDM provides the centralized command from the center to define the user policies to be used with each user. The IDM policy requests sent down to the switch are used to set up the user profile in the Policy Enforcement Engine so that the per-user ACL, QoS, and Rate-Limiting parameters can be used from the actual policy defined in IDM.

Wirespeed Performance for ACLs

At the heart of the Policy Enforcement Engine is a memory area called the Ternary Content Addressable Memory (TCAM) that is contained within the ProVision ASIC, along with the surrounding code for the Policy Enforcement Engine.

It is this specialized memory area that helps the ProVision ASIC to achieve wirespeed performance when processing ACLs for packets. In fact, multiple passes through the TCAM can be performed for packet sizes found typically in customers' production networks. For the typical network, the average packet size will tend to be about 500 bytes. When maximum lookups are enabled, the ProVision ASIC performance is optimal for an average packet length of 200 bytes or more, which includes the range of packet sizes in typical networks.

The TCAM can support approximately 3,000 data entries that may be used to represent various traffic controls, including ACLs. For most customers, this quantity of entries will be more than adequate to ensure wirespeed performance for ACL processing. Keep in mind that each ACL entry may consist of multiple criteria, such as a specific IP address and TCP or UDP port number.

In the initial release, the contents of the TCAM are common among the multiple interface modules that a switch may have installed. For example, a ProCurve Switch 5406zl may have up to six interface modules, and a ProCurve Switch 8200zl/5412zl up to 12 interface modules.

Appendix B: Power over Ethernet (PoE)

The IEEE 802.3af standard for PoE enables Ethernet switches to provide power, as well as network signals, for compliant devices over existing CAT-5 cable. Implementing devices that support integrated PoE, such as the ProCurve Switch 8200zl, 5400zl and 3500yl Series, has several benefits:

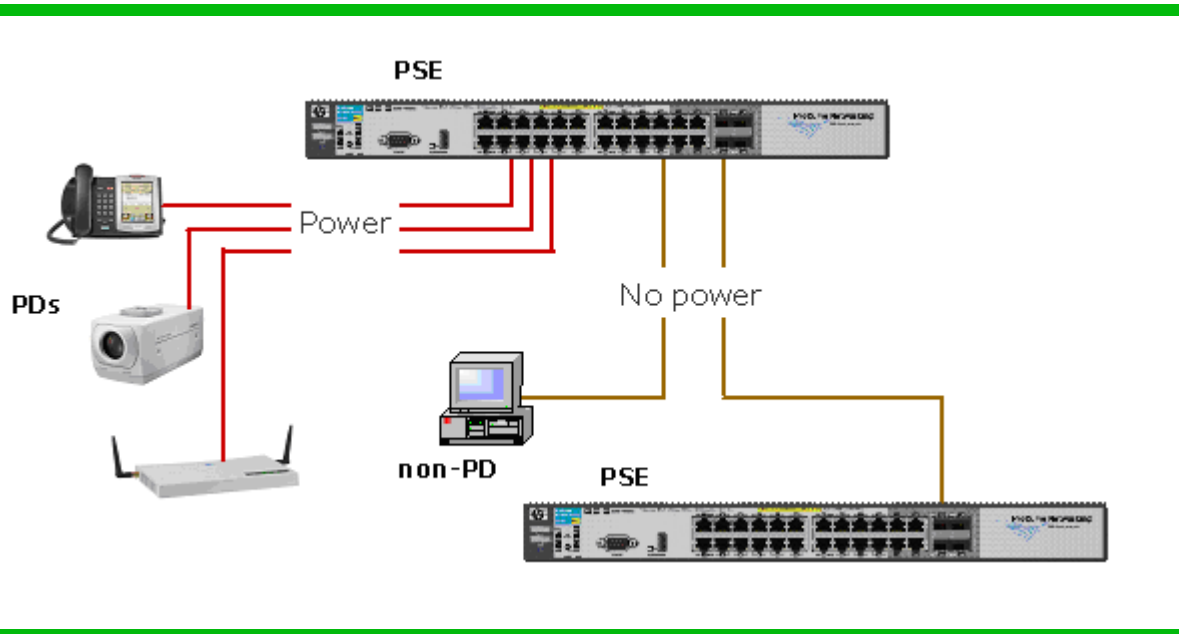
- Lower cost installations because a single cable provides power along with data connectivity
- No need to modify existing network cabling
- Provides a truly international standard for power distribution
- Enables remote reset and power-off capability

PoE device types

There are two types of devices defined for PoE implementations. The first type, called the *powered device* or PD, receives power from the second type, called the *power sourcing equipment* or PSE. Powered devices include any Ethernet device capable of receiving power through a data port such as IP telephones, IP cameras, PDAs and tablet PCs.

Power sourcing equipment, such as ProCurve switches with PoE support, must meet IEEE 802.3af specifications for voltage (47 to 57 volts DC) and wattage (up to 15.4 watts), with further limitations on devices that receive power.

ProCurve switches mentioned in this paper are PoE-enabled to support both IEEE 802.3af-compliant devices, as well as some pre-802.3af standard devices.



Power delivery options

The 802.3af standard provides two options for providing power over CAT 5 Ethernet cable. They are:

1. Unused pairs: This option takes advantage of the fact that 10Base-T and 100Base-TX signals use only two of the four twisted pairs in the cable. In this option, the pins on pairs 1/2 and 3/6 are used for data, and the pins on pairs 4/5 and 7/8 are used for power. The 4/5 pair is twisted together to form the positive supply, and the 7/8 pair is twisted together to form the negative supply.
2. Data pairs: This option provides power over pairs that also provide data, and is necessary because 1000Base-T signals use all four pairs. In this option, either the 1/2 pair or the 3/6 pair can form the positive supply.

The 802.3af standard requires Power Source Equipment (PSE) to choose between the options. It is not possible to mix and match or to use all pairs for power transmission. ProCurve switches provide power over the data pairs. However, the standard requires PDs to be able to accept power over either data pairs or unused pairs.

PoE negotiation

In most networks, PSE will be connected to some devices that support PoE and some that do not. Consequently, in order to prevent damage to non-PoE devices, the 802.3af specification includes a negotiation mechanism between PSEs and the stations connected to them. Under the specification, PSEs apply a slight voltage on the power-delivery pairs. If there is a valid PD on the cable, it will present a resistance and a capacitance.

Typically, this discovery process requires approximately 500 milliseconds. The PSE will apply the full wattage if it discovers a PD, but the flow is current-limited to prevent damage to cable and equipment under fault conditions.

After discovery, a PD must draw a minimum current for the power flow to remain active. If, for example, the PD is unplugged, the PSE will discontinue power supply over the cable. The discovery process will be repeated when the PD is returned to service.

The ProCurve switches described in this document that are PoE-enabled are:

- ProCurve Switch 8212zl: Supports up to 288 PoE ports.
- ProCurve Switch 5412zl: Supports up to 288 PoE ports.
- ProCurve Switch 5406zl: Supports up to 144 PoE ports.
- ProCurve Switch 3500yl-24G-PWR: Supports up to 24 PoE ports.
- ProCurve Switch 3500yl-48G-PWR: Supports up to 48 PoE ports.

Additional PoE Power – External Supplies

The ProCurve 620 Redundant/External Power Supply can be used with up to two Switch 3500s as an external supply providing an extra 398 W of PoE power for each switch. This extra power can be used to provide full PoE power redundancy at 15.4 W for all ports for the 3500yl-24G-PWR. It also can be used to provide full 15.4 W to all 48 ports on the 3500yl-48G-PWR, or full PoE power redundancy at 15.4 W for 24 ports on the 48-port switch.

The ProCurve Switch zl Power Supply Shelf provides a 3U high enclosure to hold up to two extra zl switch power supplies (the same type used in the 8212zl/5400zl switches) to provide additional PoE power to zl switches beyond the PoE power available through the switch's internal power supplies alone. These extra supplies also can be used to provide PoE power redundancy beyond what is possible with switch internal supplies. Large Switch zl VoIP installations are ideal environments for the ProCurve Switch zl Power Supply Shelf.

Support for Pre-802.3af Standard Powered Devices

The ProCurve Switch 8212zl, 5400zl and 3500yl Series also support a number of pre-802.3af standard Powered Devices. These include:

- Cisco 7902G, 7905G, 7912G, 7940G, 7960G IP Phones
- Cisco Aironet 350, 1100, 1200, 1230AG Access Points

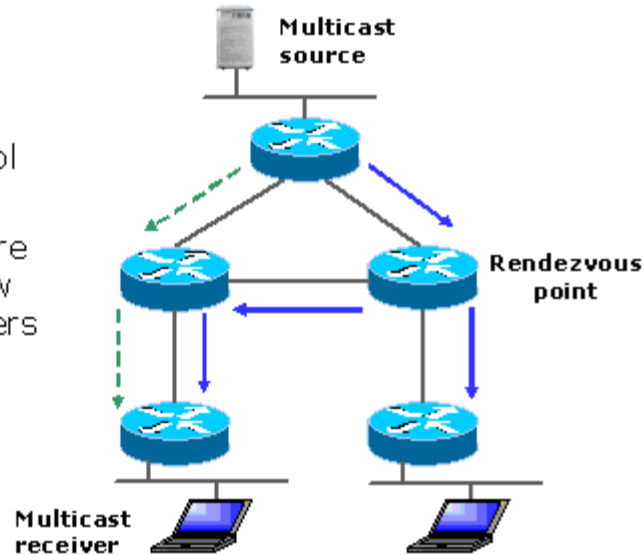
The Cisco pre-802.3af standard IP phones require the use of a straight-through cable for connectivity to the PSE. These IP phones are dependent on the polarity of the DC voltage (a violation of the IEEE 802.3af standard) and will not work properly if the Data Pairs (wires 1 and 2 and 3 and 6) are crossed over anywhere in the cable plant between the IP phone and Power Source Equipment.

Appendix C: PIM-Sparse Mode

In Protocol Independent Multicast-Sparse Mode (PIM-SM), the assumption is that no hosts want the multicast traffic unless they ask for it specifically. In contrast, PIM-Dense Mode (PIM-DM) assumes downstream router membership unless it receives an explicit prune message. PIM-SM is appropriate for wide-scale deployment for both densely and sparsely populated groups, and is the optimal choice for all production networks regardless of size and membership density.

The operation of PIM-SM centers on the use of a shared tree, with a router functioning as a Rendezvous Point (RP), the root of the tree. A shared tree prevents each router from maintaining source and group state information for every multicast source. Regardless of the number or location of multicast receivers, multicast senders register with the RP and send a single copy of multicast data through it to the registered receivers. Also, regardless of the location or number of sources, group members register to receive data and always receive it through the RP.

Shared Tree,
Multicast Protocol
Very efficient in
cases where there
are relatively few
multicast receivers



In order to receive a multicast stream, routers explicitly join the stream by sending “join” messages to the RP. This join message is analogous to a unicast router following a default route to a destination. Effectively, the function of the RP is a place for multicast sources and receivers to meet.

PIM-SM is extremely memory and CPU efficient. Since the only thing most routers need to know is how to reach the RP, memory requirements are reduced greatly. There are several methods that can be used by routers in a PIM-SM domain to learn where to find the RP. Probably the simplest mechanism is statically configuring all routers to reach the RP. However, if the routers are configured statically to an RP and the RP fails, then the multicast network is no longer functional.

Alternatively, the RP can be learned dynamically through the PIM-SM bootstrap mechanism. Since this bootstrap mechanism is dynamic, it allows for network changes and redundancy. The PIM-SM bootstrap mechanism is generally the recommended approach for simplicity and redundancy.

Appendix D: LLDP-MED

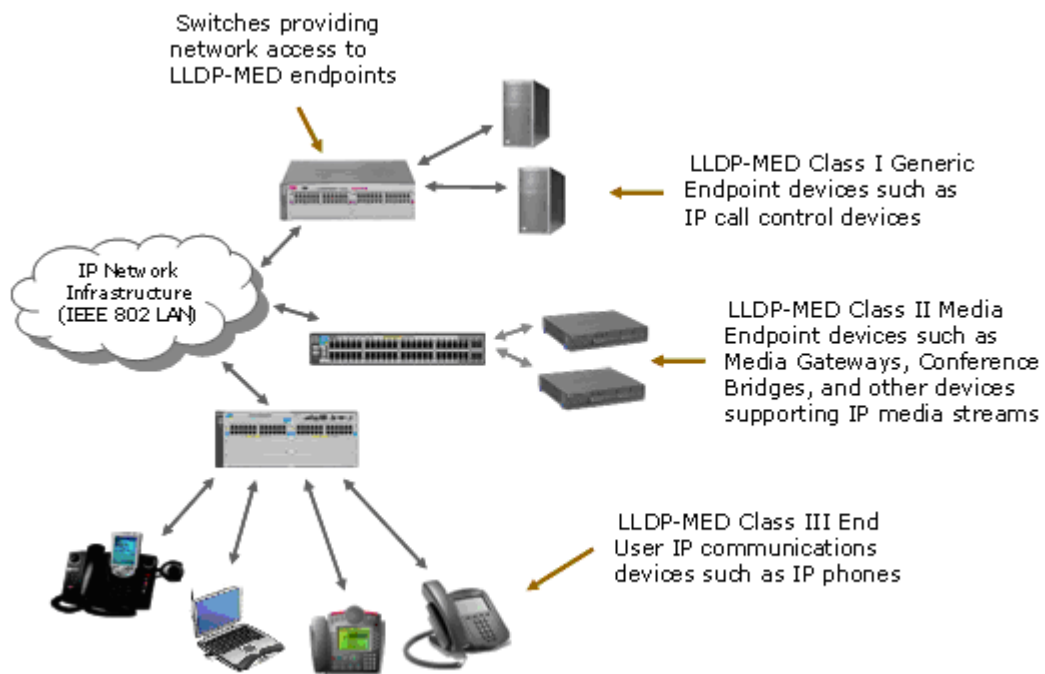
IEEE 802.1AB Link Layer Discovery Protocol (LLDP) provides a standards-based method that enables devices such as ProCurve switches to advertise themselves to adjacent devices and learn about adjacent LLDP devices. An LLDP packet transmitted by a ProCurve switch contains data about the switch and port. The switch advertises itself to adjacent (neighbor) devices by transmitting LLDP data packets out all ports on which outbound LLDP is enabled, and reading LLDP advertisements from neighbor devices on ports that are inbound LLDP-enabled. An LLDP-enabled port receiving LLDP packets inbound from neighbor devices stores the packet data in a neighbor database (MIB).

LLDP provides the base capabilities for network devices, but was not considered sufficient for IP telephony devices. As a result, in 2004, an initiative by Mitel, ProCurve, Avaya and Enterasys was undertaken to enhance LLDP so it could better support IP telephony devices. The development of LLDP-Medium Endpoint Discovery (LLDP-MED) (ANSI/TIA-1057/D6) extended the LLDP standard to support advanced features on the network edge for VoIP endpoint devices

with specialized capabilities and LLDP-MED standards-based functionality. The extensions to LLDP include the specification of additional TLV (Type, Length and Value) entries specifically for VoIP management.

LLDP-MED benefits include:

- Plug-and-play provisioning for MED-capable, VoIP, endpoint devices.
- Simplified, vendor-independent management enabling different IP telephony systems to interoperate on one network.
- Automatic deployment of convergence network policies that includes voice VLANs, Layer 2/CoS priority, and Layer 3/QoS priority.
- Configurable endpoint location data to support the Emergency Call Service (ECS), such as Enhanced 911, 999 and 112.
- Detailed VoIP endpoint data inventory readable via SNMP from the switch.
- PoE status and troubleshooting support via SNMP.
- Support for IP telephony network troubleshooting of call quality issues via SNMP.



LLDP-MED endpoint devices are located at the network edge and communicate using the LLDP-MED framework. Any LLDP-MED endpoint device belongs to one of the following three classes:

- **Class 1 (Generic Endpoint Devices):** These devices offer the basic LLDP discovery services, network policy advertisement (VLAN ID, Layer 2/802.1p priority, and Layer 3/DSCP priority), and PoE management. This class includes such devices as IP call controllers and communication-related servers.
- **Class 2 (Media Endpoint Devices):** These devices offer all Class 1 features, plus media-streaming capability, and include devices such as voice/media gateways, conference bridges, and media servers.
- **Class 3 (Communication Devices):** These devices are typically IP phones or end-user devices that otherwise support IP media and offer all Class 1 and Class 2 features, plus location identification and emergency 911 capability, Layer 2 switch support, and device information management.

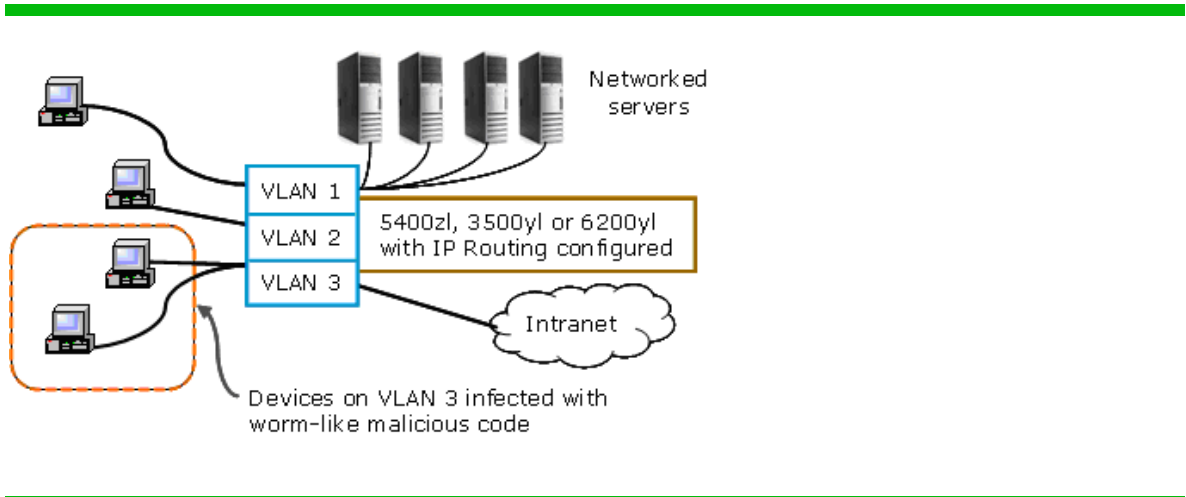
Appendix E: Virus Throttle Security

Virus throttle is based on the detection of anomalous behavior of network traffic that differs from a normal activity. Under normal activity, a computer will make fairly few outgoing connections to new computers, but instead, is more likely to connect regularly to the same set of computers. This is in contrast to the fundamental behavior of a rapidly spreading worm, which will attempt many outgoing connections to new computers. For example, while computers normally make approximately one connection per second, the SQL Slammer virus tries to infect more than 800 computers per second.

Virus throttle works by intercepting IP connection requests in which the source subnet and destination address are different. The virus throttle tracks the number of recently made connections. If a new, intercepted request is to a destination to which a connection was recently made, the request is processed as normal. If the request is to a destination that has not had a recent connection, the request is processed only if the number of recent connections is below a pre-set threshold. The threshold specifies how many connections are to be allowed over a set amount of time, thereby enforcing a connection-rate limit. If the threshold is exceeded, because requests are coming in at an unusually high rate, it is taken as evidence of a virus. This causes the throttle to stop processing requests and to instead notify the system administrator.

This applies to most common Layer 4 through 7 session and application protocols, including TCP connections, UDP packets, SMTP, IMAP, Web Proxy, HTTP, SSL, and DNS — virtually any protocol where the normal traffic does not look like a virus spreading. For virus throttle to work, IP routing and multiple VLANs with member ports must first be configured.

Note that some protocols, such as NetBIOS and WINS, and some applications such as network management scanners, notification services and p2p file sharing, are not appropriate for virus throttle. These protocols and applications initiate a broad burst of network traffic that could be misinterpreted by the virus-throttle technology as a threat.



On the ProCurve Switch 8212zl, 5400zl, 3500yl, and 6200yl Series, virus throttle is implemented through connection-rate filtering. When connection-rate filtering is enabled on a port, the inbound routed traffic is monitored for a high rate of connection requests from any given host on the port. If a host appears to exhibit the worm-like behavior of attempting to establish a large number of outbound IP connections in a short period of time, the switch responds, depending on how connection-rate filtering is configured.

Response options

The response behavior of connection-rate filtering can be adjusted by using filtering options. When a worm-like behavior is detected, the connection-rate filter can respond to the threats on the port in the following ways:

- **Notify only of potential attack:** While the apparent attack continues, the switch generates an Event Log notice identifying the offending host source address (SA) and (if a trap receiver is configured on the switch) a similar SNMP trap notice.

- **Notify and reduce spreading:** In this case, the switch temporarily blocks inbound routed traffic from the offending host source address for a “penalty” period, and generates an Event Log notice of this action and a similar SNMP trap notice if a trap receiver is configured on the switch. When the penalty period expires, the switch reevaluates the routed traffic from the host and continues to block this traffic if the apparent attack continues. During the reevaluation period, routed traffic from the host is allowed.
- **Block spreading:** This option blocks routing of the host’s traffic on the switch. When a block occurs, the switch generates an Event Log notice and a similar SNMP trap notice if a trap receiver is configured on the switch. Note that system personnel must explicitly re-enable a host that has been previously blocked.

Sensitivity

The ability of connection-rate filtering to detect relatively high instances of connection-rate attempts from a given source can be adjusted by changing the global sensitivity settings. The sensitivity can be set to low, medium, high or aggressive, as described below:

- **Low:** Sets the connection-rate sensitivity to the lowest possible sensitivity, which allows a mean of 54 routed destinations in less than 0.1 seconds, and a corresponding penalty time for Throttle mode (if configured) of less than 30 seconds.
- **Medium:** Sets the connection-rate sensitivity to allow a mean of 37 routed destinations in less than one second, and a corresponding penalty time for Throttle mode (if configured) between 30 and 60 seconds.
- **High:** Sets the connection-rate sensitivity to allow a mean of 22 routed destinations in less than one second, and a corresponding penalty time for Throttle mode (if configured) between 60 and 90 seconds.
- **Aggressive:** Sets the connection-rate sensitivity to the highest possible level, which allows a mean of 15 routed destinations in less than one second, and a corresponding penalty time for Throttle mode (if configured) between 90 and 120 seconds.

Connection-rate ACL

Connection-rate ACLs are used to exclude legitimate high-rate inbound traffic from the connection-rate filtering policy. A connection-rate ACL, consisting of a series of access control entries, creates exceptions to these per-port policies by creating special rules for individual hosts, groups of hosts, or entire subnets. Thus, the system administrator can adjust a connection-rate filtering policy to create and apply an exception to configured filters on the ports in a VLAN.

Appendix F: VRRP


Virtual Router Redundancy Protocol (VRRP) is designed to eliminate the single point of failure inherent in the static default routed environment. In a VRRP environment, two or more “virtual” routers cooperate to provide a high availability capability on a LAN. VRRP specifies an election protocol that dynamically assigns routing responsibility to one of the virtual routers on a LAN.

A virtual router consists of a set of router interfaces on the same network that share a virtual router identifier (VRID) and a virtual IP address. One router in the group becomes the VRRP Master and the other routers are designated as VRRP Backups. The VRRP Master controls the IP addresses associated with a virtual router.

The VRRP Master router periodically sends advertisements to a reserved multicast group address. The VRRP Backup routers listen for advertisements and one of the backups will assume the Master role, if necessary. A VRRP router can support many virtual router instances, each with a unique VRID/IP address combination. The election process provides dynamic failover to one of the remaining VRRP Backups should the Master become unavailable.

VRRP is an election protocol that dynamically assigns responsibility for a virtual router on a LAN

Provides High Availability for a default gateway without the need to reconfigure end hosts



The virtual IP address shared by a group of VRRP routers on a given network segment functions as the next-hop IP address used by neighboring hosts. The VRRP Master router simply forwards packets that have been received from hosts using the VRRP Master as the next-hop gateway. The existence of a VRRP master and of one or more VRRP Backups is transparent to the neighboring hosts.

The advantage gained from using VRRP is a default path with higher availability, but that does not require configuration of dynamic routing or router discovery protocols on every end-host.

VRRP on ProCurve switches is interoperable with other routers that support RFC 3768.

VRRP operational aspects include:

- Preemptive mode, which can be disabled to prevent VRRP router-flapping
- Default Advertisement interval of one second
- Default Detection time of 3.6 seconds

XRRP support on 5300xl switch

Similar in concept to VRRP, the ProCurve XL Router Redundancy Protocol (XRRP) provides the equivalent mechanism in the ProCurve Switch 5300xl Series for backup functionality. Like VRRP, XRRP presents a virtual router to the end node connections whose IP and MAC address can transition from the master ProCurve Switch 5300xl to the backup ProCurve Switch 5300xl, in the event the master ProCurve Switch 5300xl interface fails.

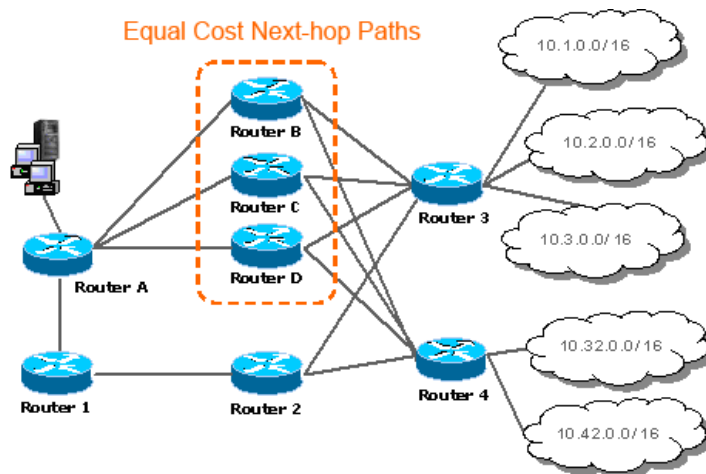
XRRP does not interoperate with VRRP, but can coexist in a VRRP environment without interference.

Appendix G: OSPF Equal Cost Multipath

In Open Shortest Path First (OSPF), if different subnet destinations in a network are reachable through multiple equal-cost paths, the router chooses a single path for each destination subnet. (A trunk is considered a path even though a trunk consists of multiple links).

With OSPF Equal Cost Multipath (OSPF-ECMP), routers support optional load-sharing across redundant paths to the same subnet. All traffic for different hosts in the same subnet may go through a different next-hop router. Multiple paths are balanced based on a destination IP address.

ProCurve's OSPF-ECMP feature is interoperable with OSPF-ECMP implementations from various vendors including Cisco, 3Com and Extreme switches. The ProCurve implementation supports up to four ECMP paths.



The example shows that there are three equal-cost, next-hop paths from Router A to the destination subnets with load-sharing across redundant paths. Router A can route packets based on a destination IP address, as follows:

Destination IP	Next Hop
10.1.1.1	Router C
10.1.1.2	Router B
10.1.2.3	Router B
10.32.1.1	Router B
10.32.1.2	Router C
10.32.1.3	Router B

Appendix H: Troubleshooting

LED Status Indicators for 8200zl

The ProCurve Switch 8212zl management module has various LED status indicators that are described in the following table.

LED	State	Indication
Power	On (green)	The switch is receiving power.
	Off	The switch is NOT receiving power.
Fault	On (orange)	On briefly at the beginning of switch self-test, after the switch is powered on or reset. If on for a prolonged time, the switch has encountered a fatal hardware failure, or has failed its self-test.
	Off	The normal state; indicates there are no fault conditions on the switch.
	Blinking (orange)	A fault has occurred on the switch, one of the switch modules, an individual port, a power supply, or a fan. The Status LED for the module or other device with the fault will flash simultaneously.
Test	On (green)	The switch self-test and initialization are in progress after you have power-cycled or reset the switch. The switch is not operational until this LED goes off. The Self-Test LED also comes on briefly when you "hot-swap" a module into the switch and the module is automatically self-tested.

LED	State	Indication
	Off	Normal operation; the switch is not undergoing self-test.
	Blinking (orange)	A component of the switch has failed its self-test. The status LED for that component, for example a switch module, and the switch Fault LED will flash simultaneously.
DIMM (marked as "DIM") Dual-Inline Memory Module on the Management Module	On (green)	Normal operation.
	Off	DIMM status is unknown.
	Blinking (orange)	If DIMM, Fault, and Self-Test LEDs are blinking, DIMM failed self-test. If DIMM and Fault LEDs are blinking, an operational fault has occurred. If blinking at a fast rate, an operational alert occurred and is unresolved.
Flash	On (green)	Flash Card status is known and fault-free.
	Off	Flash Card status is unknown.
	Blinking (orange)	If Flash, Fault, and Self-Test LEDs are blinking, Flash Card failed self-test. If Flash and Fault LEDs are blinking, an operational fault has occurred. If blinking at a fast rate, an operational alert occurred and is unresolved. For example, Flash Card is not present.
Mgmt	On (green)	A Management Module is present and fault-free.
	Off	The switch is powered off.
	Blinking (orange)	There is a fault on the Management Module.
PoE	On (green)	If any PoE modules are installed.
	Off	If no PoE modules are installed.
	Blinking (orange)	Internal PoE fault, external load fault, or denied PoE power.
Temp	On	Internal temperature is normal.
	Blinking (orange)	An over-temperature condition has been detected.
Fan	On (green)	Normal operation.
	Blinking (orange)	One or more of the switch's fans have failed. The switch Fault LED will be blinking simultaneously.
Internal Power (numbers corresponding to the power supply positions)	On (green)	A power supply is installed in the position in the back of the switch corresponding to the number, and the supply is plugged into an active AC power source. As shipped, the switch has a single power supply in position 1.
	Off	A power supply is not installed in the position corresponding to the number.
	Blinking (orange)	The power supply installed in the position corresponding to the number is not plugged into an active AC power source, or has experienced a fault. The switch Fault LED will be blinking simultaneously.
EPS	On (green)	An external power supply is connected.
	Off	An external power supply is not connected.
	Blinking (orange)	The external power supply has a fault or, is connected but not plugged into AC power.

LED	State	Indication
Modules A, B, C, D, E, F, G, H, I, J, K, L	On (green)	A module is installed in the switch module slot corresponding to the letter and the module is undergoing or has passed self-test. This also occurs when you install a module when the switch is already powered on ("hot-swap").
	Off	A module is not installed in the switch module slot corresponding to the letter.
	Blinking (orange)	The module status LED flashes very briefly when a module is being hot-swapped. If the LED flashes for a prolonged time, the module in the slot corresponding to the letter has failed self-test or encountered some other fault condition.
Modules in PoE Mode	On (green)	PoE is normal for this slot.
	Off	The module in this slot is not a PoE module.
	Blinking (orange)	If any port has an internal hardware failure, it blinks once every 1.6 seconds. If any port is denied PoE power or is detecting an external PD fault, it blinks once every 0.5 seconds.
LED Mode Select	Act (green)	Flickers to show relative activity.
	FDx (green)	Flickers to show relative activity.
	PoE (green)	Indicates which ports are supplying power. <ul style="list-style-type: none"> • If the Mode LED is on, the port is providing PoE power. • If the Mode LED is off, the port is not providing PoE power. • If the Link LED is on, the port is enabled for PoE. • If the Link LED is off, the port is disabled for PoE. • If the Link LED is blinking, the port has an error or the port is denied power due to insufficient power.
	Spd (green)	Indicates speed of operation of each port. <ul style="list-style-type: none"> • If the Port LED is off, the port is operating at 10 Mbps. • If the Port LED is blinking, the port is operating at 100 Mbps. • If the Port LED is on continuously, the port is operating at 1000 Mbps.
	Usr (green)	Reserved for future development.

Table 5: LED Status Indicators for Management /System Support Module

The LED status indicators for the ProCurve Switch zl modules (used in the 8212zl, 5406zl and 5412zl chassis) are described in the following table.

LED	State	Indication
Link	On (green)	Indicates the port is enabled and receiving a link beat signal (for the twisted-pair ports), or a strong enough light level (for the fiber-optic ports) from the connected device.
	Off	Indicates no active network cable is connected to the port, the port is not receiving link beat or sufficient light, or the port has been disabled.
	Blinking (orange)	If the port LED is blinking simultaneously with the Fault LED, the corresponding port has failed its self-test.
Mode	Depending on the mode selected, displays: <ul style="list-style-type: none"> • Network activity information • Whether the port is configured for full duplex operation • Maximum speed operation • Whether or not PoE power is being supplied 	

Appendix I: Part Numbers and Field Replaceable Units

The table below lists various part numbers and Field Replaceable Units (FRUs) for the ProCurve Switch 8212zl and 5400zl Series.

Part Number *	Component
8212zl specific parts (not shared with 5400zl Series)	
J9091A J9091-69001	8212zl Chassis (with Fan Tray) Warranty replacement for 8212zl Chassis (with Fan Tray)
J9092A J9092-69001	8200zl Management Module Warranty replacement 8200zl Management Module <i>without CF card</i>
J9093A J9093-69001	8200zl Fabric Module Warranty replacement 8200zl Fabric Module
J9094A J9094-69001	8212zl Fan Tray Warranty replacement 8212zl Fan Tray
J9095A J9095-69001	8200zl System Support Module Warranty replacement 8200zl System Support Module
5070-2983	Two-post Rack Mounting Kit for Switch 8212zl
5070-3051	Programmed CompactFlash Kit for 8200zl Management Module
5188-6699	RJ45-to-DB9 Console cable for 8212zl

5400zl specific parts	
J8726-69001	Warranty replacement 5400zl Management Module without CF card
J8697-60005	Warranty replacement Fan Tray Assembly for 5406zl
J8698-60005	Warranty replacement Fan Tray Assembly for 5412zl
5070-1056	Programmed CompactFlash Kit for 5400zl Management Module
5184-1894	DB9-DB9 Console Cable for 5400zl Series
5069-8561	Two-post Rack Mounting Kit for 5406zl
5069-8562	Two-post Rack Mounting Kit for 5412zl

Table 13: Part Numbers and FRUs for ProCurve Switch 8212zl/5400zl Series

* Warranty usually requires that the failed part be returned to HP.

The switches normally ship with a rack mounting kit that allows installation into a two-post, 19-inch data communications rack. If installation into a Hewlett-Packard 10000 Series four-post 19-inch rack or a four-post rack meeting the standard EIA unit of measurement is desired, a rack mounting kit (5070-0145) is available providing rails that give sturdy support for the switch along its entire depth. When shipping a system pre-installed in a cabinet, the Power Supplies must be removed and shipped separately due to weight considerations for the mounting kit.




Figure 15: 5070-0145 Rack Rail mounting kit for 8212zl/5400zl Series

The table below lists the part numbers of power cords that can be used with the ProCurve Switch 8200zl and 5400zl Series products. The power cords available for these switches are sized for the increased current that can be drawn meeting the needs of PoE. As a result, these power cords may not be found in a typical environment and power cords “borrowed” from other products will not work in most instances.

Country	ProCurve Switch zl Series	
	875W Power Supply	1500W Power Supply
Hong Kong, Singapore, U.K.	8120-5334	8120-6898
Australia, New Zealand	8120-5335	8121-0871
Europe	8120-5336	8120-6899
U.S., Canada, Mexico	8121-0973	8121-6903
Japan	8120-5342	8120-6903
China	8120-8385	8120-0924
India, South Africa	8120-5341	8120-0915
Switzerland	8121-5339	8121-0915
Denmark	8120-5340	8120-6897

Table 15: Power Cords for ProCurve Switch zl Power Supplies

To find out more about
ProCurve Networking
products and solutions,
visit our web site at

www.procurve.com



© 2007 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

4AA1-2548ENW Revision 1 11/2007