


ProCurve Switch 8100fl Series Reviewer's Guide V2.5



ProCurve Switch 8100fl Series Interconnect Fabric	3
Background	3
ProCurve Networking by HP	3
Product Overview	3
Performance	4
High Availability	4
Security	4
Scalability	4
Chassis and Module Details	5
Modules Common to 8100fl Series	5
Switch 8100fl Redundant Management Module (J8731A)	5
Interface Modules	6
Switch fl 10-Port 100/1000-T Module (J8734A)	6
Switch fl 10-Port Mini-GBIC (SFP) Module (J8735A)	6
Switch fl 2-port 10GbE X2 Media-Flexible Module (J8737A)	6
Fabric Switching Modules	7
Switch 8108fl Redundant Switch Fabric Module (J8729A)	7
Switch 8116fl Redundant Switch Fabric Module (J8730A)	7
Fan Trays, Power Supplies, Console Ports	7
Fan Trays (J8989-61001)	7
Power Supplies (J8732A)	8
Console Ports	8

Hardware Architecture	9
Data Plane	10
Packet Management Block	11
Packet Processors	11
Traffic Manager	12
Backplane Interface Between Traffic Manager and Switching Fabric	12
Switching Fabric	13
Data Plane Review: Packet Walkthrough	14
Control Plane	15
Management Plane	16
Software and System Features	16
Software Architecture	16
Operational Details	16
Initial Configuration	16
System Software Memories	16
Software Installation	16
Software Synchronization on Redundant Management Modules	17
Minimizing System Downtime for Software Upgrades	17
Features and Benefits	17
L2 and L3 Performance and Latency	17
10 Gigabit Performance: Full Mesh, 8 and 16 10GbE ports	18
1 Gigabit Performance: Full Mesh, 80 and 160 GbE ports	18
Notes on Latency Measurements	18
High Availability	18
Optional Redundant Switch Fabric Modules	18
Optional Redundant Management Modules	19
Optional Redundant Power Supplies	19
Link Aggregation Groups (LAGs)	19
IP Routing	19
Spanning Tree	19
Virtual Router Redundancy Protocol (VRRP)	20
Scalability	20
Passive Backplane Chassis	20
Scalable Hardware Architecture	20
Security	20
Filtering – Access Control Lists	20
Protection from DoS Attacks	21
Management Access	21
Summary	21
Appendix	22
Product Ordering Numbers	22
Supported Accessories	22
Mini-GBIC (SFP) Transceivers – Gigabit Connectivity	22
X2 10Gbps Transceivers – 10 Gigabit Connectivity	23
The ProCurve Adaptive EDGE Architecture	23
For more information	25

ProCurve Switch 8100fl Series Interconnect Fabric



Background

This reviewer's guide will help network engineers at computer trade publications, ProCurve resellers and end-user sites evaluate the merits of the ProCurve Switch 8100fl Series Interconnect Fabric.

ProCurve Networking by HP

ProCurve Networking by HP has an extensive line of networking products based on the ProCurve Adaptive EDGE Architecture™, which enables companies to effectively accommodate the enterprise network as it evolves into an anytime, anywhere resource that can adapt to changing business needs.

This ability to provide an adaptable anytime, anywhere resource is made possible through continuous command from the center of the network, coupled with added intelligence at the edge, where users connect and policies are enforced. Control to the edge provides the network infrastructure with secure, robust functionality to support today's and tomorrow's business needs, while ensuring secure, appropriate network access.

The ProCurve Switch 8100fl Series represents a new class of product – Interconnect Fabric – that allows companies to design and interconnect networks based on the Adaptive EDGE Architecture.

Product Overview

The 8100fl Series offers high-performance, high-availability, cost-effective connectivity for intelligent edge devices, while delivering a flexible, scalable, high port density Gigabit and 10 Gigabit Ethernet (10GbE) core networking solution. It complements ProCurve's traditional core product offerings and ProCurve's Intelligent Edge Switches, providing customers with investment protection and utmost choice and flexibility in designing their network.

The 8100fl Series, featuring multiport, modular switches that perform non-blocking, wirespeed, Layer 2 switching, Layer 3 routing and Layer 4 application switching, delivers exceptional functionality and cost-efficiency to the end user. The 8100fl Series offers two chassis configurations: the ProCurve Switch 8108fl and ProCurve Switch 8116fl. Both are capable of leveraging ProCurve Intelligent Edge Switch offerings designed for the ProCurve Adaptive EDGE Architecture.

The ProCurve Switch 8108fl (J8727A) is an eight-slot chassis-based routing switch delivering 119 million pps, wirespeed non-blocking performance for up to eight 10GbE ports (16 ports for 2:1 oversubscribed 10GbE connectivity) or 80 100/Gigabit Ethernet ports. The 8108fl is ideal for medium-to-large networks and provides high-performance and highly available core switching and routing for ProCurve Adaptive EDGE Architecture applications as well as for collapsed backbones, data centers and server farms.

The ProCurve Switch 8116fl (J8728A) is a 16-slot chassis-based routing switch delivering 238 million pps, wirespeed non-blocking performance for up to 16 10GbE (32 ports for 2:1 oversubscribed 10GbE connectivity) or 160 100/Gigabit Ethernet ports. The 8116fl is ideal for large networks and provides high-performance and highly available core switching and routing for ProCurve Adaptive EDGE Architecture applications as well as for collapsed backbones, data centers and server farms.

Performance

The 8100fl Series offers line-rate performance and the ability to support future technologies. The first-generation switch fabric modules support up to 10 Gbps between modules in a non-blocking crossbar switch design. Virtual output queues for each destination port prevent any head-of-line blocking issues, and separate queues and multistage replication for multicast traffic ensure unicast traffic is affected minimally by multicast streams.

In addition, because a portion of the operating system is run locally on each interface module, control plane processing is optimized and communication between the control plane and management module is minimized. Most traffic within the data plane is switched in hardware with no intervention needed by the management module. Control traffic is processed locally on each interface module by dedicated processors and forwarded to the management module through a dedicated internal channel between all interface modules and management modules.

The 8100fl Series provides for rich Quality of Service (QoS) and bandwidth management features that enhance system performance. These include ingress rate limiting, QoS assignment, sophisticated queuing, and egress management, such as minimum bandwidth guarantees and maximum bandwidth shaping.

High Availability

The 8100fl Series offers redundant management and fabric modules to provide high availability. Redundant paths between interface modules and fabric modules are supported, so the interface modules have seamless failover capabilities should the fabric modules receive errors. In addition, the management module supports image and configuration synchronization automatically. Redundant power supplies and redundancy in the fans promote maximum availability. Hardware and software architecture is designed to allow for further enhancements for hitless failover and nonstop software upgrading.

Security

All ProCurve products are built with industry-standard security protocols and offer utmost protection against malicious agents.

To prevent Denial of Service (DoS) attacks to the control and management planes, the 8100fl Series offers rate limiting on the communication channel from interface modules to the active management module. The remote management capabilities of the 8100fl Series provide secure interfaces such as Secure Shell (SSH) and Secure Copy Protocol (SCP).

Scalability

The 8100fl Series offers exceptional scalability with its flexible, chassis-based system. The chassis itself is designed with performance upgrades in mind, with the ability to improve overall system performance in future generations without a chassis upgrade, and maintain non-blocking, line-rate performance. As such, the system will grow in accordance with future business and technical requirements.

The 8100fl Series programmable application-specific integrated circuit (ASIC) design can scale and add software features through future software releases. The ASICs and distributed CPUs are designed with deep inspection capabilities, ready for such future upgrades as IPv6 hardware routing and unforeseen features that can be enabled by software upgrades to the system.

Chassis and Module Details

Each chassis model has slots dedicated for up to two redundant management modules, up to two redundant fabric switching modules, and up to three power supplies.




Figure 1: 8100fl Series chassis overview

Switch 8108fl (J8727A)	Switch 8116fl (J8728A)
8-slot Chassis	16-slot Chassis
Management Module (J8731A)	Management Module (J8731A)
8-slot Fabric Module (J8729A)	16-slot Fabric Module (J8730A)
Two (2) Fan Trays (J8989-61001)	Two (2) Fan Trays (J8989-61001)
1 Power Supply (J8732A)	Two (2) Power Supplies (J8732A)
Two (2) Power Supply Blanks	1 Power Supply Blank
Nine (9) Blank Slot Covers	Nineteen (19) Blank Slot Covers

Each ordered chassis comes complete with the above items. A choice of interface modules, extra management modules, extra fabric modules, and extra power supplies can be ordered as separate line items. See the Appendix on page 22 for part numbers and details.

Modules Common to 8100fl Series

Switch 8100fl Redundant Management Module (J8731A)

The 8100fl Series allows for two management modules in a single chassis. One is included with each chassis and an optional second management module can be installed for redundancy. Management modules can be inserted in slot MMA (Management Module A) and slot MMB (Management Module B). The management modules have a built-in 10/100Base-T Ethernet port, which is not a part of the switching data plane and is used for management purposes only. This management port is configured as MDI, requiring the use of a crossover cable if connected directly to most PCs with 10/100 Ethernet ports. If connected to another switching device that does not provide Auto-MDIX detection, a straight-through cable will be required. The management modules communicate with other modules in the chassis through a 100 Megabit full-duplex Ethernet control plane. These modules have a PCMCIA card slot for future use (not enabled in Version 1 software). The module has two LEDs on the bottom, one to indicate system

status (a green/orange bicolor LED to indicate booting, running or fault conditions) and one to indicate whether it is the active or inactive management module for the system (green LED).




Figure 2: Modules common to both 8108fl and 8116fl chassis

Management modules have a PowerPC MPC7410 CPU running at 400 MHz, 1 GB of SDRAM (Synchronous Dynamic RAM) for expansion of an uncompressed operating system and lookup tables. In addition, a 512 MB Compact Flash device is used for nonvolatile storage of system software and other system parameters.

Interface Modules

Each interface module has a PowerPC IBM405 CPU running at 266 MHz for local processing and lookups using 256 MB of SDRAM for expansion of uncompressed operating system software and local lookup tables. Each module has a packet buffer of 128 MB of RL-DRAM (Reduced Latency Dynamic RAM), which is used to store packets organized into Virtual Output Queues on a per egress port basis.

Switch fl 10-Port 100/1000-T Module (J8734A)

This module has 10 ports of 100/1000Base-T for up to line-rate Gigabit connectivity. Each RJ-45 port has embedded LEDs: a bicolor LED to indicate link state and speed (green for Gigabit, orange for 100 bps) and another to indicate port activity. The module has a bicolor green/orange LED on the bottom to indicate system status (booting, running or fault conditions).

Switch fl 10-Port Mini-GBIC (SFP) Module (J8735A)

This module delivers 10 ports of full-duplex, line-rate Gigabit connectivity. It supports ProCurve Mini-GBIC accessories (B-versions listed in the Appendix on page 22), including the 1000Base-T Mini-GBIC device (J8177B). Each port has two LEDs, one to indicate link state and one to indicate port activity. The module has a Status LED, a bicolor green/orange LED on the bottom to indicate booting, running or fault conditions.

Switch fl 2-port 10GbE X2 Media-Flexible Module (J8737A)


This single two-port, X2 form factor 10 Gigabit module accepts X2 10Gbps transceivers, a 2:1 oversubscribed, full-duplex 10 Gigabit module. Each port has link and activity LEDs. The module has a bicolor green/orange LED on the bottom to indicate system status (booting, running or fault conditions).

The throughput capacity of the J8737A is 10Gbps, regardless of whether one or both ports are linked. This provides connectivity to multiple 10Gbps links, and automatically balances ingress traffic between the two ports.

This J8737A module¹ uses the same X2 transceivers used in other ProCurve products:

- J8440A CX-4 copper (0.5 meter to 15 meter reach with CX-4 cable)
- J8439A Optical Media Converters (OMC) for CX-4 ports (up to 300 meter reach)
- J8436A SR (Multimode fiber for 2 meters to 300 meter reach)
- J8437A LR (Singlemode fiber for 2 meters to 10 km reach)
- J8438A ER (Singlemode fiber for 2 meters to 30 km/40 km reach)²

Fabric Switching Modules



Switch 8108fl Redundant Fabric Module J8729A



Switch 8116fl Redundant Fabric Module J8730A

Figure 3: Fabric Switching Modules

In contrast to the modules that are common to both chassis, the fabric modules are unique to each system. Unlike other modules, there are no front accessible ports on fabric modules. Each fabric module has a green “Active” LED and a bicolor green/orange “Status” LED to indicate booting, running or a fault condition.

Fabric modules have a local CPU, a PowerPC IBM405 running at 266 Mhz, 256 MB of SDRAM and 16 MB of SRAM. Fabric modules boot up and download a copy of their operating system from the active management module.

Switch 8108fl Redundant Switch Fabric Module (J8729A)

This fabric module for the eight-slot chassis can be inserted into either of the two center slots and provides for near-hitless failover (0.2 second failover) for switched or routed traffic when used with a second fabric module.

Switch 8116fl Redundant Switch Fabric Module (J8730A)

Taller in size than the 8108fl Redundant Switch Fabric Module, this fabric module for the 16-slot chassis can be inserted into either of the two center slots and provides for near-hitless failover (0.2 second failover) for switched or routed traffic, when used with a second fabric module.

Fan Trays, Power Supplies, Console Ports

Common components to both the eight-slot ProCurve Switch 8108fl and 16-slot ProCurve Switch 8116fl are the fan trays and power supplies.

Fan Trays (J8989-61001)

Each system comes with two fan trays. Each fan tray contains three variable-speed 8” fans and each fan is monitored by the system. In the event of a single fan failure an SNMP trap and

¹ The earlier J8733A fixed-port LR 10G module was discontinued in early 2006. The J8736A one-port 10G X2 Media Flexible module was discontinued in May 2007, in favor of this two-port J8737A.

² ProCurve 10-GbE X2-SC ER Optic (J8438A) was introduced in the Spring 2006. Links longer than 30 km for the same link power budget are considered engineered links. Attenuation for such links needs to be less than the minimum specified for B1.1 or B1.3 singlemode fiber.

event log entry are generated. Each fan tray is designed to cool its respective half of the chassis (containing fabric, management and interface modules).

Each module (management, interface and fabric) contains an onboard temperature sensor. The management module will generate an SNMP trap and event log entry when the temperature passes a user-configurable warning threshold (initially set at the factory for an onboard temperature of 72° C). In the event of a module passing a factory-configured critical threshold (78° C), power to the affected module is shut down to prevent damage to the components.

Power supplies contain their own individual fans and are unaided by the airflow provided by these fan trays.

Additional spare fan trays can be ordered directly from an authorized ProCurve Spare Parts Reseller or in North America from www.hp.com/buy/parts. Search by part number: J8989-61101.

Power Supplies (J8732A)

Each chassis can support up to three power supplies. Each system comes complete with the minimum number of supplies for a fully loaded chassis. The eight-slot chassis comes with one power supply and space for two more, allowing for a fully redundant power supply configuration. The 16-slot chassis comes with two power supplies and room for one more, allowing for an N+1 redundancy configuration, sufficient for a single power supply failure. Power supplies are auto-sensing for 100-240 VAC operation. The rated output of each power supply is dependent on the input voltage (1200 watts @100 VAC or 1500 watts @ 240 VAC). Each base chassis is supplied with three country-specific power cords and has C19 connectors at the product end (notable for rack cabinets with integrated power distribution systems).




Figure 4: 8100f Series console and management ports

Console Ports

Each chassis contains two RJ-45 serial console ports as shown in Figure 4, one for Management Module A and one for Management Module B. These ports are used to directly connect an RS-232 serial management console to the switch. The management module that is active (indicated by the "Active" LED or the CLI command "show redundancy") dictates which console port should be utilized. The user can employ the console port only for RS-232 out-of-band communication; it cannot be used for a Telnet connection (as it is a serial RS232, not an Ethernet connection). For network connections (Telnet, SSH, TFTP, SCP) to the management module, use the 10/100Base-T port located on each of the management modules. An RJ45-to-DB9 adapter is included with every chassis for console port access (using any standard straight-through network cable).

Hardware Architecture

This section is an excerpt from the [ProCurve Switch 8100fl Series Interconnect Fabric Technical Brief](#) (2/2006) located at www.procurve.com in the Reference Library.

A ProCurve Switch 8100fl Series chassis consists of one or two management modules, one or two fabric modules and one to 16 interface modules. When a second management module is present, the management modules are redundant components that provide failover of system management. One is active and the other is standby. Configurations and images are synchronized and each module is actively monitoring the state of the other.

Similarly, when a second fabric module is present, hardware and software components are monitoring the state of each. The standby fabric module always is maintaining the state of the data path even though it is not actively forwarding traffic. This allows for rapid failover because the secondary fabric module does not need to re-establish state if the primary fabric module fails.

In terms of system logic, the 8100fl Series is represented by three major subsystems: the data plane, the control plane and the management plane.

The data plane contains the elements that forward network traffic. The control plane dynamically configures and monitors the data plane and implements network protocols. The management plane provides user and network management interfaces and statically configures and monitors the entire system.

Figure 5 provides a system-wide view of all three planes. The data plane is confined to elements on the interface modules and fabric modules. Note that there is a portion of the data plane that interfaces with the local CPU on the interface modules, particularly for network control traffic such as BPDUs or OSPF updates. The control plane is distributed throughout all modules in order to provide optimum processing. The management plane is located mostly on the management modules, though some monitoring components (for statistic counters) are located on other modules. Note that the management plane has completely out-of-band interfaces – separate from the control and data planes for access to configuration and system control. These separate interfaces include the management port on the active management module and the internal Ethernet channel between the interface, fabric and management modules, which does not follow the same path as the data traffic.

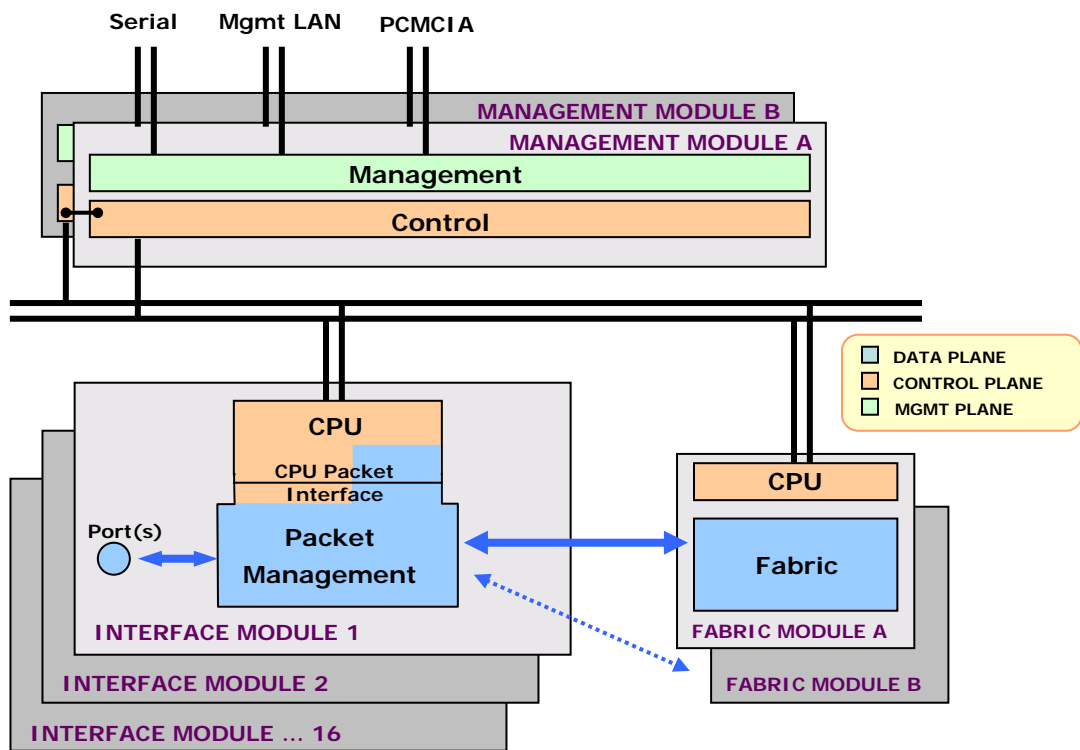


Figure 5: 8100fl Series architecture: Data/Control/Management Plane

Data Plane

The 8100fl Series data plane provides all of the processing and forwarding of packets throughout the system, which are implemented primarily in hardware via ASICs. Only rarely — and for specific Internet Protocol (IP) functionality, such as fragmentation — will packets be intercepted and forwarded via software. In such cases, all forwarding is handled locally on the ingress interface module, not by a management module.

The data plane components are located on the 8100fl Series interface modules and fabric modules. Each interface module has redundant data plane connections to both fabric modules. Because the redundant fabric modules are in either an active or standby state, only one of the connections is active. However, a standby fabric module is always maintaining the dynamic state of the active module for rapid failover (0.2 seconds in the event of a failure).

Figure 6 details the components of the 8100fl Series data plane. On the interface module, the packet management block consists of packet processors — one for ingress and one for egress — and a traffic manager. On the fabric module, the fabric consists of a crossbar switch and a bandwidth manager. The data path is completely full duplex. Packets simultaneously transit the components of the data plane in both directions with no performance degradation. The ports on an interface module include components not shown, such as physical layer (PHY) optics and media access control (MAC).

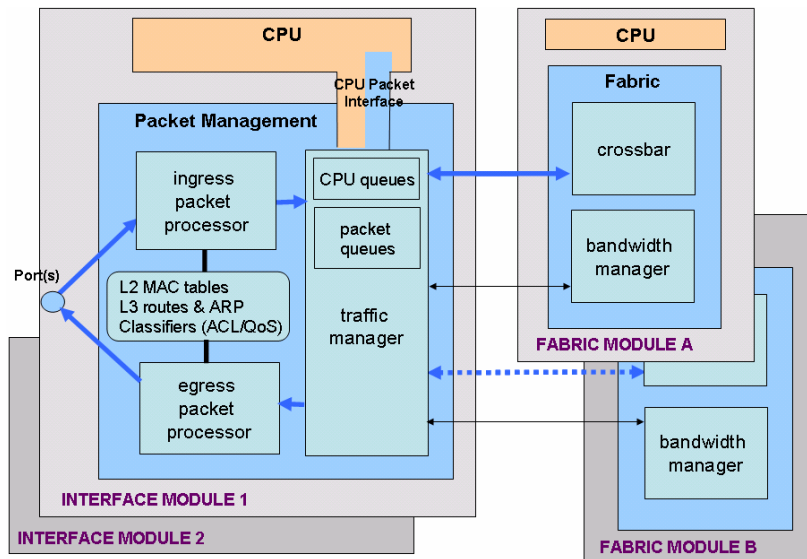


Figure 6: Components of the 8100fl Series data plane

Packet Management Block

The packet management block is responsible for receiving and sending packets from ports and either managing their access to the fabric or sending them to the local CPU for software processing. This block determines the destination of all forwarded packets by performing line-rate routing and bridging. In the event of congestion on the egress port, packets are queued in very large, high-speed packet buffers. These queued packets can be treated with highly granular QoS features.

Packet Processors

The ingress packet processor ensures the security of inbound packets by processing their headers at line rate, enforcing and assigning their VLAN, and preparing them for bridging and routing processing.

All Layer 2 bridging and Layer 3 routing forwarding decisions are made in hardware by the ingress packet processor. Every interface module has this information (distributed by the management module); therefore, local bridging and routing functions are distributed throughout the system. Because forwarding decisions are made locally on the interface modules and not by the central management module, system performance is optimized.

In the event that new Layer 2 MAC addresses must be learned for bridge updates, the hardware automatically refers the MAC addresses to the CPU for processing.

Packets are matched to rich access control entries by the ingress packet processor. The first release allows 1,000 entries with resultant actions to accept or deny. The design allows for a variety of potential actions on matched Access Control Lists (ACLs), such as rate limit, assign QoS, policy route and mirror.

The egress packet processor updates various Layer 2 and Layer 3 packet headers, such as VLAN tags, MAC addresses, 802.1p and IP-TOS Class of Service (CoS), IP TTL, etc. It also performs final replication of packets if needed for multicast.

The functionality of the ASIC-based packet processors is firmware upgradeable. This enables future proofing and investment protection because new features can be added over time. For example, the ingress packet processors are capable of defining robust access control entries that control the assignment of QoS classes or the means by which packets are rate limited with highly granular controls.

The packet processors perform all functions at line rate (10 Gbps).

Traffic Manager

The traffic manager contains large, high-speed packet buffers that are utilized in the event that the destination (egress) port is congested. However, an egress port will become congested only if high levels of traffic from multiple ingress ports are sent to a common egress port and the sum of the traffic is greater than the media speed of that port. Because packets are buffered before the fabric on ingress, the architecture uses advanced technology, called “virtual output queues,” which allocates queues for every egress port in order to prevent head-of-line blocking on the input port. The virtual output queue architecture allows for rich bandwidth management and QoS features.

The hardware architecture can accommodate up to eight queues per egress port. The current release of software (V2.0) provides five forwarding queues per port and allows for Differentiated Services (DiffServ) per-hop behaviors, such as three assured forwarding classes, an expedited service class and default class. Each queue has weighted random early detection (WRED) with three levels of drop precedence per queue.

The size of each queue is allocated dynamically by hardware as buffer space is required. Each interface module provides 128 MB of memory for this purpose. This allows queues with more congestion to obtain a larger share of the buffer when needed, resulting in an optimum use of the available resources. In summary, the total buffer space per egress queue, system-wide, will vary between 64 KB to 25 MB (2.5 MB on average).

The selection of the queue is determined by the ingress packet processor based on packet header processing.

Packets destined for the CPU, such as protocol packets for the control plane, are treated by dedicated QoS queues and rate-limiting features. This provides a secure interface to the CPU and prevents DoS attacks on the control plane, affecting regular switch data plane traffic.

In addition, the traffic manager has separate multicast queues to avoid unicast traffic blocking and provide efficient replication of multicast.

Backplane Interface Between Traffic Manager and Switching Fabric

Interface modules are connected physically via the passive chassis backplane to the fabric modules. In order to provide redundancy, each interface module is connected via separate channels to each fabric module. The channel to a fabric module consists of two types of interfaces: a packet data interface and a control interface, as shown in Figure 7: Backplane Interface between Traffic Manager and Switching Fabric on page 13.

The effective throughput of the switching fabric is 10 Gbps full duplex (20 Gbps net) per interface module. The switching fabric uses a 2x speedup design, providing a raw switching capacity of $8 \times 20 \times 2 = 320$ Gbps for the eight-slot fabric module and $16 \times 20 \times 2 = 640$ Gbps for the 16-slot fabric module. This 2x speedup allows a packet to enter and exit the switch fabric in both directions at 10 Gbps wire rate speeds.

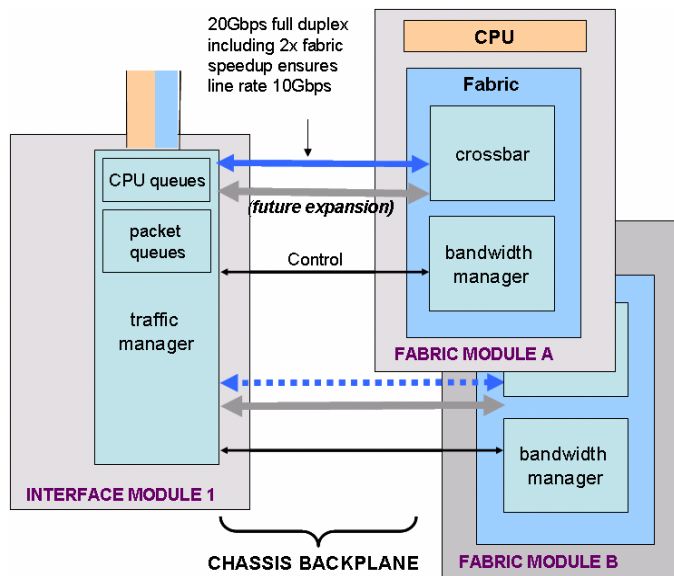


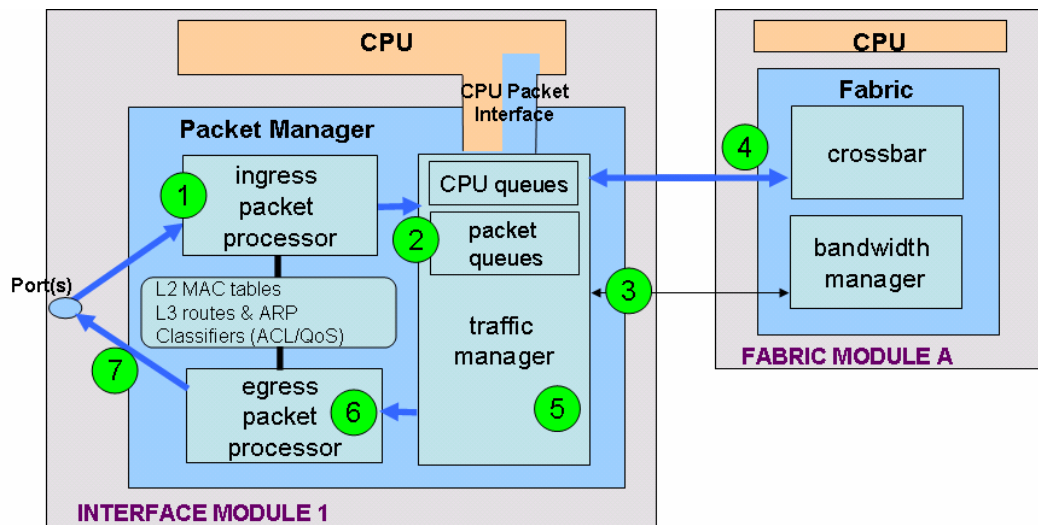
Figure 7: Backplane Interface between Traffic Manager and Switching Fabric

For investment protection, the chassis backplane already is built with additional interfaces to every interface and fabric module to accommodate future capacity expansion. This will allow for line-rate forwarding of higher port count interface modules in the future.

Switching Fabric

A bandwidth manager coordinates with all interface module traffic managers for global switch scheduling. Each traffic manager provides continuous status on all queues to bandwidth managers on both the active and standby fabric modules. The bandwidth manager determines which traffic manager is granted access to the switch crossbar via sophisticated bandwidth allocation algorithms. The algorithms allow for features such as guaranteed minimum bandwidth provisioning in 1 Mbps increments and also, weighted fair queuing or strict priority scheduling. The ASIC design can support future capabilities, such as egress traffic-shaping (maximum bandwidth), through a software update. The crossbar provides the non-blocking fabric to interconnect all traffic managers for all interface modules.

For efficiency and minimal use of crossbar bandwidth, the fabric also performs first-stage packet replication for multicast packets sent to different interface modules (one packet into fabric and one packet out to each interface module). The second stage of multicast replication is conducted by the receiving traffic manager for the set of egress ports, and the final stage of multicast replication is conducted by the egress packet processor to a single egress port. This multistage replication of multicast provides the most efficient use of system-wide resources and bandwidth.



1. Ingress packet processing performed
2. Traffic Manager queues packet
3. Bandwidth Manager coordinates scheduling
4. Crossbar fabric interconnects Interface Modules
5. Traffic Manager on egress card receives packet
6. Egress packet processing performed
7. Packet transmitted to physical interface

Figure 8: Data Plane Packet Walkthrough Overview

Data Plane Review: Packet Walkthrough

Figure 8 shows the numbered steps as data packets traverse the switch.

(1) The packet is received by an interface module port via physical layer interfaces (PHY), such as optics and media access control (MAC), and then sent to the ingress packet processor.

- Ingress header processing (e.g., VLAN assignment)
- Forwarding lookup (e.g., Ethernet bridging, IP routing)
- In-line classification (e.g., ACLs, filters, policy for QoS, and future rate limiting)
- Policing (future rate limiting with three-color marker and 1 Mbps granularity)
- Assign CoS (based on 802.1p, DiffServ, etc.)
- Link aggregation and ECMP (Equal Cost Multi-Path) trunk selection

(2) The traffic manager receives QoS prioritized packet from the input packet processor and enqueues the packet.

- Packets assigned to virtual output queues, up to five CoS queues per egress port
- Separate multicast queues avoid unicast traffic blocking
- Dynamic memory assignment to individual queues from large memory pools
- Congestion management (WRED) performed on individual queues based on packet drop-precedence ("color") with three profiles per queue

(3) Bandwidth manager coordinates with all interface module traffic managers for global switch scheduling.

- Each traffic manager provides continuous status on all queues to active and standby bandwidth managers on each fabric module

- Bandwidth manager determines which traffic manager is granted access to crossbar fabric via sophisticated bandwidth allocation algorithm
- Guaranteed minimum bandwidth provisioning in 1 Mbps increments, also weighted round robin and strict priority scheduling

(4) Crossbar provides a non-blocking fabric to interconnect all traffic managers.

- 20 Gbps bidirectional connectivity (effective 10 Gbps full duplex) to each traffic manager with future capacity expansion built into chassis
- Performs first-stage packet replication for multicast packets sent to different interface modules

(5) Traffic manager on egress interface module receives packet from switch fabric.

- Replicates multicast packets to individual ports, if required

(6) Traffic manager sends packets to egress packet processor.

- Replicates multicast packets to each physical interface, if required (e.g., individual VLANs)
- Modifies packet header appropriately for link type (e.g., VLAN tags, MAC addresses, 802.1p and IP TOS, IP TTL)

(7) Egress packet processor sends packet to the egress port MAC and PHY for transmission on the link.

Control Plane

The 8100fl Series control plane provides a channel for network protocol packets. It also configures and monitors data plane components. The interface with the data plane is through a secure and QoS-controlled packet interface to allow the capture of packets intended for network protocols or management. The control plane interfaces with the management plane for configuration, monitoring and image updates. Control plane components are distributed among the management modules and the interface modules. The modules are redundantly interconnected to each management module for high availability. In addition, the management modules are directly and redundantly connected to facilitate failover features of the management modules.

The control plane software on the management module, among other things, runs routing protocols, such as OSPF and RIP, to learn and distribute routes and maintains a routing information base (RIB). It also distributes all the active routes from the RIB to all interface modules. A proprietary protocol is used within the system to ensure all interface modules are always synchronized with the latest information in the RIB on the management module. The control plane software on the interface modules populates the relevant data plane hardware tables with the routes distributed from the management module. This ensures the data plane traffic is hardware-forwarded by the local ASICs. The interface modules also maintain a copy of all the routes in software tables. This allows the interface module CPU to handle all exception traffic locally without having to send the packets to the management module.

The control plane software on the management module also ensures that Layer 2 MAC addresses for bridging are distributed to interface modules. The control plane software on the interface modules populates the appropriate hardware tables with these MAC addresses, which ensures line-rate, hardware-based forwarding.

Also, the control plane software on the management module translates user-specified policies and actions (ACLs, QoS policies, etc.) into data plane-specific structures, and then transfers them to the interface modules, which, in turn, populate the appropriate local hardware tables.

The management module also keeps track of the health of all interface modules, fabric modules, the standby management module and all the processes within each of those modules. It takes appropriate actions to handle failures on any of these modules. In addition, it keeps track of the state of all the data ports within the system and ensures all interface modules always have the latest state information about all other modules within the system.

Management Plane

The management plane provides overall system management, including configuration and monitoring. It includes redundant external interfaces, such as serial and LAN. The management LAN and serial interfaces are completely separate and out-of-band from the interface module ports. This allows for a very secure and reliable dedicated interface to the management plane (regardless of the state of the data plane). The management plane also has internal management protocol interfaces such as Telnet, SSH and TFTP, among others, that can be reached in-band via network connections on the regular interface module ports.

Software and System Features

Software Architecture

The 8100fl Series features a distributed software system, designed to ensure high availability and performance. Every card, including the interface modules and redundant fabric modules, runs a real-time operating system.

The distribution of software to multiple CPUs optimizes tasks and CPU load to relevant cards. The management module CPU provides overall system management and does not forward data packets. The interface module CPU provides exception packet processing and local hardware control.

Operational Details

Initial Configuration

The Switch 8100fl Series is designed for the core of the network, not the edge. Rather than being a “plug-and-play” switch, the 8100fl requires some initial configuration to be properly connected to other switches or devices. By default, the ports are disabled in a “shutdown” state. Once configured, the individual ports must be enabled by issuing a “no shutdown” command. Remote management processes, such as Telnet or SSH, must be enabled, and virtual terminal lines (VTY) must be created before anything other than a direct serial console session can be initiated.

System Software Memories

The 8100fl Series management module has a 512 MB Compact Flash memory (device name “flash:”) used for intermediate storage of files during software upgrade procedures as well as storage for error logs generated in the case of a software error. These logs can be offloaded and sent to ProCurve Support for further analysis. Operational logs also are stored on this device, if no syslog server has been configured.

Each management module reserves an area on this Compact Flash device for two banks of system software – Bank-A and Bank-B. Each can hold a different version of the system software. These different versions of software can be run using the same configuration file. Version 2 software allows the storage of multiple configurations on the “flash:” device. Future software enhancements will add the capability to store a specific configuration file and associate it with a specific system software image.

Each interface and fabric module has 256 MB SDRAM for loading a local image of the operating system, as well as keeping locally relevant lookup information. This local SDRAM is managed by the local CPU on each module.

Software Installation

To upgrade the system software on an 8100fl Series, the system image is first transferred to the Compact Flash on the active management module. In-band or out-of-band network transfers can use either TFTP or SCP. Once the image is copied to the flash device, it is installed in either Bank-A or Bank-B on each management module. The system is then configured to boot from one or the other bank. (This setting is stored in nonvolatile RAM for the next reboot cycle and initially factory-set for Bank-A). The operating system is extracted from the designated bank when the management module reboots. Once the management module is running, each interface module receives its software over the management plane.

Software Synchronization on Redundant Management Modules

The “image install” command has options to install software on individual management modules as well as specific banks, although users cannot install software onto the running bank. System software installation by default is targeted to the opposite flash bank to the one currently in use. For example, if the management module was booted from Bank-A, installation of new software will be to Bank-B and vice versa.

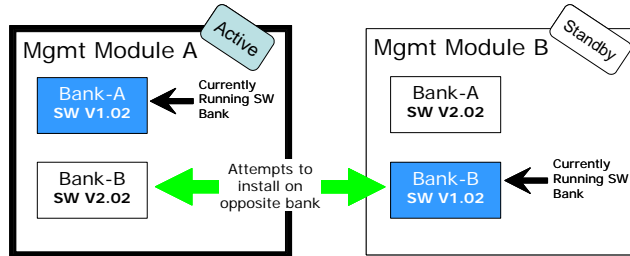


Figure 8: Management Module Software Bank Synchronization

When a system has two management modules installed, the default version of the “image install” command attempts to install the software on the same A or B bank on both management modules as shown in Figure 8. Typically, both management modules would be running from the same bank; both from Bank-A or both from Bank-B. If there is a mismatch of the two management modules, an error message is displayed.

Assuming Management Module B is the standby module (not active) the command to reboot Management Module B to Bank-A is:

```
boot system management b bank-a
```

Because Management Module B is the standby module, no interruption to the Switch 8100fl service will occur by rebooting this standby module (i.e., the active management module is still running the system). Once Management Module B is up and running on Bank-A, the “image install” command can be issued again to properly synchronize both management modules with the same software; in this case, into Bank-B.

Minimizing System Downtime for Software Upgrades

By upgrading a standby management module and rebooting it first – i.e., keeping this second management module as the standby, now running the newer software – a scheduled system downtime will now be limited to only a few minutes. The 8100fl software requires interface modules to be reset whenever a management module switchover occurs. By upgrading the software on the standby module, users can reboot the second management module while the switch continues to operate on the previous version of software. Once the second management module is up and running, users can then switch over to it – and reboot the interface modules to the newer software – minimizing the total system downtime.

Future software improvements will add hitless failover and hitless software upgrading capabilities.

Features and Benefits

L2 and L3 Performance and Latency

- 10 Gbps full-duplex bandwidth between modules
- Line-rate, non-blocking performance
- L2 and L3 hardware forwarding
- Deep 128 MB dynamically allocated packet buffers per interface module

The following measurements were performed by ProCurve using test equipment manufactured by Ixia Communications (www.ixiacom.com). In these tests, eight or 16 ports were used for 10 Gigabit testing, and 80 or 160 ports were used for Gigabit testing.

“Full Mesh” for this performance test is defined as:

- Port 1 receives packets destined to hosts on all other ports (2, 3, 4, 5, ...)
- Port 2 receives packets to all other ports (1, 3, 4, 5, ...)
- Port 3 receives packets to all other ports (1, 2, 4, 5, ...) and so on

10 Gigabit Performance: Full Mesh, 8 and 16 10GbE ports

Packet Size	L2 & L3 Throughput	Full Mesh % drops	10 Gbps, 8-ports		10 Gbps, 16-ports	
			L2 Latency (μ secs)	L3 Latency (μ secs)	L2 Latency (μ secs)	L3 Latency (μ secs)
64	100 %	0%	8.4	8.5	8.4	8.5
128	100 %	0%	8.5	8.6	8.5	8.8
256	100 %	0%	8.6	8.7	8.6	8.9
512	100 %	0%	8.7	8.8	8.7	8.9
1024	100 %	0%	8.8	8.9	8.8	9.0
1280	100 %	0%	8.9	9.0	8.9	9.2
1518	100 %	0%	9.0	9.2	9.1	9.2

With the introduction of the J8737A 2-port 10GbE module, since the module is 2:1 oversubscribed, these 10 Gigabit performance values only apply to the use of one port. When both ports are linked, the module automatically allows for up to an aggregate of 10Gbps through either port and balances ingress traffic between both ports. The throughput from ports through the switch fabric is a maximum of 10Gbps.

1 Gigabit Performance: Full Mesh, 80 and 160 GbE ports

Packet Size	L2 & L3 Throughput	Full Mesh % drops	1 Gbps, 80-ports		1 Gbps, 160-ports	
			L2 Latency (μ secs)	L3 Latency (μ secs)	L2 Latency (μ secs)	L3 Latency (μ secs)
64	100 %	0%	23.3	23.8	24.0	24.1
128	100 %	0%	23.9	24.6	24.0	24.3
256	100 %	0%	24.3	25.1	24.4	24.8
512	100 %	0%	25.3	25.9	25.8	26.1
1024	100 %	0%	26.0	27.2	27.1	28.3
1280	100 %	0%	27.2	28.4	28.4	29.4
1518	100 %	0%	28.2	29.2	28.9	30.5

Notes on Latency Measurements

The latency times 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 RL-DRAM 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.

High Availability

Optional Redundant Switch Fabric Modules

Redundant fabric switch modules are front accessible, user-replaceable modules. Traffic from interface modules is sent simultaneously to both fabric modules in the event of a failure of the primary switch fabric for fast failover (0.2 seconds) as detailed in the *Hardware Architecture* section.

Optional Redundant Management Modules

A redundant management module allows for recovery of a failed primary management module and assists in the software upgrading of an 8100fl Series chassis.

The failover time is one to three minutes, which is the time it takes to reset interface modules when a redundant management module takes over upon failure of the active management module. Improvements to this failover time will be available in future software updates. See *Minimizing System Downtime for Software Upgrades* on page 17, for more information on how to reduce the overall downtime of a switch using redundant management modules.

Optional Redundant Power Supplies

Up to three redundant power supplies allow power feeds to come from different power main circuits. The eight-slot model requires one power supply to handle a fully loaded configuration, with an extra power supply providing full power redundancy. The 16-slot model requires two power supplies to handle a fully loaded configuration, with a third power supply providing an N+1 redundancy configuration. Each of the power supplies are load-sharing; adding more lessens the load on an individual power supply and lengthens the life of a given power supply.

Link Aggregation Groups (LAGs)

Referred to as “port trunking” in other ProCurve products, support for Link Aggregation Groups (LAGs) is provided for up to 40 and 80 groups on the 8108fl and 8116fl, respectively. A LAG can be configured with up to eight member links and can span modules. Links must be of the same speed and full-duplex operation, but can be of different media types (copper and fiber). Traffic can be load-balanced over the LAG based on:

- L2 source/destination address pairs
- L3 (default) source/destination address pairs
- L4 using the L3 SA/DA and L4 source/destination port fields

IP Routing

The 8100fl Series performs IP routing at wirespeed in ASIC hardware. In the initial release of the 8100fl Series, IP interfaces can be configured on up to the following:

- 72 routed physical ports
- 180 routed VLANs
- 40(8108fl) / 80(8116fl) routed LAGs

Dynamic Routing protocols supported are:

- RIP (Version 1, V1 compatible and Version 2) with ECMP
 - Up to 55 RIP interfaces
 - Up to four paths to another router (can be limited to one)
 - Plain Text and MD5 authentication
 - Redistribution and import/export policies to control the redistribution of routes from/to other systems
- OSPF v2 with ECMP
 - Up to 55 adjacencies
 - Up to 16 equal-cost paths and static multipath
 - Support for NSSA (Not So Stubby Areas) RFC 1587
 - Plain Text and MD5 authentication
 - External route summarization
 - Redistribution and import/export policies to control the redistribution of routes from/to other systems

Spanning Tree

The 8100fl Series provides support for 802.1s Multiple Spanning Tree Protocol (MSTP) with backward compatibility to RSTP (802.1w) and STP (802.1D).

Virtual Router Redundancy Protocol (VRRP)

The 8100fl Series supports Virtual Router Redundancy Protocol (VRRP) as defined in RFC 2338 and is compatible with other industry-standard implementations of VRRP. VRRP is one form of high availability in a Layer 3 environment by having two routing switches back each other up when presenting themselves to the network as a default gateway. In the event of a connection failure with one of the routing switches, the other routing switch transparently takes over the routing function.

Version 2 software supports up to 180 routable VLANs per interface (port or multiport LAG), with up to 15 VRRP instances per VLAN for up to 2700 VRRP IP addresses per connection to the downstream distribution switches.

VRRP features include:

- Disablement of pre-empt mode to prevent VRRP router flapping
- Configurable advertisement interval (default = 1 sec)
- Default detection time: 3.6 seconds

VRRP is not compatible with XRRP – XL Router Redundancy Protocol used on the ProCurve 5300xl Series switches. Backup routers must be configured to backup routers running compatible router redundancy protocols – 5300xl to 5300xl, and 8100fl to 8100fl or another router also running compatible RFC 2338 VRRP. RFC 2338 specified authentication types for traffic between VRRP routers, and a later RFC 3768, removed all authentication types. Switch 8100fl version 3 software (expected in 1H 2007) will be RFC 3768 compliant.

Scalability

Passive Backplane Chassis

The 8108fl and 8116fl chassis themselves have passive backplanes, with no active components between the interface modules and the fabric switching module. To provide for future increases in capacity, extra backplane links are provided to allow 20 Gbps full-duplex communication between next-generation interface and fabric modules. Backward compatibility of first-generation interface modules is assured, so an investment in today's 8100fl Series modules will continue to work with next-generation fabric modules.

Scalable Hardware Architecture

- 100K IP routes and MAC addresses, scaled up to 200K ³
- Up to 4,000 Multicast Routes ⁴, scaled up to 8,000 ⁵
- Ready for IPv6 hardware routing

The 8100fl Series hardware architecture is designed with scalability in mind. To be able to perform IPv6 routing in hardware, the ASICs that inspect the packet traffic must be designed with capabilities to inspect the 128-bit source/destination addresses (compared to 32-bit IPv4 IP addresses). Future software versions will enable this IPv6 routing capability, using existing first-generation modules.

Security

Filtering – Access Control Lists

Filtering of traffic or setting policies is performed by the hardware ASIC Packet Managers⁶ for a rich set of inspection criteria. The filtering is performed at wirespeeds and can be applied to Console or Telnet/SSH access. The 8100fl Series supports the following types of ACL filtering:

- Inbound and/or outbound per interface
- Layer 2 ACLs – filter based on source or destination MAC address
- Layer 3 ACLs – filter based on

³ Enabled in a future software release (beyond Version 2.x)

⁴ Feature available with software Version 2.5, target availability, Summer 2006

⁵ Enabled in a future software release (beyond Version 2.x)

⁶ See Hardware Architecture on page 9

- Source or Destination IP address
- Source or Destination TCP/UDP port number
- DSCP value
- IP protocol type (e.g., ICMP, IGMP, TCP, UDP, VINES)
- Fragmented packets

Protection from DoS Attacks

Network traffic travels over separate communications channels from control traffic (BDPUs, OSPF LSA, etc.) that the management module uses to manage the system. The channels to the management module are rate limited to prevent DoS attacks from impacting the switching performance of the system.

Management Access

- Out-of-band management port, separate from the data flow through the switch
- Console line-level passwords
- Telnet/SSH passwords
- Assignment of ACLs to line-level access
- TACACS+ authorization
- RADIUS authorization
- Separate rate-limited queues for control packets to the management module to prevent DoS attacks

Out-of-band management access allows for a separate management LAN to the 8100fl Series, so normal network traffic will not impact access to the management functions of the network administrator. In the case of in-band access to the switch through one of the regular interface module ports, internal rate-limited queues segregate management traffic from normal network traffic.

Summary

The ProCurve Switch 8100fl Series Interconnect Fabric delivers a new way to build networks with intelligence and interconnection from the core to every edge. A high port density Gigabit and 10 Gigabit Ethernet core networking solution, the 8100fl Series offers exceptional performance, security and availability for interconnecting intelligent edge devices. With cost-effective scalability and utmost design and deployment flexibility, the 8100fl Series ensures customers' networking investments are protected well into the future.

Appendix

Product Ordering Numbers

Product #	Name	Notes
J8727A	ProCurve Switch 8108fl	8-slot chassis, includes management & fabric module, fan trays and one power supply
J8728A	ProCurve Switch 8116fl	16-slot chassis, includes management & fabric module, fan trays and two power supplies
Extra components to order: (Base systems include Management and Fabric module)		
J8729A	ProCurve Switch 8108fl Redundant Fabric Module	Order as a second fabric module for an 8-slot system
J8730A	ProCurve Switch 8116fl Redundant Fabric Module	Order as a second fabric module for a 16-slot system
J8731A	ProCurve Switch fl Redundant Management Module	Order as a second management module for either an 8-slot or a 16-slot system
J8732A	ProCurve Switch fl Redundant Power Supply	Order as a second or third power supply for either an 8-slot or a 16-slot system (auto switching 100-240VAC)
J8734A	ProCurve Switch fl 10-port 100/1000-T Module	Single-slot, 10-port RJ45 connector for use with Cat5E or better, up to 100M reach
J8735A	ProCurve Switch fl 10-port mGBIC Module	Single-slot, 10-port mGBIC (SFP) using SX, LX, LH B-version ProCurve mGBIC modules
J8736A	ProCurve Switch fl 10GbE X2 Media Flexible Module	Single-slot, single-port X2 form-factor 10Gbps transceiver module Order CX-4, SR, LR or ER transceivers separately
J8989-61101	ProCurve fl Fan Tray	Spare part item; base systems include full complement of two fan trays Order from an authorized HP Spare Parts Reseller ⁷

Supported Accessories

Mini-GBIC (SFP) Transceivers – Gigabit Connectivity

The Switch fl 10-Port Mini-GBIC (SFP) Module (J8735A) supports the use of the following Gigabit Mini-GBIC accessories (B-version only):

- ProCurve Gigabit-SX-LC Mini-GBIC (J4858B), for use with multimode fiber cable of 62.5/125 µm or 50/125 µm (core/cladding) diameter, LC connector, graded-index, low metal content, for distances up to 550 m depending on Modal Bandwidth characteristics.
- ProCurve Gigabit-LX-LC Mini-GBIC (J4859B), for use with either multimode or singlemode fiber. Multimode fiber cable of 62.5/125 µm or 50/125 µm (core/cladding) diameter, LC connector, graded-index, low metal content for distances up to 550 m, depending on Modal Bandwidth characteristics. Singlemode fiber cable, 1310 nm wavelength, low metal content, for distances up to 10 km.
- ProCurve Gigabit-LH-LC Mini-GBIC (J4860B), for use with singlemode fiber cable, 1310 nm wavelength, LC connector, 1310 nm, low metal content, for distances up to 70 km.
- ProCurve Gigabit 1000Base-T Mini-GBIC (J8177B), for use with Category-5E or better, RJ45 connector, 100-ohm differential 4-pair unshielded twisted pair (UTP) or shielded twisted pair (STP) balanced, for distances up to 100 meters.

⁷ In the USA, go to www.hp.com/buy/parts. Search by part number: J8989-61101

Earlier versions of ProCurve mini-GBICs (J4858A, J4859A and J4860A) and non-genuine ProCurve Gigabit fiber SFPs are not supported for use in the Switch 8100fl Series. Unsupported optics will fail self-test and an error log message and SNMP trap will be generated.

X2 10Gbps Transceivers – 10 Gigabit Connectivity

The Switch fl 10GbE X2 Media Flexible Module (J8736A) supports the use of the following X2 transceivers:

- ProCurve Switch 10-GbE X2-CX4 Transceiver (J8440A), using CX-4 certified copper cable (for 0.5 meter to 15 meter reach)
- ProCurve 10-GbE CX4 Media Converter (J8439A), when connected to a CX4 port (J8440A), provides up to 300 meter reach using 12-strand MTP connector ribbon cable (requires two J8439A, one at each end connected to a CX4 port)
- ProCurve 10-GbE X2-SC SR Optic (J8436A), using multimode fiber for up to 300 meter reach
- ProCurve 10-GbE X2-SC LR Optic (J8437A), using singlemode fiber for up to 10 km reach
- ProCurve 10-GbE X2-SC ER Optic (J8438A), using singlemode fiber for up to 30 km/40 km reach⁸

The ProCurve Adaptive EDGE Architecture

The ProCurve Adaptive EDGE Architecture uses ProCurve Intelligent Edge Switch products to enable network services at the edge (“control to the edge”). The Switch 8100fl Series Interconnect Fabric products provide an integral role in this new type of network design. Because it is a standards-based, open-design strategy, it can accommodate legacy/third-party networking products to allow easy migration and flexibility. The Adaptive EDGE Architecture also focuses on enabling dynamic and automatic control and configuration (“command from the center”). The services provided at the edge can include:

- Security
- Mobility
- Convergence
- Resiliency
- Other services detailed in the Adaptive EDGE Architecture Toolbox

The Adaptive EDGE Architecture is not dependent upon:

- A specific ProCurve product
- Layer 3 at the edge
- RADIUS
- Identity Driven Management (IDM)

The Adaptive EDGE Architecture focuses on flexibility and an Adaptive EDGE Network is tailored to the customer’s needs. While meeting current requirements, it should be designed with the capacity for growth to take advantage of future applications and technologies. Implemented with all possible features, an Adaptive EDGE Network is a transparent utility network like today’s power and telephone networks, enabling users to access services and resources conveniently and transparently.

An Adaptive EDGE Network can include any appropriate ProCurve product, along with legacy products or products from other vendors, in a combination that best fulfills the customer’s needs. In addition to using a flexible range of products, an Adaptive EDGE Network can be implemented in a variety of topologies. For instance, the network can incorporate Layer 3 routing at the edge, or it can be based upon a Layer 3 core and Layer 2 edge switches.

⁸ ProCurve 10-GbE X2-SC ER Optic (J8438A) available Spring 2006. Links longer than 30 km for the same link power budget are considered engineered links. Attenuation for such links needs to be less than the minimum specified for B1.1 or B1.3 singlemode fiber.

It can deploy RADIUS and 802.1X, or use the customer's current security technologies and strategies. In short, an Adaptive EDGE Network is a flexible, standards-based network that uses intelligent edge devices to address business and technology needs.

To learn more about the ProCurve Adaptive EDGE Architecture, go to www.procurve.com, and click on "Reference Library," then "White Papers."

For more information

To learn more about ProCurve networking solutions, contact your local ProCurve sales representative or visit the company's web site at www.procurve.com.

For a list of ProCurve Elite Partners that can provide ProCurve solutions, go to www.procurve.com and click "How to buy ProCurve"

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.

4AA0-2624ENW Rev 2, 04/2007