

Configuring the 8100fl
Interconnect Fabric Switch
Rev. 6.11

HP Training
Student guide

© Copyright 2006 Hewlett-Packard Development Company, L.P.

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

This is an HP copyrighted work that may not be reproduced without the written permission of HP. You may not use these materials to deliver training to any person outside of your organization without the written permission of HP.

Configuring the 8100fl Interconnect Fabric Switch

June 2006

Overview

Introduction	Overview-1
Course objectives	Overview-1
Prerequisites	Overview-1
Course module overviews	Overview-2
Course agenda	Overview-3
Additional Information	Overview-4

Introducing the ProCurve Switch 8100fl

Objectives	1-1
ProCurve Switch 8100fl series	1-2
ProCurve WAN devices and interconnect switches	1-4
ProCurve LAN edge devices	1-5
ProCurve software and WLAN devices	1-6
The interconnect fabric in the Adaptive EDGE Network	1-7
Comparing the interconnect fabric with the traditional core switch	1-8
8116fl chassis overview	1-9
8100fl chassis slot locations	1-10
Switch fl redundant management module	1-11
Switch fl redundant fabric module	1-13
Switch fl interface modules	1-14
Features and benefits of 8100fl interface module architecture	1-15
Module 1 summary	1-17

Managing and Monitoring the ProCurve Switch 8100fl

Objectives	2-1
Initial system configuration	2-2
8100fl packaging overview	2-3
8108fl master carton	2-4
8100fl carton components	2-5
8100fl console ports	2-6
Navigate CLI levels	2-7
Save and view configuration changes	2-9
Copy configuration to TFTP server	2-10
Enable remote management	2-11
Enable remote management: Enable Telnet and virtual terminal	2-12

Enable remote management: Enable out-of-band management.....	2-13
Enable password protection	2-14
8100fl default port state	2-16
Configure remote in-band access.....	2-17
Configure time services.....	2-18
View 8100fl logs.....	2-19
Configure syslog	2-20
System monitoring and software update	2-21
Chassis: Fan Trays	2-22
Chassis: Power supplies	2-23
Monitor current temperature and thresholds	2-24
Monitor status of fans and power supplies.....	2-25
Display current software version	2-26
Obtain technical support information.....	2-27
Redundant management module operation and monitoring	2-28
Normal redundant management module behavior	2-30
Management module software banks.....	2-31
Rebooting management modules	2-32
Normal redundant fabric module behavior	2-33
Software update procedure.....	2-34
Transfer software to flash memory and verify	2-36
Install new software image	2-37
Verify image installation.....	2-38
Boot system with new image	2-39
Module 2 summary	2-40
Module 2 lab overview.....	2-43

Basic Switch 8100fl Provisioning

Objectives	3-1
Link Aggregation on the 8100fl.....	3-2
Configure Link Aggregation Groups (LAGs).....	3-4
View and modify LAG attributes.....	3-5
View and modify LAG port member status.....	3-6
Default VLAN membership for ports and LAGs.....	3-7
IP interface definition.....	3-8
8100fl VLAN support	3-10
Access mode example	3-12
VLAN creation and access port assignment	3-14
Display VLAN membership	3-15
Define IP interface for VLAN.....	3-16
View VLAN access mode ports and interface configuration.....	3-17
VLANs and LAGs: Comparison with other ProCurve switches	3-19
Example of trunk switchport mode	3-20
VLAN creation and trunk mode assignment for LAG.....	3-21

VLAN creation and trunk mode assignment for port.....	3-23
Display VLAN membership for trunk mode ports	3-24
Display VLAN membership for trunk mode ports	3-24
“Native” VLAN on the 8100fl trunk.....	3-25
Define VLAN interfaces and display status.....	3-26
Display IP interface and route tables	3-27
Module 3 summary	3-28
Module 3 lab overview.....	3-31

Provisioning Network Redundancy on the Switch 8100fl

Objectives.....	4-1
8100fl support for Spanning Tree and VRRP	4-2
MSTP/VRRP redundancy solution	4-3
Benefits of combining MSTP and VRRP	4-4
MSTP topology example	4-6
MST instance 1: Active path affects VRRP role	4-8
MST instance 2: Active path affects VRRP role	4-9
Common Spanning Tree and IST instance	4-10
Edge switch capabilities drive MSTP configuration.....	4-12
Configure and monitor MSTP.....	4-13
Spanning Tree configuration overview.....	4-14
Define VLANs and configure ports	4-16
IST configuration example.....	4-17
View Spanning Tree statistics on CIST Root	4-19
View Spanning Tree statistics on CIST Backup Root	4-20
Define parameters common to switches in MST region.....	4-21
Configure Bridge Priority for each instance	4-23
Parameters at port or LAG interface configuration context.....	4-24
View Spanning Tree statistics: Root of MST instance 1	4-25
View Spanning Tree statistics: Root of MST instance 2	4-26
Troubleshoot MST configuration.....	4-27
Configure and monitor VRRP.....	4-28
Automatic default gateway failover with VRRP	4-29
Plan for VRRP load sharing.....	4-30
8100fl VRRP configuration overview	4-32
VRRP configuration example: Configure VRIDs on 8100fl_A	4-33
VRRP configuration example: Configure VRIDs on 8100fl_B	4-34
View VRRP status	4-35
Module 4 summary	4-36
Module 4 lab overview.....	4-39

Configuring IP Routing on the Switch 8100fl

Objectives	5-1
IP routing on the 8100fl	5-2
Define port-based IP interfaces	5-4
Impact of port-based interfaces on Default VLAN	5-6
View IP interfaces and route table	5-7
Steps for RIP configuration	5-8
The <i>network</i> statement	5-9
Network numbering strategies	5-10
Compare RIP configuration procedures	5-12
3400cl IP route table	5-14
8100fl IP route table	5-15
Add a RIP router	5-16
Static route alternative	5-17
Static entries in IP route table	5-18
Choose the best path	5-19
8100fl OSPF support	5-21
OSPF configuration steps	5-23
Configure single-area OSPF: Switch 8100fl_1A	5-25
Configure single-area OSPF: Switch 8100fl_1B	5-26
View OSPF neighbor table	5-27
View OSPF interface state	5-28
Define OSPF interface as “passive”	5-29
OSPF interface parameters	5-30
OSPF in IP route table	5-32
Display link-state database	5-33
Define Area Border Router	5-34
Impact of ABR address summarization	5-35
Autonomous System Boundary Router	5-36
Define additional OSPF area types	5-37
Module 5 summary	5-39
Module 5 lab overview	5-43

Learning Check Answers

Appendices

- Appendix A: 8100fl Command Reference
- Appendix B: 8100fl Reviewer’s Guide

Introduction

Configuring the 8100fl Interconnect Fabric Switch v6.11 demonstrates and describes the deployment of the ProCurve Switch 8100fl as an interconnect fabric to support intelligent edge devices. The course emphasizes hands-on activities in the configuration of the 8100fl to support Adaptive Edge Networks.

Course objectives

After completing this course, you will be able to:

- Describe the role of the interconnect fabric in the Adaptive EDGE network
- Explain why the ProCurve Switch 8100fl is an appropriate interconnect fabric for enterprise networks
- Describe the functions and features of the switch fl modules
- Configure device security on the 8100fl switch
- Maintain and monitor the performance of the 8100fl switch
- Configure MSTP on the 8100fl switch
- Configure VRRP on the 8100fl switch
- Define VLANs and associated ports
- Configure VLAN interfaces and physical interfaces on the 8100fl
- Configure link aggregation on the 8100fl
- Configure RIP and OSPF on the 8100fl

Prerequisites

Adaptive EDGE Fundamentals

IP Routing Foundations

Routing Switch Essentials

Course module overviews

Module 1, “Introduction,” describes the 8100fl, describes its role as an interconnect fabric switch, and places it in the context of other ProCurve products.

Module 2, “Managing and Monitoring the ProCurve Switch 8100fl,” describes basic procedures for configuring the 8100fl for deployment in an enterprise environment.

Module 3, “Basic Switch 8100fl Provisioning,” introduces procedures for defining VLANs, IP interfaces, and aggregated links on the 8100fl.

Module 4, “Provisioning Network Redundancy on the Switch 8100fl,” shows how to configure a network redundancy solution that combines MSTP and VRRP.

Module 5, “Configuring IP Routing on the Switch 8100fl,” describes procedures for configuring basic IP routing features on the 8100fl, including static routes, RIP, and OSPF.

Course agenda

Day 1

Module 1: Introducing the ProCurve Switch 8100fl

Module 2: Managing and Monitoring the ProCurve Switch 8100fl

Lab 2.1: Initial Configuration and Monitoring

Lab 2.2: Updating Software on Redundant Management Modules

Module 3: Basic Switch 8100fl Provisioning

Lab 3: Configuring VLANs and Link Aggregation

Module 4: Provisioning Redundancy

Lab 4: Configuring Redundant Links and Default Gateways

Module 5: Configuring IP Routing

Lab 5: Configuring IP Routing

Additional Information

Additional information



- The HP Certified Professional (HPCP) program is a world-class certification program benchmarked around the world to ensure validation of the technical and sales competencies and expertise needed to plan, deploy, support and service HP technology and solutions
- ProCurve participates in the Sales and Integration Tracks within HPCP
- This course prepares you for a required core examination for the Accredited Systems Engineer (ASE) certification
- The exam number for this course is HP2-013
- For more information on HPCP, go to www.hp.com/certification
- For more information on HP ProCurve Training and Certification, go to <http://www.hp.com/rnd/training/>

Configuring the 8100fl is part of a series of courses on networking technologies and ProCurve products. For more information, visit the ProCurve Web site, www.procurve.com.

Introducing the ProCurve Switch 8100fl

Module 1


Objectives

This module introduces the ProCurve Routing Switch 8100fl series. After completing this module, you will be able to:

- List the features and benefits of the ProCurve Switch 8100fl series
- Explain how the 8100fl provides an interconnect fabric for the Adaptive EDGE Network
- Compare and contrast the switches in the 8100fl series
- Describe the functions and architecture of the 8100fl series modules
- Describe the features and benefits of the 8100fl switching architecture

ProCurve Switch 8100fl series

ProCurve Switch 8100fl series



Highly available, wire-speed interconnect fabric for intelligent edge devices using

- 10 GbE
- 100/1000T
- Mini-GBIC

High degree of flexibility and redundancy, supports:

- Routed or switched links
- Up to 1,024 IEEE 802.1Q compliant VLANs
- Link aggregation
- Multiple Spanning Tree
- VRRP
- IGMP

Dynamic routing protocols with ECMP support for OSPF

Rev 6.11

Student Guide: 1–2

3

As ProCurve's first interconnect fabric switch, the ProCurve Switch 8100fl series provides enterprises with a wire-speed, high availability core designed specifically to interconnect intelligent edge devices.

The 8100fl series is available in two models:

- The ProCurve Switch 8108fl provides eight interface slots, enabling it to support up to 80 gigabit ports or eight 10-GbE ports.
- The ProCurve Switch 8116fl provides sixteen interface slots, enabling it to support up to 160 gigabit ports or 16 10-GbE ports.

Both models support redundant management and switch fabric modules. All available modules will be described later in this module.

Flexible routing and switching

In keeping with its role as an interconnect fabric, the 8100fl is a true routing switch, handling Layer 2 and Layer 3 forwarding at wire speed. The 8100fl supports virtual (VLAN) IP interfaces and physical (port-based) interfaces, even allowing administrators to mix both types of interfaces on a single switch.

Supported routing protocols include RIP and OSPF. Equal Cost Multipath (ECMP) routing is supported for OSPF.

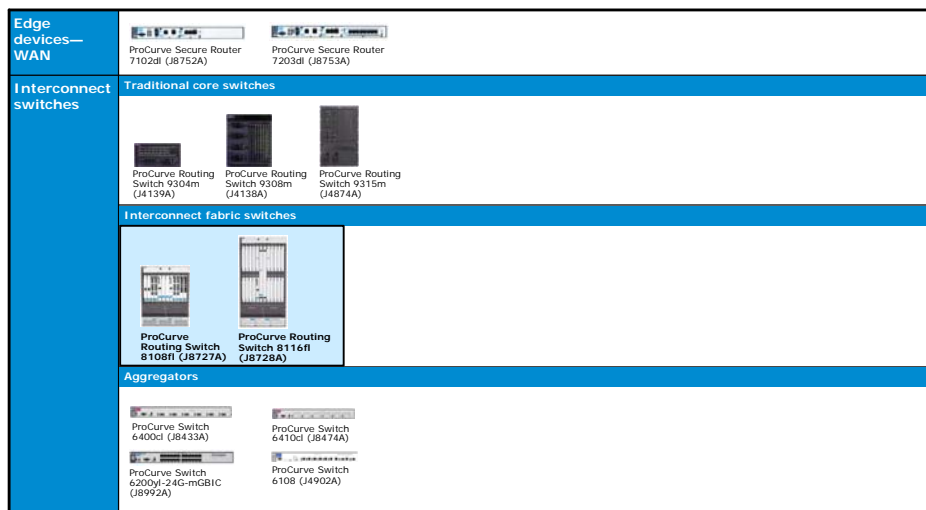
High capacity and availability

As well as offering wire-speed forwarding for gigabit and 10-GbE links, the 8100fl offers extremely flexible support for link aggregation.

To provide high levels of network availability, the 8100fl supports two versions of Spanning Tree and the Virtual Router Redundancy Protocol (VRRP). The 8100fl supports both Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP), making it completely compatible with other ProCurve products. To support complex routed environments, each of the 8100fl's IP interfaces can support up to 15 VRRP instances.

ProCurve WAN devices and interconnect switches

ProCurve WAN devices and interconnect switches



Rev 6.11

Student Guide: 1–4

4

The ProCurve product line includes switches, routers, and other devices to fill almost any role in the contemporary network. The next three pages will introduce the entire ProCurve line, including WAN routers, WLAN devices, and software, as well as a complete line of LAN switches.

Because it is specifically designed to be an interconnect fabric, the ProCurve Switch 8100fl series occupies a unique position in the ProCurve product line. As shown, the other interconnect products are the 9300m traditional core routers and aggregators, including the ProCurve Switch 6200yl-24G-mGBIC.

ProCurve LAN edge devices

ProCurve LAN edge switches



Edge devices—LAN

Intelligent Edge switches

ProCurve Switch 5406zl (J8697A)

ProCurve Switch 5406zl-48G (J8699A)

ProCurve Switch 5412zl (J8698A)

ProCurve Switch 5412zl-96G (J8700A)

ProCurve Switch 3500yl-24G-PWR* (J8692A)

ProCurve Switch 3500yl-48G-PWR* (J8693A)

ProCurve Switch 5304xl (J4850A)

ProCurve Switch 5304xl-32G (J8166A)

ProCurve Switch 5348xl (J4849B)

ProCurve Switch 5308xl (J4819A)

ProCurve Switch 5308xl-48G (J8167A)

ProCurve Switch 5372xl (J4848B)

ProCurve Switch 3400cl-24G (J4905A)

ProCurve Switch 3400cl-48G (J4906A)

*Power over Ethernet

Edge switches—managed

ProCurve Switch 4202vl-48G (J8771A)

ProCurve Switch 4202vl-72 (J8772A)

ProCurve Switch 4204vl (J8770A)

ProCurve Switch 4208vl (J8773A)

ProCurve Switch 4208vl-64G (J8774A)

ProCurve Switch 4208vl-96 (J8775A)

ProCurve Switch 4104gl (J4887A)

ProCurve Switch 4140gl (J8151A)

ProCurve Switch 4148gl (J4888A)

ProCurve Switch 4108gl (J4865A)

ProCurve Switch 4160gl (J8152A)

ProCurve Switch 4108gl Bundle (J4861A)

ProCurve Switch 2824 (J4903A)

ProCurve Switch 2848 (J4904A)

ProCurve Switch 2810-24G (J9021A)

ProCurve Switch 2810-48G (J9022A)

ProCurve Switch 2626 (J4900B)

ProCurve Switch 2626-PWR* (J8164A)

ProCurve Switch 2650-PWR* (J8165A)

ProCurve Switch 2600-S-PWR* (J8762A)

ProCurve Switch 2510-24 (J9019A)

ProCurve Switch 2512 (J4812A)

ProCurve Switch 2524 (J4813A)

ProCurve Switch 2650 (J4899B)

*Power over Ethernet

Web Managed

ProCurve Switch 1800-8G (J9029A)

ProCurve Switch 1800-24G (J9028A)

Edge switches—unmanaged

ProCurve Switch 2708 (J4896A)

ProCurve Switch 2724 (J4897A)

ProCurve Switch 2312 (J4817A)

ProCurve Switch 2324 (J4818A)

ProCurve Switch 2124 (J4868A)

ProCurve Switch 408 (J4097B)

The ProCurve LAN offerings range from unmanaged edge switches that offer basic workgroup connectivity to intelligent edge switches with advanced features that enable the Adaptive EDGE Architecture.

ProCurve software and WLAN devices

ProCurve software and WLAN devices



Wireless Edge Services	ProCurve Wireless Edge Services xl Module (J9001A) ProCurve Redundant Wireless Services xl Module (J9003A) ProCurve Radio Port 210 (J9004A) ProCurve Radio Port 220 (J9005A) ProCurve Radio Port 230 (J9006A)
Secure Access	ProCurve Access Control Server 745wl (J9038A) ProCurve Switch xl Access Controller Module (J8162A)
Network Management Software	ProCurve Manager 2.1 ProCurve Manager Plus 2.1 (J8778A, J9009A, J8991A, J8779A) ProCurve Identity Driven Manager 2.0 (J9012A, J9013A, J9014A) ProCurve Mobility Manager 1.0 (J8990A)

Rev 6.11

Student Guide: 1–6

6

The ProCurve product line includes an extensive array of WLAN devices designed to provide secure, reliable wireless access for enterprise users.

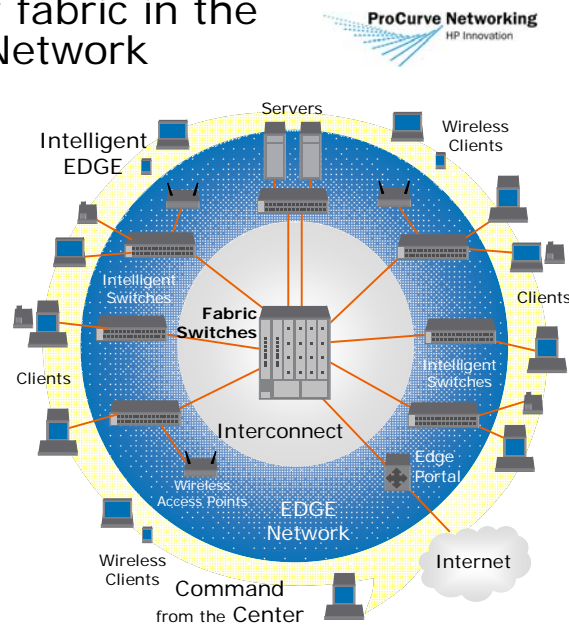
To provide “Command from the Center, Control to the Edge,” ProCurve offers ProCurve Manager 2.1, ProCurve Manager Plus 2.1, ProCurve Identity Driven Manager 2.0, and ProCurve Mobility Manager 1.0. These cutting-edge products provide network administrators with the power and flexibility they need to protect and maximize the contemporary network.

The interconnect fabric in the Adaptive EDGE Network

The interconnect fabric in the Adaptive EDGE Network

An Adaptive EDGE Network is a flexible, standards-based network that uses intelligent edge devices to address business and technology needs

The interconnect fabric provides a highly reliable, high-speed interconnect for intelligent edge switches



Rev 6.11

Student Guide: 1-7

7

The ProCurve Adaptive EDGE Networking Architecture is designed specifically to meet the changing needs of today's enterprise networks. The growth of the Internet and the availability of an ever-growing array of applications and services have presented the enterprise network with new requirements for security, mobility, and convergence. The Adaptive EDGE Architecture enables enterprises to fulfill these requirements by deploying intelligent edge devices in networks designed to meet business needs.

The ProCurve Switch 8100fl series plays a crucial role in the Adaptive EDGE Architecture by providing a highly reliable, high-speed interconnect for the intelligent edge switches. The rest of this course will describe scenarios and configurations that enable the 8100fl to exhibit optimal performance in this role.

Comparing the interconnect fabric with the traditional core switch

Comparing the interconnect fabric with the traditional core switch



	Traditional Core Switch	Interconnect Fabric Switch
Principle roles in network design	<ul style="list-style-type: none"> • Central point for network interconnect, security, traffic management • Performs most routing • Performs complex traffic filtering and control 	<ul style="list-style-type: none"> • Highly reliable, high-speed interconnect for intelligent EDGE switches • Honors all decisions made at the edge
Key features and functions	<ul style="list-style-type: none"> • Robust routing with multiple protocol support 	<ul style="list-style-type: none"> • Bandwidth and high availability • Simple and easy to deploy and manage
Limitations	<ul style="list-style-type: none"> • Scalability limited by bandwidth of processor sub-systems • Complex and difficult to configure 	<ul style="list-style-type: none"> • Will not have feature set of a core routing switch because functionality is expected at network edge
Key benefits	<ul style="list-style-type: none"> • Familiar and pervasive • Available from many vendors 	<ul style="list-style-type: none"> • Lower cost and complexity • Highly scalable performance

Rev 6.11


Student Guide: 1–8

8

Although the interconnect fabric and the traditional core switch occupy similar physical locations in a network topology, they provide dramatically different feature sets. They can often both be deployed effectively in a single enterprise, especially in situations where organizations are migrating toward the Adaptive EDGE Architecture.

In keeping with its role as interconnect fabric, the 8100fl has a substantially different feature set than the traditional core router. The 8100fl provides high levels of bandwidth and availability, but also provides a relatively simple feature set. For instance, the 8100fl does not support IPX, AppleTalk, or the large number of router and ACLs typically supported by core routers. In this way, the feature set of the 8100fl assumes many routing and security duties will be performed by edge devices.

8116fl chassis overview



The architecture and design of the 8100fl are significantly different from other ProCurve switches. Some unique features of the 8100fl chassis are:



- RJ-45 ports for serial console access. Located on the top brow of the chassis, the ports are wired for RS-232. Access requires a special adapter included in the 8100fl basic kit. Each port is hard-wired to one of the two management module slots.
- Redundant, replaceable fabric modules
- Removable fan trays, each of which supports three cooling fans

8100fl chassis slot locations


8100fl chassis slot locations




Center slots are **Fabric Module** slots
 Surrounded by **Management Module** slots
 Interface Module slots **numbered**
 left to right

 = management module
 = fabric module

Rev 6.11



Student Guide: 1–10



10

The 8100fl features dedicated slots for three different types of modules—fabric modules, management modules, and interface modules.

As a high availability interconnect fabric switch, the 8100fl locates its switching fabric on redundant modules, not on the chassis backplane. Note that the fabric module slots on the 8116fl extend the entire height of the switch's 20 rack unit chassis, making it twice the height of the fabric module for the 8108fl.

The 8100fl also features dedicated slots for management modules. All of the slots in the chassis are connected to the management module through an internal 100MB connection.

The slots for interface modules are numbered from left to right and identified by labels at the bottom of each slot.

The 8116fl chassis includes two “expansion” slots located just above the management modules. These slots are not connected to the switch fabric and cannot support any type of fl switch module. As shipped to customers, the 8116fl includes blank covers for these slots.

Switch fl redundant management module

Switch fl redundant management module

10/100 Base-T management port

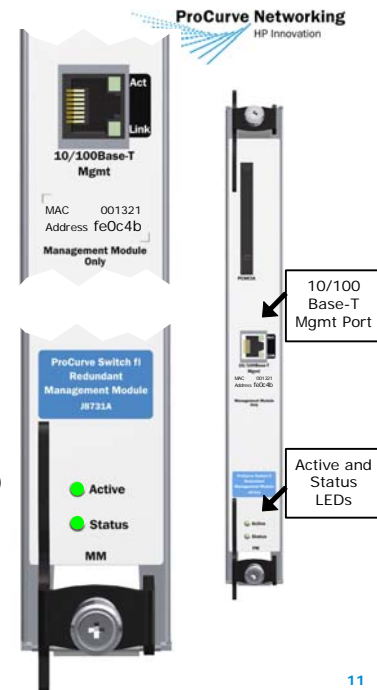
- MDI host port
 - Use crossover cable for direct connection to PC
- Traffic is isolated from user networks, not routed or switched
- May be used for out-of-band management after an IP address has been assigned and Telnet has been enabled

Active LED

- Is illuminated on the primary (active) management card
- Is not illuminated on the secondary (standby) management card

Status LED

- Flashes during software boot process
- Remains solid when bootup is complete



Rev 6.11

Student Guide: 1-11

11

The 8100fl management module includes an Ethernet port that can be enabled and configured to support an “out-of-band” IP interface. This port is most useful when you need to transfer files such as updated software from a computer to the switch’s flash memory. Because the computer and the switch’s management port are host ports, operating in MDI-X mode, you must use a Cat 5 crossover cable to connect them.

The management port is considered “out of band” because traffic received through the port is used only by its IP host process. The traffic is never forwarded through the switch fabric to interface modules. Conversely, traffic received through the interface modules is never forwarded through the management port.

In an environment that supports multiple 8100fl switches, you can connect each 8100fl management port to a Layer 2 switch, creating an isolated, out-of-band management network. While the management network can include one or more computers, it must not include switches or routers that are connected to the data network.

The management port is not immediately available when the switch is at default settings. Procedures to activate the out-of-band management port will be described in Module 2.

Just beneath the management port, the MAC address of the management module is displayed in brackets. Each management module has a unique MAC address. This is not the MAC address that determines the addresses to be assigned to individual ports. Instead, the MAC addresses that the switch assigns to each port are inherited from a pool of 512 addresses assigned to the chassis.

The bottom of each management module contains two LEDs. The “Active” LED indicates whether the module is functioning as the primary module. In a system with redundant management modules, the Active LED will be lit on only one module. The “Status” LED blinks during system boot or reload. The LED will be solid green when the boot procedure is completed.

Finally, the top of the module includes a PCMCIA slot that is not supported in the current software version.

Switch fl redundant fabric module

Switch fl redundant fabric module

Redundant fabric modules provide for nearly hitless failover for switched traffic


- Both switch fabrics receive a copy of all control information
 - Enables failover of less than 0.2 seconds

Two LEDs

- Status LED indicates booting, running, and fault conditions
- Active LED indicates whether module is active or non-active

Different fabric modules for 8108fl and 8116fl

- 8108fl fabric module (J8729A) is half the size of 8116fl module (J8730A)



Rev 6.11

Student Guide: 1–13

12

On the 8100fl, the fabric module processes all traffic for all interface modules.

Although the system functions properly with only one module, the installation of a redundant module provides for a very high level of availability, enabling a nearly hitless failover time of less than 0.2 seconds. On the 8100lf, all control data is sent to both fabric modules, enabling the secondary module to begin forwarding packets correctly immediately after switchover. Similarly, the interface cards detect the failure of the primary module and immediately begin sending packets to the secondary.

As noted earlier, the 8108fl and 8116fl use different fabric modules:

- 8108fl Redundant Switch Fabric Module (J8729A)
- 8116fl Redundant Switch Fabric Module (J8730A)

The module for the 8108fl is half the height of the 8116fl module, but their functionality is identical.

Switch fl interface modules

Switch fl interface modules



Interface modules compatible with 8108fl and 8116fl

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

10 ports providing
100/1000-T connectivity



Rev 6.11

**Switch fl 10-port
Mini-GBIC Module
(J8735A)**

Supports "B version"
Mini-GBICs:
Gigabit 1000Base-T
Mini-GBIC (J8177B)
Gigabit-SX-LC
Mini-GBIC (J4858B)
Gigabit-LX-LC
Mini-GBIC (J4859B)
Gigabit-LH-LC
Mini-GBIC (J4860B)



Student Guide: 1-14

**Switch fl One-port
10 GbE X2 Module
(J8736A)**

Supports the same
transceivers as
3500yl/5400zl:
10 GbE X2-SC LR
Optic (J8437A)
10-GbE X2-CX4
Transceiver (J8440A)
10-GbE X2-SC SR
Optic (J8436A)
10-GbE X2-SC ER
Optic (J8438A)
10-GbE CX4 Media
Converter (J8439A)



13

Both 8100fl modules support three interface modules:

- Switch fl 10-port 100/1000-T Module (J8734A)
- Switch fl 10-port Mini-GBIC Module (J8735A)
- Switch fl One-port 10-GbE X2 Module (J8736A)

On both the 8108fl and the 8116fl, the modules can be used in any combination to populate all interface module slots.

Features and benefits of 8100fl interface module architecture

Features and benefits of 8100fl interface module architecture



Architectural features that contribute to performance include:

- Separate ingress and egress packet processors on each interface module
 - Enable the 8100fl to handle traffic at the full line rate
- Traffic Manager on each interface module
 - Facilitates transfers to and from switch fabric
- Independent CPU, RAM, and packet buffer space on each interface module
 - CPU interfaces with management module over dedicated management bus, enabling ingress processors to perform local packet lookups without involving the management module
- Multicast replication minimizes congestion by ensuring each packet in a multicast stream is replicated as necessary with the correct tag
 - Traffic Manager stores copy of each packet
 - Egress processors on interface modules perform tagging for multiple VLANs

Rev 6.11

Student Guide: 1–15

14

The architecture of the 8100fl interface modules includes features that contribute to the switch's performance as an interconnect fabric. The architectural highlights include:

■ **Packet processors**

Every interface module features separate ingress and egress Packet Processors that handle traffic at full line rate. The ingress Packet Processor performs lookups on inbound packets and appends information that will be used by other components in the switch to forward the packet to the switch fabric. The egress Packet Processor handles packets that arrive from the chassis-wide switch fabric, preparing the packets for outbound transmission through ports on the module.

■ **Traffic Manager**

The Traffic Manager on each interface module receives packets from the ingress Packet Processor and facilitates their transfer to the chassis-wide switch fabric. The Traffic Manager also receives packets from the switch fabric and sends them to the egress Packet Processor.

■ **CPU, RAM, and Packet buffer**

To enhance performance and avoid congestion, each 8100fl interface module includes its own CPU, RAM, and packet buffer space. The CPU uses a dedicated 100 Mbps bus to interact with the CPU on the management module. This provides the module's ingress Packet Processor with the information required to perform local packet lookups without involving the Management Module.

- **Multicast replication**

This feature ensures that each multicast packet will cross the switch fabric only once, regardless of the number of ports and VLANs to which it must ultimately be forwarded. In the replication process, the Traffic Manager on each interface module stores each multicast packet in an identified location. Each port's egress processor copies the packet as many times as is required and applies the appropriate VLAN tagging to each copy before forwarding it.

Module 1 summary

Module 1 summary



In this module, you learned about the features and benefits of the ProCurve Switch 8100fl

Topics included:

- The role of the interconnect fabric in the Adaptive EDGE Network
- The features of the 8100fl that make it suitable for deployment as an interconnect fabric
- The functions and architecture of 8100fl series modules
- Advantages of the 8100fl switching architecture

Rev 6.11

Student Guide: 1–17

15

Module 1 introduced the ProCurve Switch 8100fl series, which is the first ProCurve product designed specifically to serve as an interconnect fabric for edge-oriented enterprise networks.

As well as describing the switch's role in the contemporary network, the module described its modules, chassis, and the advantages of its switching architecture.

Module 2 will describe basic configuration procedures.

Learning check

Module 1

1. What is the role of the interconnect fabric in the Adaptive EDGE Network?
.....
.....
2. Name three features of the ProCurve Switch 8100fl that make it well suited for deployment as an interconnect fabric.
 - a.
 - b.
 - c.
3. What is the difference between the two models of 8100fl?
.....
.....
4. Describe the three types of interface modules available for the 8100fl.
 - a.
 - b.
 - c.

Managing and Monitoring the ProCurve Switch 8100fl

Module 2

Objectives

After completing this module and the accompanying hands-on activity, you will be able to:

- Configure a ProCurve Switch 8100fl using the CLI
- Connect to the 8100fl series switch using the serial console port or out-of-band management port
- Enable out-of-band management on an 8100fl series switch
- Enable remote management of an 8100fl series switch
- Describe the process for saving and viewing changes to the configuration files of an 8100fl series switch
- Describe the functioning of the redundant management fabric modules on the 8100fl
- Perform 8100fl configuration, monitoring, and maintenance tasks including:
 - Configure device security
 - Configure time and syslog services
 - Monitor and manage the system environment
 - Perform software update
 - Back up system configuration

Initial system configuration

Initial system configuration



Initial system configuration

- **8100fl packaging**
- **Console port location**
- **CLI hierarchy**
- **Configuration storage architecture**
- **Device security**

System monitoring and software update

Rev 6.11

Student Guide: 2-2

3

Module 2 will begin with a discussion of the processes and tools for performing initial system configuration on a ProCurve Switch 8100fl. Specific topics will include:

- 8100fl packaging
- Console port location
- CLI hierarchy
- Configuration storage architecture
- Device security

The second part of Module 2 will describe processes and tools for system monitoring and for updating software images.

8100fl packaging overview

8100fl packaging overview



Installation of the 8100fl requires two or more people



Country-specific power cord packed just inside the lid of the master carton



Rev 6.11

Student Guide: 2-3

4

The 8108fl and the 8116fl chassis are heavier and bulkier than most other ProCurve products. For instance, the 8108fl weighs just less than 220 pounds when all slots are populated. A fully populated 8116fl weighs nearly 340 pounds. Consequently, the unpacking and installation of either model will require two or more workers.

Furthermore, customers accustomed to other ProCurve products will find that the 8100fl packaging is substantially different. For instance, the country-specific power cord is inserted into each 8100fl package just before shipping. As shown, it can be found just inside the lid of the master carton.

The next two slides will provide more detail on the 8100fl packaging.

8108fl master carton

8108fl master carton



Master carton contains basic J8727A and included components



Power Cord will go on top here

Rev 6.11

Student Guide: 2-4

5

The master carton for the 8108fl holds multiple boxes for the chassis and all included components. Because the carton is purpose-built for shipping on a pallet, it cannot typically be removed from the pallet for delivery to the switch's ultimate destination. Instead, the chassis, modules, power supplies, and fans typically are unpacked at a receiving facility before being delivered for installation.

8100fl carton components

8108fl carton components



Rev 6.11

Student Guide: 2-5

6


As shown above, the master carton for the 8108fl contains the following components:

- Switch chassis
- One power supply
- One management module
- One fabric module
- Two fan trays
- Rack mounting kit
- Documentation and console adapter

The basic 8116fl package includes all of the same components as the 8108fl plus an additional fan tray.

8100fl console ports

8100fl console ports



Rev 6.11

Student Guide: 2-6

7

The 8100fl features two console ports, located on the top brow of the switch. Each of the ports is hard-wired to one of the management modules. The console port on the left side of the chassis is hard-wired to Management Module Slot A, which is on the left side of the chassis. The console port on the right side of the chassis is hard-wired to Management Module Slot B, which is on the right side of the chassis.

The console ports of the 8100fl employ a different architecture than other ProCurve switches. Most notably, the ports use an RJ45 connector, but are wired as RS232 serial ports. To access a port, you must use an RJ45-to-DB9 adapter included with the switch.

When an 8100fl is at factory defaults, you can access the switch for configuration or monitoring by connecting to the console port associated with the active management module. As described in Module 1, the Active LED will be solid green on the active module. No password is assigned to this port at default settings. The process for assigning a password will be discussed later in this module.

During the first console-port session, you can enable remote management using Telnet through an in-band network or through the switch's dedicated out-of-band management port. The processes for this configuration will be described later in this module.

Navigate CLI levels

Navigate CLI levels



	8100fl series	3400cl/5300xl
<i>enable</i>	EXEC level (>) to Privileged EXEC level (#)	Operator level (>) to Manager level (#)
<i>exit</i>	Return to previous level; for instance from Privileged EXEC to EXEC	Move from Manager level to Operator level
<i>configure</i>	Move from Privileged EXEC level to global configuration context	Move from Manager level to global configuration context
<i>end</i>	Exit any configuration context to Privileged EXEC level Some configuration levels are nested within other levels	Exit any configuration context to Manager level All configuration levels are directly under global configuration level
<i>logout</i>	Exit EXEC or Privileged EXEC level and terminate session	Exit Manager or Operator level and terminate session

Rev 6.11

Student Guide: 2-7

8

The 8100fl supports the same navigation commands as other ProCurve switches, including the 3400cl and 5300xl.

Like the User EXEC and Operator levels on other ProCurve switches, the EXEC level on the 8100fl is the point of entry for configuration and management tasks. The EXEC level supports a limited set of commands, including a subset of the *show* commands available at the Privileged EXEC level.

By entering the *enable* command, the user moves to the Privileged EXEC level. This level allows the full set of *show* commands, as well as some one-time execution commands, such as those that reboot the switch and copy image and configuration files between flash memory and a TFTP server.

On the 8100fl, the global configuration level permits you to enable routing protocols and to configure box-wide features such as host name, SNMP support, and ACLs. To enter the configuration context, issue the *configure* command in the Privileged EXEC. For user convenience, the 8100fl will accept *configure terminal*, but it is not required.

The process for exiting 8100fl configuration levels also is similar to other ProCurve switches. You can always use the *exit* command to return to the previous level. In the Privileged EXEC, you can use the *disable* command (the reverse of *enable*) to return to the EXEC level. *Disable* and *exit* have the same effect at the Privileged EXEC level.

The CLI also supports a few shortcuts for the *exit* command that are similar to other ProCurve switches. From any configuration context, you can issue *end* to exit the configuration mode entirely and return to the Privileged EXEC. To disconnect, issue *logout* in the Privileged EXEC or EXEC levels. *Logout* is not valid in configuration contexts.

Finally, the 8100fl supports tab completion and the command shortcuts available on other ProCurve switches. Like other ProCurve switches, the CLI of the 8100fl will accept any unambiguous string of characters as a shortcut. For instance, *int e 1/1* is a shortcut for *interface ethernet 1/1*, which is the command to enter the configuration context for gigabit port 1/1.

Save and view configuration changes

Save and view configuration changes



- Like other ProCurve switches, the 8100fl stores the running configuration in volatile RAM and the startup configuration in non-volatile flash memory
- Most commands for saving and viewing configuration changes are the same as on the 3400cl/5300xl

	8100fl series	3400cl/5300xl
<i>write memory</i>	Copy running configuration to startup	Copy running configuration to startup
<i>show run</i>	View current running configuration	View current running configuration
<i>show config</i>	View current startup configuration	View current startup configuration
<i>erase startup</i>	Erase startup configuration Switch must be rebooted with separate command	Erase startup configuration Switch reboots automatically

Rev 6.11

Student Guide: 2-9

9

Most commands used to alter configuration files are the same as on other ProCurve switches. The *write memory*, *show run* and *show configuration* commands can be executed at the Privileged EXEC level or any configuration context. Note that you must explicitly reboot the 8100fl after erasing the startup configuration. The 3400cl/5300xl automatically reboots when the configuration is deleted.

Copy configuration to TFTP server

Copy configuration to TFTP server



Copy command must be executed at Privileged EXEC level

Syntax to identify server and filename is different from other ProCurve switches:

- Copy the running-config and startup-config to the TFTP server:

```
8100fl# copy running-config tftp://10.1.1.104/backup1.cfg
```

```
8100fl# copy startup-config tftp://10.1.1.104/backup2.cfg
```
- Copy configuration file from the TFTP server to the startup-config:

```
8100fl# copy tftp://10.1.1.104/backup2.cfg startup-config
```

Rev 6.11

Student Guide: 2-10

10

Like other ProCurve switches, the 8100fl supports TFTP for the transfer of configuration and image files. The syntax for using this feature on the 8100fl is slightly different than on other devices, but the basic commands are the same. For example, the *copy* command initiates a file transfer and has a basic structure—*copy <source> <destination>*—similar to the *copy* command on other ProCurve switches.

On the 8100fl, however, you must use the full URL, such as “tftp://10.1.1.104/backup2.cfg,” to identify the server and filename for transfer.

You can also use Secure Copy (SCP) to transfer files. The SCP syntax is: “scp://[username@]location/directory/filename.” You will be prompted for a password by the SCP server.

Before transferring files to or from the 8100fl, you must enable remote management, a process that will be described in the next few pages.

Enable remote management

Enable remote management



Remote management for the 8100fl is disabled at default settings

To enable remote access:

- Enable the Telnet server
- Enable support for at least one virtual terminal session
 - Maximum of 10 numbered 0 through 9
- Assign an IP address and activate at least one interface

Establish Telnet sessions through:

- Any router interface
- Connection to out-of-band Ethernet management port

Rev 6.11

Student Guide: 2-11

11

Unlike other ProCurve switches, the 8100fl will not act as a Telnet server until you explicitly enable the Telnet protocol and configure support for a virtual terminal session. On other switches, the Telnet service is enabled as soon as the first IP interface becomes available.

To support Telnet management sessions, the 8100fl must be configured with at least one IP interface. This can be one of the 8100fl's router interfaces, part of the data network, or an address assigned to the dedicated out-of-band Ethernet management port.

Enable remote management: Enable Telnet and virtual terminal

Enable remote management: Enable Telnet server and virtual terminal



Enable the switch's Telnet server:

```
8100fl(config)# ip telnet
8100fl(config-telnet)# no shutdown
```

Enable support for one virtual terminal session:

```
8100fl(config-telnet)# line vty 0
8100fl(config-line)# exit
8100fl(config)#
```

To support additional virtual terminal sessions:

```
8100fl(config)# line vty 1
8100fl(config-line)# exit
8100fl(config)#
```

Rev 6.11

Student Guide: 2-12

12

The process for enabling Telnet access to the 8100fl management interface has two steps:

1. **Enable the Telnet service**

On the 8100fl, the *telnet* configuration command is preceded by *ip* because the Telnet service is part of the 8100fl's IP host process. At default settings, the IP Telnet service is disabled. This default state must be reversed using the *no shutdown* command.

2. **Enable support for at least one virtual terminal session**

By default, no virtual terminals are enabled on the 8100fl. You can configure up to 10 separate terminals, numbered 0 through 9, which enables the 8100fl to support multiple configuration sessions simultaneously. You can enable a virtual terminal with the *line vty [x]* command shown above. You cannot create multiple lines with a single command. Each new Telnet client receives the lowest-numbered terminal available at login time.

The steps can be performed in any order.

The procedures for enabling SSH on the 8100fl are similar to the Telnet procedures. For instance, you enable SSH by issuing the *ip ssh* command. See the "Security Configuration" section of the *Management and Configuration Guide* for more information.

Enable remote management: Enable out-of-band management


Enable remote management: Enable out-of-band management



Assign IP address to Ethernet management port:

```
8100fl(config)# interface management 0
8100fl(config-interface-management)# ip address 192.168.0.201/24
8100fl(config-interface-management)# no shutdown
```

- In system with redundant management modules:
 - Interface configuration parameters apply to port on active module
 - After switchover event, move cable from inactive to active module
 - Each management module has different MAC address
- Use a crossover cable to connect a computer to active management port



Rev 6.11

Student Guide: 2-13

13

Like all 8100fl interfaces, the Ethernet port on the management module is disabled by default. To activate it, you must enter its configuration context (*interface management 0*), assign an IP address, and issue the *no shutdown* command.

In a system with redundant management modules, configuration parameters entered in the *management 0* interface apply to the port on the active management module. No parameters are entered directly into the backup management module, which automatically synchronizes its startup configuration with the active module. Management module synchronization and switchover will be described in more detail later in this module.

In the event of a management module switchover, the configuration of the management port applies to the interface on the module that becomes active.

Enable password protection

Enable password protection



The 'enable' password protects the Privileged EXEC from unauthorized access

```
8100fl(config)# enable secret procure
```

Line-level passwords protect serial console and virtual terminal entry points from unauthorized access

```
8100fl(config)# line console 0
8100fl(config-line)# password 0 hp
8100fl(config-line)# line vty 0
8100fl(config-line)# password remote
8100fl(config-line)# line vty 1
8100fl(config-line)# password remote
8100fl(config-line)# write memory
8100fl(config-line)# end
8100fl# logout
```

To authenticate to the Privileged EXEC from a Telnet session:

```
Password: remote (system will not echo the characters)
8100fl> enable
Password: procure (system will not echo the characters)
8100fl#
```

Rev 6.11

Student Guide: 2-14

14

The 8100fl supports configuration of passwords to protect the serial console, remote terminal sessions, and access to the Privileged EXEC level.

As shown above, the same syntax is used to assign line-level passwords to the serial console and to all configured virtual terminals. To assign a virtual terminal password, enter the terminal's configuration context with the command *line vty*, followed by the terminal number. It is possible to assign a different password to each virtual terminal. However, it is advisable to configure the same password for all virtual terminals to prevent confusion. When users start a Telnet or SSH session, they are assigned the available terminal with the lowest number. Administrators and users cannot control terminal assignment.

The *enable* password protects the Privileged EXEC from unauthorized access.

After the configuration steps shown on the slide are completed, a user starting a console or Telnet session will be prompted for passwords before entering the EXEC and Privileged EXEC levels.

Optional password encryption

By default, the 8100fl encrypts passwords as they are shown in the running configuration and startup configuration files. In the example, the Privileged EXEC password will be encrypted because the *enable secret procure* command does not specify an encryption option. To force a password to be stored in clear text, use the 0 encryption option, as shown in the password applied to the serial port with *password 0 hp*.

A third encryption option enables you to indicate that the string being configured as a password has *already* been encrypted. This option is invoked by supplying the value of “5” for the encryption parameter, as in *password 5 \$1\$ZyK.\$8NHx2DJ*. Exercise caution when using this option, as it can make it difficult for users to authenticate. For instance, if *password 5 hp* were entered into the configuration, the serial connection password would not be “hp,” but the value output by encryption of an unknown string. The password would be shown in the configuration files as “hp,” which the system would consider to be the encrypted value of the true password.

8100fl encryption options

Option	Command value	Command example
Stored as encrypted string in configuration files	System default	<i>password hp</i>
Stored as clear text in configuration files	0	<i>password 0 hp</i>
String entered as password has already been encrypted	5	<i>password 5 \$1\$ZyK.\$8NHx2DJBsiGQyhTBmUakz1</i>

The 8100fl also supports aaa authentication using local usernames or RADIUS/TACACS+ servers. See “Security Configuration” in the *Management and Configuration Guide* for more information.

8100fl default port state

8100fl default port state



Default 8100fl port settings:

- All ports disabled or “shut down”
- Spanning Tree is disabled on all ports
- All ports are members of VLAN 1, the default VLAN

```
8100fl# show vlan
VLAN      Name      Status  Ports      Type
1          Default  active  Gig1/1
           Gig1/2
           Gig1/3
           Gig1/4
           Gig1/5
           Gig1/6
           Gig1/7
           Gig1/8
           Gig1/9
           Gig1/10
```

To access port interface configuration contexts:

```
8100fl(config)# interface gigabitethernet 1/1
8100fl(config)# interface tengigabitethernet 3/1
```

Rev 6.11

Student Guide: 2-16

15

By default, all ports on 8100fl interface cards are disabled or “shut down” and will remain in that state unless explicitly enabled. This configuration, although unique among ProCurve switches, is appropriate for an interconnect fabric because it allows technicians to connect cables to ports before the links are logically enabled. This can be especially useful when configuring redundant links because it eliminates concerns about creating loops while the configuration is incomplete.

By default, Spanning Tree is disabled on the 8100fl. Spanning Tree is enabled per port instead of per VLAN or system-wide, as with other ProCurve switches. Spanning Tree on the 8100fl will be discussed in detail in Module 4.

As with other routing switches, all ports on the 8100fl are untagged members of VLAN 1 at default settings.

To learn the VLAN configuration of an 8100fl, issue *show vlan*, which displays information about every VLAN and all of its assigned ports.

Port interface configuration contexts

To enter configuration parameters for ports on the 8100fl, you must enter the port configuration context, using the commands shown below.

Type of port	Full command	Shortcut
100/1000	interface gigabitethernet 1/1	int g 1/1
100/1000	interface ethernet 1/1	int e 1/1
10-GbE	interface tengigabitethernet 1/1	int t 1/1

Configure remote in-band access

Configure remote in-band access




Activate physical interface

```
8100fl(config)# interface e 1/1
8100fl(config-interface-gig1/1)# no shutdown
```

Assign IP address to VLAN 1 interface

```
8100fl(config)# interface vlan 1
8100fl(config-interface-vlan1)# ip address 10.1.1.1/24
```



If Telnet is enabled on the 8100fl and a virtual terminal session is available, this host can initiate a Telnet session

Rev 6.11

Student Guide: 2-17

16

To configure an 8100fl for in-band management, enable at least one of the switch's ports and associate an IP address with the appropriate VLAN or physical interface.

In the example above, the first port on slot 1 of the 8100fl (*interface e 1/1*) is connected to an edge switch that provides connectivity for a management host in VLAN 1. Because port e 1/1 is disabled by default, it must be specifically enabled with *no shutdown*.

In the example, the administrator creates a virtual interface (*interface vlan 1*) that is associated with the VLAN 1 broadcast domain. With this configuration in place, an authorized user can access the 8100fl through port e 1/1 or through any port that is enabled in the future and is a member of VLAN 1.

Alternatively, the administrator could assign the address directly to the port. However, this option decreases the flexibility available in future configurations. When an IP address is assigned to a port on the 8100fl, the port is removed from VLAN 1 and placed in its own broadcast domain. No other ports can be added to the broadcast domain. The port cannot be made a tagged member of any VLANs unless the IP address applied directly to it is removed.

Configure time services

Configure time services



To set the 8100fl clock manually:

```
8100fl(config)#clock set <HH:MM:SS> <day> <month> <year>
```

To configure the 8100fl to synchronize with a time server:

```
8100fl(config)#ntp server <IP address>
```

To view current clock settings:

```
8100fl(config)#show clock
%13:57:11 UTC Mon Jan 23 2006
```

The system clock of the 8100fl can be set manually or can be configured to synchronize its clock with a Network Time Protocol (NTP) server.

When an 8100fl is booted for the first time, the system clock is unset. If an NTP server is identified in the startup configuration, the switch will synchronize its clock with the server. If the system is set manually and the system is shut down, the time setting will be maintained for about two hours. After that period, it will be necessary to reconfigure the time.

View 8100fl logs

View 8100fl logs



The 8100fl logs system events in a buffer. Events include:

- Successful and unsuccessful logon attempts
- Physical system failures
- File system operation failures

To view the logs for all types of events:

```
8100fl# show logging
```

Can be filtered to only show specific types of events, including:

- Authentication
- SNMP messages
- OSPF
- RIP
- Spanning Tree

Access help to learn filter parameters:

```
8100fl# show logging ?
```

Rev 6.11

Student Guide: 2-19

18

By default, the 8100fl logs system events, such as authentication attempts and physical system failure, in a system memory buffer.

To display the current log, enter *show logging*, as shown above. You can limit the output to information about specific types of events by using a filter. For instance, to display log entries related to OSPF, issue *show logging ospf*. The *show logging* command also will accept *all* as a filter, but the effect is the same as *show logging* with no argument.

Configure syslog

Configure syslog



To configure logging to an external syslog server:

```
8100fl(config)# logging host <ip_address>
8100fl(config)# logging source <ip_address_of_8100fl_interface>
```

Events sent to syslog server may be filtered by event type

- For example:

```
8100fl# logging host <ip_address> ospf
```

As with other ProCurve switches, 8100fl system events can be logged to an external syslog server. Note that the configuration of a syslog server requires two commands, one to specify the address of the server and one to specify the router interface that will be used as a source address for syslog messages.

The output to the syslog server can be filtered according to event type, using the same filters used in *show logging*.

System monitoring and software update

System monitoring and software update



✓ *Initial system configuration*

System monitoring and software update

- Fans and power supplies
- Software update
- Status of management and fabric modules

Rev 6.11

Student Guide: 2–21

20

The rest of the Module 2 will describe tools for monitoring 8100fl system performance and for updating system software.


Chassis: Fan Trays

Chassis: Fan Trays



The Fan Trays each contain three fans

- No single fan failure is catastrophic
- The failure of a *whole* Fan Tray makes half of the system unusable
- Slot Covers are important for proper airflow/pressure control



Rev 6.11

Student Guide: 2-22

21

The 8108fl and 8116fl both ship with two fan trays, both of which are required for proper functioning. Each fan tray contains three fans. The fan tray can maintain proper temperatures with one failed fan, as the other fans create sufficient positive pressure to cool the chassis.


However, the system will not function properly if an entire fan tray fails because each fan tray is responsible for cooling half of the chassis. Failure of a fan tray will render half of the system unusable and generate Temperature Warnings. If the temperature climbs high enough, the system will shut down.

Chassis: Power supplies

Chassis: Power supplies



- Up to three power supplies supported
- Bus connected
 - Any one power supply can power an 8108fl chassis
 - No priority order for power supplies
 - However, there is a one-to-one power cord-to-power supply connection



Rev 6.11

Student Guide: 2–23

22

The 8100fl supports up to three redundant power supplies that are installed through slots at the bottom of the chassis. The 8108fl can function properly with any one of the supplies connected. The 8116fl can function properly with two supplies connected.

The 8108fl includes one power supply with the basic kit; a second power supply is necessary to ensure power redundancy. The 8116fl includes two power supplies. A third power supply is necessary to ensure power redundancy in an 8116fl with more than eight populated slots.

Monitor current temperature and thresholds

Monitor current temperature and thresholds



View the current temperatures for all installed modules:

```
8100fl# show environment temperature
```

```
Interface-Modules :
Slot    Temperature
-----
IM-1    37
```

```
Fabric-Modules :
Slot    Temperature
-----
FM-A    42
FM-B    37
```

```
Management-Modules :
Slot    Temperature
-----
MM-A    53
MM-B    51
```

View temperature thresholds that trigger system events:

```
8100fl# show environment thresholds
```

```
Temperature Limits (deg C) :
Type      Normal    Warning    Critical    Shutdown
-----
Hot-Point 40        72        74         78
```

Rev 6.11

Student Guide: 2-24

23

Like all devices of its kind, the 8100fl is sensitive to overheating. To view the current temperature for all installed modules, enter the *show environment temperature* command.

The system has a set of default thresholds that will trigger the generation of system events in the 8100fl log. For example, if the temperature of any module reaches 72°C, the system will generate a message indicating the temperature is approaching “warning” level. The system will generate another event if the temperature reaches the “critical” temperature threshold of 74°C. The system will shut down if the temperature reaches 78°C.

Although it is not usually necessary or recommended, you can change the thresholds using a *set-temperature* command at the global configuration level. See “Performance Monitoring” in the *Management and Configuration Guide* for more information.

Monitor status of fans and power supplies

Monitor status of fans and power supplies



Display status of power supplies:

```
8100fl# show environment power
```

Power module	Present	Status
Module-1	Yes	Ok
Module-2	Yes	Ok
Module-3	No	N/A

Display operational status of all six fans in both Fan Trays:

```
8100fl# show environment fans
```

Tray	Unit	Config RPM	Actual RPM	Hw Version	Fw Version
Module-A	1	4000	3851	3	4
Module-A	2	4000	3859	3	4
Module-A	3	4000	3923	3	4
Module-B	1	4000	3824	3	4
Module-B	2	4000	3900	3	4
Module-B	3	4000	3813	3	4

Rev 6.11

Student Guide: 2-25

24

Similar to the *show chassis* command on the 9300m and 9408sl, the *show environment* command enables you to view the status of power supplies and fans.

Power supply status

The output of *show environment power* presents comprehensive information about the status of installed power supplies. As shown, the “Present” column in the output indicates whether a power supply is installed in each slot. The “Status” column indicates whether the installed power supply is receiving AC power. The Status column will display “Failed” for any slots where a power supply is installed, but has failed or is not plugged in.

Fan status

The output of *show environment fans* indicates the status of all fans in each of the 8100fl’s two fan trays. As shown, the output indicates whether the measured RPM for each fan is within configured limits.

Display current software version

Display current software version



Display system uptime, module inventory, and current software version:

```
8100fl_1B# show version
ProCurve Networking Switch 8100fl Series System Software
Version CY.02.03.0010
Copyright (c) 1998-2005 by ProCurve Networking.
Compiled on Sun Jan 29 17:17:33 PST 2006
Bootloader Version CY.02.02.0004
Switch uptime is 38 minutes, 6 seconds
System restarted by cold reset
System image file is ms-CY.02.03.0010.ver

ProCurve 8108fl chassis
2 Management-Modules
2 Fabric-Modules(8108fl)
1 Std 10 Port 100/1G Copper
```

As well as showing the current software version, the *show version* command displays the system uptime and an inventory of currently installed modules.

Obtain technical support information

Obtain technical support information



Obtain configuration and process status information that is of interest to technical support personnel

8100f1_1B# **show tech-support**

Output includes detailed information on system status, including:

- Running configuration
- Module status
- Route tables and routing configuration
- Environment variables
- IP interfaces
- LAGs

Output can be filtered to display only:

- Basic information about system version, status of modules, and running configuration
- Layer 2 information
- Layer 3 information

Rev 6.11

Student Guide: 2-27

26

The *show tech-support* command enables you to obtain system information that will be useful to technical support personnel in the event of a support issue.

When entered with no filters, the command outputs multiple pages of detailed system information. While this output is difficult to interpret online, it can be captured to a log and sent to Technical Support personnel for analysis.

You can also use command parameters to limit the output.

Redundant management module operation and monitoring

Redundant management module operation and monitoring



When a second management module is installed, it assumes the role of standby module

- Startup-config on active or “primary” management module is copied to standby or “secondary” module
- Running-config, files in flash memory, and memory resident tables are NOT synchronized between primary and secondary modules

To display current status of modules that support redundancy:

```
8100fl# show redundancy
Slot  Module Description      Model      Redun    Switch    Status
-----
MM-A   Management-Module          J8731A     primary  auto      enabled
MM-B   Management-Module          J8731A     secondary auto      enabled
FM-A   Fabric-Module(8108fl)      J8729A     secondary auto      enabled
FM-B   Fabric-Module(8108fl)      J8729A     primary  auto      enabled
```

To cause a management module redundancy switchover:

```
8100fl# redundancy switchover management-module
Switchover may cause service interruption, proceed? [no]:
```

Rev 6.11

Student Guide: 2-28

27

When a second management module is installed in an 8100fl with a functioning first module, the first module retains its active role and the second module becomes the standby module.

The startup configuration stored in the active module is copied to the standby module as soon as it becomes available. During switch operation, the configuration files on the modules are synchronized each time an administrator saves the running configuration. However, no other files in the primary module’s flash memory are synchronized to the secondary module. The running configuration stored on the active module is not synchronized to the standby module. If the secondary module becomes the primary due to a switchover event, the module loads the startup configuration from flash memory into RAM and uses it as the running configuration.

When the standby module becomes active during switchover, the management port maintains the configured IP address. However, the MAC address of the management interface becomes the address displayed on the face of the active management module.

This change in MAC address does not affect the interface ports because their MAC address is based on the 8100fl backplane MAC address, not on the management module address.

If the 8100fl is physically accessible, you can determine which module is active by viewing the “Active” LEDs on both modules. If the switch is not accessible, you can issue the *show redundancy* command to learn the status of fabric and management modules. To view the management modules alone, issue *show redundancy management-module*.

To force a switchover, enter *redundancy switchover management-module*. Note that you can use this command to force a switchover of either management or fabric modules. The *switchover* command can help to minimize downtime during software update of the management module or while replacing a failed fabric module. These processes will be described later in this module.

Normal redundant management module behavior

Normal redundant management module behavior



Neither management module slot has priority over the other

- Primary/secondary module roles determined by election process during cold start
- Executing a switchover causes secondary module to become primary and resets all fabric and interface modules

Each management module

- Runs under its own software, independent of the other module's software
- Has two distinct software banks, which may contain the same or different software versions

Use *boot system* to reboot primary or secondary management module using either software bank

Rev 6.11

Student Guide: 2–30

28

When an 8100fl with redundant management modules is powered on, the modules participate in an election process that determines which will be the primary module. Neither management module slot has priority. You cannot configure the switch to use a particular module as primary.

Switchover disrupts service for up to two minutes. During switchover, the new primary module forces a reset of the fabric and interface modules and downloads card-specific portions of its software into the appropriate modules. At the start of the switchover, the fabric and interface modules go off-line and their Status LEDs blink. When each module comes online, its LED turns solid green. The console displays a message saying that the fabric and interface modules are up and the system is back in working order.

As with other ProCurve switches, the *reload* command restarts the switch. The *boot system* command enables you to boot either management module and to specify a software bank. The usage and options for this command will be discussed later in this module.

Management module software banks

Management module software banks



To display the management module information, including each bank's software version:

```
8100fl# show module management a
Slot                               :MM-A
Ports                             :0
Module-Type                       :Management-Module
Model                             :J8731A
Admin                             :enabled
Power                             :power on
Status                            :OK
Running Software                  :Bank-B
Bank-A Software                   :CY.02.03.0010
Bank-B Software                   :CY.02.03.0010
Bootloader Version                :CY.02.02.0004
Major HW Version                  :2
Minor HW Version                  :1
CPLD Version                      :cpld 1 9, cpld 2 5, cpld 3 4
Serial Number                     :US515SP103
Manufacture Code                  :00
Manufacture Date                  :19/5
CLEI Code                         :0
CPU Memory Size                   :1024 MB
```

Rev 6.11

Student Guide: 2-31

29

The 8100fl management module stores two software images in separate banks known as Bank A and Bank B. Neither area is designated as primary or secondary, which means that either image can be used to boot the system. The output from the *show module* command displays information on the software version stored in each bank and about the bank used to most recently boot the switch.

To view comprehensive information on all system modules, issue *show module all*.

Rebooting management modules

Reboot management modules



Target secondary management module with *boot system* command:

```
8100fl# boot system management-module a bank-A
```

- Secondary module reboots to software bank A
- No impact on primary management module
- No service interruption

Target primary management module with *boot system* command:

```
8100fl# boot system management-module b bank-A
```

- Primary module reboots to software bank A
- Secondary module becomes primary, resets fabric and interface modules
- Approximately two-minute service interruption due to management module switchover

Rev 6.11

Student Guide: 2-32

30

The *boot system* command enables you to reboot an 8100fl management module using the system image stored in either bank. As described earlier, the “A” and “B” labels assigned to 8100fl management modules do not indicate priority. The primary module is the module that won the election process during the last cold start. As described earlier, enter *show redundancy* to learn which module is primary.

The primary use for this command is to change the active software bank on a management module, especially after copying a new software image to one of the banks. The rebooting of the secondary management module does not affect system performance. When the primary module reboots, it will surrender its primary status. The secondary module will become primary and reset the fabric and interface modules.

The extent of the interruption caused when the primary module fails depends on the status of the secondary module. If the secondary module is available, the system will experience a management module switchover. If the secondary module is not available, the system will cold restart.

Normal redundant fabric module behavior

Normal redundant fabric module behavior



- Fabric modules do not require specific configuration or direct software update
- Neither fabric module slot has priority
 - The active module is the one that won the election process during cold start
- To cause a fabric module redundancy switchover:


```
8100fl# redundancy switchover fabric-module
Switchover may cause service interruption, proceed? [no]: y
```
- Fabric module switchover does not require reboot of line cards or any other modules
 - Does not result in downtime

Rev 6.11

Student Guide: 2–33

31

The 8100fl fabric module does not require any configuration and typically operates without user intervention. The module does not require direct software updates because it receives its software from the management module during system operation.

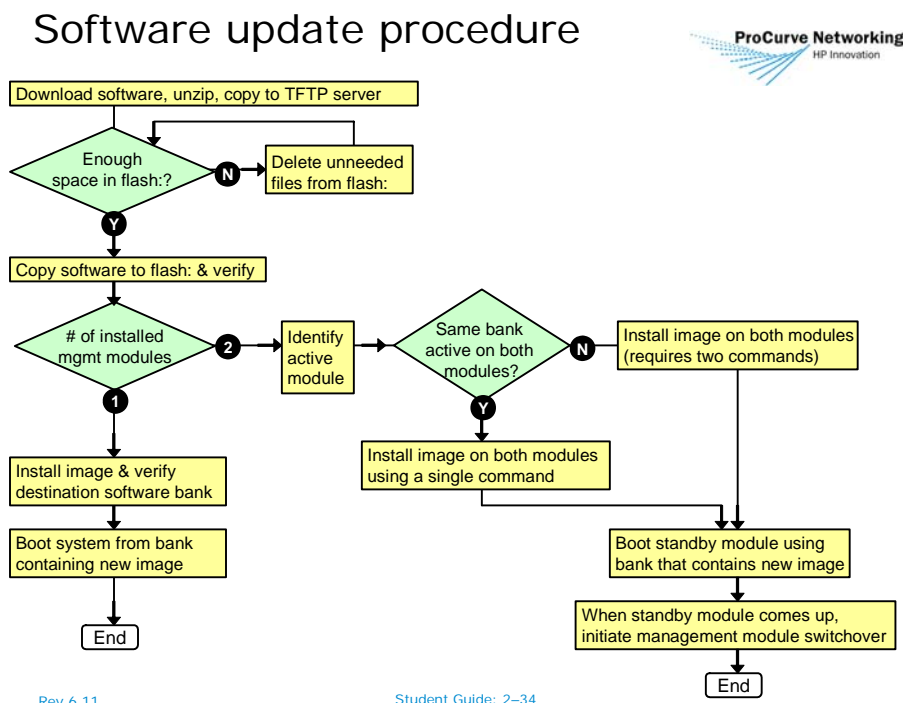
If a fabric module must be replaced in a system with redundant modules, you can minimize disruption by forcing a fabric switchover with the *redundancy switchover fabric-module* command. Because all control data is copied to both fabric modules during the forwarding process, a fabric module switchover does not cause any perceptible service interruption. The measured interruption time typically will be under 200 ms. After the switchover is complete, the failed module can be replaced with no further impact on system performance.

In a system with a single fabric module, the module's failure will disable the switch from forwarding any traffic.

You can force an individual fabric module to reload its software by issuing the *reload fabric-module [a/b]* command. In a system with redundant fabric modules, reload of the active fabric module causes a switchover event. A reload of both modules, or of the only module in a non-redundant system, will cause a service interruption.

While the fabric module is booting up, the module's Status LED blinks green and the module does not forward any traffic.

Software update procedure



The procedure used to update software on the 8100fl is similar to the procedures for other ProCurve switches. However, the process for the 8100fl is more complex because the switch supports redundant management modules with two banks for software images on each management module.

Procedure for single module

The procedure for updating software on an 8100fl with a single management module includes the following steps:

1. Download and unzip the distribution file and place the image on a computer that can act as a TFTP or SCP server.
2. Determine whether enough free space is available in the switch's flash memory to hold the image file. The images are approximately 30 MB in size.
3. Copy the image to flash memory and verify the transfer.
4. Install the image on the management module.
5. Verify which software bank received the new software.
6. Reboot the system using the bank that contains the updated software.

Procedure for redundant management modules

The update of a system with redundant management modules requires the same basic steps. However, there are more options for completing the software update. The procedure illustrated on the flow diagram illustrates a software update method that minimizes system downtime. However, other procedures may be appropriate for specific situations.

When installing the image on the modules, you can install them with a single command if the modules are using the same active bank ID. If they are running software using different bank IDs, you must install the image on each module with a separate command.

If you follow the process for a single module and simply reboot the primary management module after the image is installed, the *boot system* command will initiate a switchover. The secondary management module will become the primary. However, because the module was not explicitly booted with the new image, it will be running the previous version when it finishes restarting.

The procedure illustrated above recommends that you begin by booting the secondary management module using the software bank that contains the new image. The module will boot with the new software and remain in the secondary role. After the secondary management module has booted and you have confirmed that its status is “OK” in the *show module* output, you can initiate a management module switchover using the command *redundancy switchover management-module*. This will cause the secondary management module to become the primary management module. The management module will download the new software into the line and fabric cards, causing the entire system to begin operating under the new software.

The management module that was the primary module at the beginning of the scenario will become the secondary module. To complete the update process, you must boot the secondary module using the bank that holds the new image. You may wish to delay this step until you have confirmed that the new software is functioning satisfactorily on the primary management module.

If the secondary module is running the previous version of software, you can return to that version at any time by executing another management module switchover. After both modules have been rebooted with the new image, you can revert to the previous version by following the same procedure you used for the update. Boot the standby module using the bank that contains the previous software and then execute a switchover.

As with any software update, you should read the release notes included with the software distribution. In some cases, you may need to perform other procedures, such as updating Boot ROMs on individual modules. If such procedures are required, they will be fully documented in the release notes.

Transfer software to flash memory and verify

Transfer software to flash memory and verify



Use TFTP to transfer an image from a server to 8100fl flash memory:

```
8100fl# copy tftp://10.1.1.104/CY.02.03.0010 flash:
```

Verify that the file has been successfully copied:

```
8100fl# dir flash:
Directory of flash:/
```

```
-rw- 29763878 Thu Jan 1 03:32:42 1970 CY.02.03.0010.swi
-rw- 218 Sun Jan 25 07:11:49 1970 sample.cfg
-rw- 524296 Thu Jan 1 03:39:32 1970 FabBootROM.CY.02.02.0024
-rw- 524296 Thu Jan 1 03:39:38 1970 IntBootROM.CY.02.02.0024
-rw- 29761958 Thu Jan 1 01:23:35 1970 CY.02.03.0004
-rw- 520840 Thu Jan 1 03:38:57 1970 MgmtBootROM.CY.02.02.0004
```

Rev 6.11

Student Guide: 2-36

33

To copy a software image to flash memory, use the *copy* command with the URL of the TFTP server specified by name or IP address, as shown above. Specify *flash:* as the destination for the file. While other ProCurve switches use *flash* as a keyword, *flash:* is a device name on the 8100fl.

To copy the same image using SCP, use
scp://remoteUser@10.1.1.104/CY.02.03.0010 flash:

Unlike other ProCurve switches, the 8100fl does not support the *show flash* command. Instead, the 8100fl uses the *dir* command combined with the *flash:* device name, in the same way as a Windows- or UNIX-based computer system.

Use *dir flash:* before copying the image to the 8100fl to ensure the flash storage has sufficient free space. You should also use *dir flash:* after copying the file to verify that it is the correct size.

Install new software image

Install new software image



If switch has a single management module:

```
8100f1# image install flash:CY.02.03.0010
```

If switch has redundant modules with different active software bank IDs:

```
8100f1# image install flash:CY.02.03.0010 management-module a
```

```
8100f1# image install flash:CY.02.03.0010 management-module b
```

If switch has redundant management modules with the same active software bank ID:

```
8100f1# image install flash:CY.02.03.0010
```

```
Software installation has started, may take a while...
Upgrading is in progress --- Type q to stop polling upgrade status
..... Upgrade succeeded
8100f1#* <date and time>:MM-A:MCHSM-W-CM_EVENT_WARNING: starting
software upgrade.
*<date and time>:MM-A:MCHSM-W-INSTALL_SUCCEEDED: Installation
succeeded
```

Rev 6.11

Student Guide: 2-37

34

After the new image is copied to flash memory, it must be installed on each management module using the *image install* command. If the switch has a single management module, follow the *image install* command with the device name *flash:* and the name of the file that contains the image, such as *flash:CY.02.03.0010*. You can use the same command if the system supports redundant management modules and both are using the same active bank ID. However, if the modules have different active software bank IDs, the image must be installed separately on each module, using *image install flash:<file-name> management-module [a/b]*.

The *image install* command always installs the image in the inactive bank. You cannot target a specific bank with the *image install* command.

The installation process may require a minute, but you can perform other configuration tasks while waiting. As shown above, you can stop polling the image installation process and return to the CLI prompt by entering *q*.

Verify image installation

Verify image installation



After the installation is complete, use *show module* to verify each bank's software version

```

Slot                               :MM-A
Ports                              :0
Module-Type                        :Management-Module
Model                             :J8731A
Admin                             :enabled
Power                             :power on
Status                            :OK
Running Software                   :Bank-B
Bank-A Software                    :CY.02.03.0010
Bank-B Software                    :CY.02.03.0005
Bootloader Version                 :CY.02.03.0004
Major HW Version                   :2
Minor HW Version                   :1
CPLD Version                       :cp1d 1 9, cp1d 2 5, cp1d 3 4
Serial Number                      :US515SP103
Manufacture Code                   :00
Manufacture Date                   :19/5
CLEI Code                         :0
CPU Memory Size                    :1024 MB

```

New image was
installed in Bank-A
(the *inactive* bank)

Rev 6.11

Student Guide: 2-38

35

After the installation is complete, use *show module* to confirm that the new image is installed in the management module's inactive software bank. The module still must be rebooted to load the new image.

Boot system with new image

Boot system with new image



The *boot system* command offers many options

```
8100fl# boot system ?
bank-A          - Boot from image in bank_A
bank-B          - Boot from image in bank_B
chassis         - Whole chassis
management-module - Management-Module

8100fl# boot system bank-A ?
<cr>

8100fl# boot system bank-A
Proceed with boot system? [no]: y
System/Module will reboot with new software...
Router1#*Fri Jan  9 09:05:34 1970:MM-A:MCHSM-W-SWAP_SUCCEEDED: Boot
system completed - system/module rebooting...
*Fri Jan  9 09:05:34 1970:MM-A:MCHSM-W-CM_EVENT_WARNING: MM-B is
reloading.
```

When the system comes back up, the management module, fabric module, and all interface modules will use the new software

The final step in the update process is to activate the new software by rebooting the management module. To complete this step, use the *boot system* command with the correct software bank specified, as shown above. You cannot use the *reload* command in this situation

Module 2 summary

Module 2 summary



In this module, you learned how to perform basic configuration tasks on the ProCurve Switch 8100fl

Topics included:

- 8100fl CLI navigation
- Saving and viewing changes to configuration files
- Backing up configuration files
- Enabling remote management, both in-band and out-of-band
- Updating 8100fl software images
- Configuring passwords for management access
- Monitoring system performance

Rev 6.11

Student Guide: 2–40

37

Module 2 described the tools and processes for performing basic configuration of the ProCurve Switch 8100fl. Specific topics included:

- CLI navigation
- Configuring device security
- Saving changes to configuration files
- Backing up configuration files
- Enabling remote management
- Updating software images

Learning check

Module 2

1. What special equipment is necessary to connect to the console port of a ProCurve Switch 8100fl?

.....
.....

2. What is the default status of ports on the 8100fl?

.....
.....

3. Where is the software image installed on an 8100fl?

- a. Fabric module
- b. Chassis backplane
- c. Management module

4. After a new software image is installed on an 8100fl, what step is necessary to load the software into memory?

.....
.....

Module 2 lab overview

Module 2 lab overview



Lab has three groups of two students

You will work with a partner in one group

- Each group will configure two 8100fl switches and two 5300xl switches
- For most labs, groups are independent
- Groups will interconnect through core router in Module 5

Module 2 lab activity has two parts:

- 2-1: Initial Configuration and Monitoring
 - Device security
 - Remote access
 - Backup configuration
- 2-2: Updating Software on Redundant Management Modules
 - Copy image to flash:
 - Install image on management modules
 - Boot management module
 - Force management module switchover

Rev 6.11

Student Guide: 2–43

40

Each of the remaining modules is accompanied by a hands-on activity designed to give you practical experience with the features of the ProCurve Switch 8100fl.

As you will see in the accompanying *Lab Activity Guide*, the activities lab consists of three groups of two students. You and a partner will be assigned to one group and will configure two 8100fl switches and two 5300xl switches to interconnect and to implement features described in the course.

In the final activity, in Module 5, all three groups will be interconnected through a core router. To support this interconnected topology, the labs implement an addressing scheme that ensures that each host and network has a unique address throughout the activities. All lab addresses are in the range of 10.0.0.0/24. For all activities, however, the second octet should be used to designate the group number. Addresses in Group 1 will be 10.1.x.x./24. Group 2 addresses will be 10.2.x.x./24. Group 3 addresses will be 10.3.x.x./24.

Module 2 activities

The lab activities for Module 2 will provide you with an opportunity to perform initial configuration and monitoring tasks on the ProCurve Switch 8100fl. As shown above, these tasks include configuring device security, remote access, and backing up your 8100fl configuration.

In the second part of the lab, you will update the software on an 8100fl using the procedure described in the module for updating a system with redundant management modules.

Detailed instructions are provided in your *Lab Activity Guide*.


Objectives

After completing this module and the accompanying hands-on activities, you will be able to:

- Describe the link aggregation features of the ProCurve Switch 8100fl
- Compare and contrast the link aggregation features of the 8100fl with the aggregation features of other ProCurve switches
- Describe the types of IP interfaces available on the 8100fl
- Given a set of customer requirements, define VLANs on the 8100fl
- Given a network design document, configure ports to be members of one or more VLANs
- Configure link aggregation on the 8100fl
- Confirm link aggregation and VLAN configurations using *show* commands

Link Aggregation on the 8100fl

Link aggregation on the 8100fl



Use link aggregation to provide additional bandwidth to edge switches

- Collection of aggregated links is called a *Link Aggregation Group (LAG)*
- Each LAG controlled by a numerically identified *aggregator* that handles traffic distribution
- Ports associated with an aggregator by its numeric ID
- Up to 80 aggregators may be defined on a switch
- Maximum number of ports in LAG is eight

Rev 6.11

Student Guide: 3-2

3

To fulfill its role as an interconnect fabric, the ProCurve Switch 8100l offers extensive support for link aggregation. This technology enables the 8100fl to provide adequate uplink bandwidth for edge switches that do not have sufficient capacity in any one link.

In the example above, the ProCurve Switch 2848 supports 40 gigabit clients with 40 gigabit edge ports and four gigabit uplink ports. To provide enough capacity for this switch's 40 clients (40 x 1000 x estimated 10% utilization = 4 gbps), you can associate the 2848 switch's four fiber uplinks with a numbered "trunk." On the 8100fl, you associate ports e 1/9, e 1/10, e 6/9, and 6/10 with an "aggregator." The ProCurve Switch 3400cl-48G does not require a port trunk because it provides an uplink that is connected to a 10-GbE interface on the 8100fl.

Aggregation on the 8100fl

The link-aggregation technologies on the 8100fl are completely compatible with other ProCurve switches. However, the terminology and configuration processes on the 8100fl differ slightly from other ProCurve switches.

For instance, on the 8100fl, a collection of aggregated link is called a Link Aggregation Group (LAG). The term "trunk," used on other ProCurve switches to describe an aggregation group, is used on the 8100fl to describe a VLAN mode. (This topic will be discussed in detail later in this module.)

The 8100fl enables an administrator to create up to 80 LAGs. Up to eight ports can be assigned to a LAG. LAGs can also be configured with no port members. The ports assigned to an 8100fl LAG can span any number of interface modules. Furthermore, the ports can use different media, but must be of the same speed.

Like other ProCurve switches, the 8100fl supports dynamic link aggregation using IEEE 802.3ad Link Aggregation Control Protocol (LACP). However, manual link configuration usually allows the administrator/installer better control over link utilization and configuration than is possible with LACP. For more information on configuring LACP, see the “Link Aggregation Configuration” in the *Management and Configuration Guide*.

Configure Link Aggregation Groups (LAGs)

Configure Link Aggregation Groups (LAGs)



General configuration steps for LAG:

- Create aggregator with numeric ID
- Associate ports with the LAG by its ID
- Reverse the default *shutdown* state for ports in the LAG if ports are not already enabled

```
8100fl(config)#aggregator 1
8100fl(config-lag-1)#interface e 2/9
8100fl(config-interface-gig2/9)#lag 1
8100fl(config-interface-gig2/9)#no shutdown
8100fl(config-interface-gig2/9)#interface e 2/10
8100fl(config-interface-gig2/10)#lag 1
8100fl(config-interface-gig2/10)#no shutdown
8100fl(config-interface-gig2/10)#interface e 6/9
8100fl(config-interface-gig6/9)#lag 1
8100fl(config-interface-gig6/9)#no shutdown
8100fl(config-interface-gig6/9)#interface e 6/10
8100fl(config-interface-gig6/10)#lag 1
8100fl(config-interface-gig6/10)#no shutdown
```

Rev 6.11

Student Guide: 3–4

4

The strategy for configuring link aggregation on the 8100fl is very different from other ProCurve switches. While other ProCurve switches use the same statement to define a numeric identifier for the trunk and to add ports to the trunk, the 8100fl requires you to first create an “aggregator” and assign a numeric identifier between 1 and 80.

The creation of an aggregator moves the CLI to the aggregator context. In the example, the administrator moves to the port configuration context by entering *interface e 2/9* and associates the port with the LAG by entering the *lag* command, followed by the ID (1) of the aggregator created earlier. The process is repeated for all ports in the LAG.

Note that each port also is enabled with the *no shutdown* command. This step could be completed at any convenient time. For instance, the administrator might not enable the ports on the 8100fl until after configuring link aggregation on the switch that will terminate the other end of the aggregated links. This would enable the administrator to avoid inadvertently creating a loop, and possible broadcast storm, while the configuration in progress.

View and modify LAG attributes

View and modify LAG attributes



```
8100fl# show lag 1 attributes
LAG Name       : Lag1
Admin status   : Up
Trunk status    : Access
Native VLAN    : 1
VLAN membership in : 1 VLAN (1)
STP status     : Disabled
Bridging Mode  : Address Bridging
Aggr Mode      : Layer 3
Protocol       : none

8100fl(config)# interface lag 1
8100fl(config-interface-lag1)# aggr-mode ?
  l3-based      - Layer-3-based (IPv4 SA, DA) link assignment
  l4-based      - Layer-4-based (IPv4 SA, DA, protocol, source port,
                  destination port) link assignment
  mac-based     - Layer-2-based (SMAC, DMAC, VID) link assignment
8100fl(config-interface-lag1)#aggr-mode l4-based
```

Rev 6.11

Student Guide: 3-5

5

Variations of the *show lag* command display information about the LAG's attributes and members. When ports are assigned to a LAG on the 8100fl, they operate as a single logical link that can carry any combination of switched or routed traffic. In this example, no VLANs have been defined, and all ports are in VLAN 1. The process for associating LAGs with user-defined VLANs will be described later in this module.

The aggregator that controls a LAG is responsible for distributing outbound traffic among port members by hashing information in each packet and using the resulting value to select an outbound port. Like other link aggregation implementations, this ensures that packets from a flow will use the same port as long as the port remains up and a member of the LAG. If the port becomes unavailable, the aggregator will reallocate the packet flow to another port. The aggregator also is responsible for receiving incoming traffic and facilitating its forwarding based on header content and other switch configuration parameters.

At default settings, the aggregator selects a link based on a packet's source and destination IP address. However, you can change the aggregation mode so that the aggregator will assign traffic to links based on Layer 2 or Layer 4 information.

In the example, an administrator uses the *aggr-mode l4-based* command to configure the LAG to use Layer 4 attributes for traffic distribution. This configuration can be useful when two hosts (such as server and backup server) exchange large amounts of information.

View and modify LAG port member status

View and modify LAG port member status



To view the status of a LAG's port members:

```
8100fl(config)# show lag 1 member-ports
```

Lag Id	Designated Port	Member Ports	Status
Lag1	Gig2/9	Gig2/9	enabled/up
		Gig2/10	enabled/up
		Gig6/9	enabled/up
		Gig6/10	enabled/up

Ports may be easily removed from a LAG:

```
8100fl(config)# int e 6/10
8100fl(config-interface-gig6/10)# no lag 1
```

```
8100fl(config)# show lag 1 member-ports
```

Lag Id	Designated Port	Member Ports	Status
Lag1	Gig2/9	Gig2/9	enabled/up
		Gig2/10	enabled/up
		Gig6/9	enabled/up

An aggregator cannot be deleted if it has at least one port member:

```
8100fl(config-interface-gig6/10)# no aggregator 1
```

```
0w0d: %SYS-7-CONFIG_RESULTS: command failed
8100fl(config)#
```

Rev 6.11

Student Guide: 3-6

6

After defining a LAG, you can dynamically add or delete ports. During these changes, the aggregator ensures that traffic continues to be seamlessly forwarded over the links currently in the LAG.

Before deleting an aggregator, you must remove all references to the aggregator from the configuration. An aggregator cannot be deleted if it has port members. Additionally, a port cannot be moved from one LAG directly into another LAG. Only ports that are currently not LAG members may be added to a LAG.

Before removing ports from a LAG, disconnect redundant links or disable the ports using *shutdown*. A broadcast storm can result when redundant links are left in place when a LAG is removed.

Note that the output for each LAG specifies a “designated port.” When forwarding multicast traffic, the switch will use only the designated port. When OSPF is configured, if the switch uses multiple paths with LAG interfaces, traffic will be forwarded only through the designated port in each LAG.

Default VLAN membership for ports and LAGs

Default VLAN membership for ports and LAGs



Ports and LAGs are members of Default VLAN unless otherwise assigned

LAGs, but not LAG port members, may be assigned to VLANs

```
8100fl# show vlan
VLAN      Name      Status   Ports      Type
1          Default  active   Gig1/1
          Gig1/2
          Gig1/3
... [output omitted]
          Gig2/7
          Gig2/8
          Gig3/1
... [output omitted]
          Gig6/7
          Gig6/8
          Gig6/10
          Tengig7/1
          Tengig8/1
          Lag1
```

Diagram illustrating the replacement of individual ports by a LAG (Lag1) in the VLAN membership list:

- Gig2/9 and 2/10 replaced by Lag1
- Gig6/9 replaced by Lag1

Rev 6.11

Student Guide: 3-7

7

When 8100fl ports are added to a LAG, you can no longer configure their VLAN membership. Instead, you assign VLAN membership to the LAG and the LAG's ports are configured accordingly. As shown above, the ports that are assigned to a LAG do not appear as individual interfaces in the *show vlan* output. In the example, ports e 2/9, 2/10, and 6/9 have been replaced by Lag1, which is the final entry in the list of VLAN 1 members.

Like other ProCurve switches, the 8100fl enables you to define VLANs and assign ports and LAGs as tagged or untagged members of the VLANs. However, as will be discussed later in this module, the configuration procedures and terminology used on the 8100fl are quite different from those used on other ProCurve switches.

IP interface definition

IP interface definition



ProCurve 8100fl supports two approaches to IP interface definition:

- VLAN interfaces
 - Ports and LAG interfaces may be configured to carry untagged traffic for a single VLAN
 - Ports and LAG interfaces may be configured to carry traffic for multiple VLANs (one untagged, multiple tagged)
 - IP addresses are assigned within VLAN interface context
- Port-based interfaces
 - Each port or LAG carries traffic for a single network (does not carry tagged traffic)
 - Multiple port-based interfaces cannot belong to the same broadcast domain
 - IP addresses are assigned within port or LAG configuration context

The 8100fl automatically routes among its locally connected networks

Rev 6.11

Student Guide: 3–8

8

As an interconnect fabric, the ProCurve Switch 8100fl must support edge devices with a wide array of capabilities. To meet this need, the 8100fl enables you to define IP interfaces for VLANs or for individual ports and LAGs. Furthermore, the 8100fl enables you to configure both types of interfaces on a single switch.

In contemporary network designs, VLAN interfaces often are preferred because they offer high levels of flexibility and control, including the ability to enable a single link to carry traffic for multiple broadcast domains. Accordingly, the 8100fl offers complete support for the 802.1Q tagging standard.

Because routing switches are most likely to use VLAN interfaces, the rest of this module will focus on the process for configuring VLANs, assigning their port members, and associating them with an IP interface. In some cases, however, port-based interfaces can be appropriate or necessary. For example, if all of the hosts in a given network are connected to a single Layer 2 edge switch, it may be simpler to assign an IP address directly to the port that serves the edge switch. To configure a port-based interface on the 8100fl, simply enter the *ip address* command in the configuration context for a port or LAG.

Concurrent support for both interface types

The 8100fl can support any combination of port-based and VLAN interfaces concurrently, based entirely on edge requirements. Some routing switches recommend against mixing port-based and VLAN-based interfaces because of the potential for mismatched broadcast domains. However, the 8100fl does not share this limitation. Specific features of the 8100fl that allow mixed interface types will be covered in Module 5.

8100fl VLAN support

8100fl VLAN support



- 8100fl supports tagged and untagged VLAN port membership
- Fully compatible with other ProCurve switches, but uses different VLAN terminology and configuration procedures
- VLANs are created within global configuration context

```
8100fl(config)# vlan 10
```

```
8100fl(config-vlan)#
```

- Ports are associated with VLANs within port or LAG *interface* configuration context

Description	8100fl term	Related command in interface configuration context
Port that carries untagged traffic for a single VLAN	Access port	switchport mode access vlan <id>
Port that carries tagged traffic for multiple VLANs	Trunk port	switchport mode trunk
List or range of VLANs associated with a trunk port (VLAN tag is added to traffic)	Trunk VLAN	switchport trunk-vlans <id>,<id>
VLAN to which untagged traffic received over a trunk mode port should be assigned	Native VLAN	switchport trunk-native-vlan <id>

Rev 6.11

Student Guide: 3-10

9

The VLAN support on the 8100fl is completely compatible with the 802.1Q standard and with other ProCurve switches. However, the 8100fl uses different terminology and configuration procedures.

This is particularly true of the terminology and commands used to describe and configure VLAN port assignment. On the 8100fl, CLI commands do not explicitly mention “tagged” and “untagged.” Instead, the 8100fl uses switch port modes to distinguish tagged from untagged behavior for a given port. To support VLANs, ports on the 8100fl must be configured for one of two modes: *access mode* or *trunk mode*. Additionally, on the 8100fl, you must enter the port configuration context to associate a port with a VLAN. On other ProCurve switches, this task is accomplished in the VLAN configuration context.

Access mode

Access mode enables a port to carry untagged traffic for a single VLAN. In general, access mode is most appropriate for edge switches that directly support hosts. While the 8100fl is not well suited to service as an edge switch, access mode may be appropriate for connecting to legacy devices such as hubs or switches that do not support VLANs and QoS.

Trunk mode


Trunk mode enables a port to carry tagged traffic for any or all of the switch's user-defined VLANs. Because of its flexibility, trunk mode is the most likely configuration for an 8100fl that provides default gateway service for clients connected to VLAN-enabled Layer 2 edge switches.

When configuring trunk mode, you associate individual ports or LAGs with all VLANs whose traffic will be carried with 802.1Q tags. Untagged traffic is associated with the "native VLAN." By default, this is VLAN 1, but you can configure an access port to associate untagged traffic with any VLAN. Like the "untagged" VLAN on other ProCurve switches, the native VLAN is often associated with the network through which the switch is managed.

The next few slides will describe several design scenarios and present the configuration procedures required to implement them.

Access mode example

Access mode example



Legacy example using switches to provide higher speed uplink for hubs

Edge switches in two wiring closets:

- Each have a 100Base-TX uplink to 8100fl
- Are not configured to support VLANs
- Support hosts in the same IP address range (10.1.2.0/24)

8100fl ports that support the edge switches

- Must enable default gateway service for hosts in the network 10.1.2.0/24
- Must be in the same broadcast domain/VLAN
- Will not carry tagged traffic

Edge devices will be managed through IP addresses in the same address range as the hosts

Rev 6.11

Student Guide: 3-12

10

In this example, a ProCurve Switch 8108fl must provide default gateway service to hosts who are connected to legacy hub stacks in two wiring closets. Each stack is aggregated by a ProCurve Switch 2524. No user-defined VLANs are defined on the edge switches, and all ports are untagged members of VLAN 1. Because the links to the edge switches must carry only one VLAN, the appropriate ports on the 8108fl can be configured for access mode.

While this configuration is functional, it is not usually recommended for contemporary networks. In fact, it is usually only suitable for legacy environments with the following features:


1. Edge switch and host requirements place multiple switch ports in the same broadcast domain. If each of the switches and its hosts was in a different broadcast domain, port-based interfaces might be appropriate.
2. No VLANs are configured at the edge, which means all hosts connected to the edge switches are in the same network.
3. No convergence technologies are required. In the example, the hubs do not support QoS, which means it cannot be effectively deployed on the switches. QoS usually requires trunk-mode ports because most QoS solutions place time-sensitive traffic into separate VLANs from best-effort traffic.

Device security issue

In the example, the edge switches could be subject to security breaches because the IP addresses used for switch management must be in the same address range as the hosts. Although users could be prevented from accessing the switches, they still could be subject to broadcast storms and denial of service attacks.

VLAN creation and access port assignment

VLAN creation and access port assignment



Create VLAN 2 in the global configuration context:

```
8100fl(config)# vlan 2
```

Set the switchport mode to "access" for the appropriate ports, associate VLAN 2, and reverse the default shutdown state:

```
8100fl(config-vlan-2)# int e 6/1
8100fl(config-interface-gig6/1)# switchport mode access vlan 2
8100fl(config-interface-gig6/1)# no shutdown
8100fl(config-interface-gig6/1)# int e 6/4
8100fl(config-interface-gig6/4)# switchport mode access vlan 2
8100fl(config-interface-gig6/4)# no shutdown
```

Rev 6.11

Student Guide: 3-14

11

On the 8100fl, you cannot create a VLAN and assign ports to it in a single command, as you can on other ProCurve switches. You must create the VLAN, as shown above, before assigning ports. The CLI will display an error message if you attempt to associate a port with a VLAN that does not exist.

Unlike other ProCurve switches, the 8100fl's VLAN configuration context supports very few commands. You can assign a 12-character name and 180-character description to the VLAN. To enable support for jumbo frames, you can specify a Maximum Transmission Unit (MTU) of up to 9216 bytes.

To configure a port for access mode, you must specify the single VLAN to which this port will belong. The command *switchport mode access* will return an error because it is incomplete if no VLAN is specified.

Like other Q-compliant switches, the 8100fl supports user-defined VLANs with IDs between 2 and 4094. In the example, the switch ports were configured for membership in VLAN 2, but this ID has no special significance.

Display VLAN membership

Display VLAN membership



After ports e 6/1 and e 6/4 have been assigned as access mode switch ports for VLAN 2, they are no longer members of VLAN 1

```
8100fl# show vlan
VLAN      Name           Status  Ports           Type
1          Default       active  Gig1/1
          Gig1/2
          Gig1/3
... [output omitted]
          Gig6/2
          Gig6/3
          Gig6/5
          Gig6/6
          Gig6/7
          Gig6/8
          Gig6/10
... [output omitted]
2          VLAN-2       active  Lag1
          Gig6/1
          Gig6/4
```

Gig6/1 and 6/4
moved to VLAN 2

To view the status of a single VLAN, specify the VLAN ID:

```
8100fl# show vlan id 2
VLAN      Name           Status  Ports           Type
2          VLAN-2       active  Gig6/1
          Gig6/4
```

Rev 6.11

Student Guide: 3-15

12

As with other ProCurve switches, you can use the *show vlan* command to confirm VLAN configuration changes. In the example shown, ports 6/1 and 6/4 have been configured as access mode ports for VLAN 2. Consequently, they are no longer included in VLAN 1. Functionally, this is the same as designating the ports as untagged members of VLAN 2 on another ProCurve switch.


After ports e 6/1 and e 6/4 have been assigned to VLAN 2 and their shutdown state has been reversed, hosts in both switch stacks are in the same broadcast domain. The 8108fl forwards traffic between the two edge switches using Layer 2 information. However, because the 8108fl does not have an IP interface associated with VLAN 2, it cannot route traffic from these hosts to any other networks. The process for defining IP interfaces will be described on the next page.

In the example above, the administrator has used the *show vlan* command to examine the VLAN configuration of all ports on the 8108fl. However, as with other ProCurve switches, you can limit the command output by including a VLAN ID as a parameter, as in *show vlan id 2*.

Note that the prompt in the example indicates that the administrator entered the *show vlan* command in the Privileged EXEC context. It is not necessary to exit the configuration context to issue *show* commands, but it is recommended to avoid configuration conflicts. While the 8100fl supports multiple Telnet sessions, it supports only one configuration session. If you are in the configuration context when another user tries to enter the context, your configuration session will be terminated.

Define IP interface for VLAN

Define IP interface for VLAN



Create interface for VLAN 2 in the global configuration context:

```
8100fl(config)# interface vlan 2
8100fl(config-interface-vlan2)# ip address 10.1.2.1/24
```

View brief status of the newly defined IP interface:

```
8100fl(config-interface-vlan2)#show ip interface vlan 2 brief
Interface          IP-Address      Status  Protocol
Vlan2              10.1.2.1/24    Up      Up
8100fl(config-interface-vlan2)#
```

Rev 6.11

Student Guide: 3-16

13

Because the edge switches in this example operate at Layer 2, they cannot provide default gateway service for connected hosts. Consequently, the 8100fl must be configured to act as their default gateway. To enable this configuration, you must associate an appropriate IP address and mask, such as 10.1.2.1/24, with VLAN 2 on the 8100fl.

As shown, the configuration of an IP interface for a VLAN requires two steps:

1. Create a “VLAN interface” that is associated with the VLAN by its VLAN ID. This is analogous to configuring a virtual interface on the 9300m and 9408sl.
2. Configure an IP address for the interface. In this example, hosts in VLAN 2 will use the address 10.1.2.1 as their default gateway. By assigning this address in the VLAN interface context for VLAN 2, you can allow the 8100fl to provide default gateway service for hosts connected to both Layer 2 edge switches.

To display details on the newly defined IP interface, use the command *show ip interface vlan 2*. For a less verbose display, use *show ip interface vlan 2 brief*.

With this configuration in place, the 8100fl will be prepared to route traffic between VLAN 2 and other IP interfaces that will be configured later. No special configuration is required to enable IP routing. In its current state, of course, the 8100fl will not perform routing duties because only one IP interface has been defined.

View VLAN access mode ports and interface configuration

View VLAN access mode ports and interface configuration



VLAN configuration requires three objects that are associated by use of a common **VLAN ID**

```
8100fl# show running-config
...
vlan 2
!
interface GigabitEthernet6/1
  no shutdown
  switchport mode access vlan 2
!
interface GigabitEthernet6/4
  no shutdown
  switchport mode access vlan 2
!
...
!
interface Vlan2
  ip address 10.1.2.1/24
...
!
end
8100fl#
```

Annotations:

- VLAN must exist before it can be associated with interface(s) (points to `vlan 2`)
- VLAN must be associated with all interfaces that will carry its traffic (points to `switchport mode access vlan 2` in both interface configurations)
- IP address is assigned within VLAN interface configuration context (points to `ip address 10.1.2.1/24`)

Rev 6.11

Student Guide: 3-17

14

As shown, the running configuration of the 8100fl now includes information relating to the configuration of VLAN 2 and its IP interface. Specifically, the output shows:

1. The VLAN and its ID, which must be configured before it can be associated with an interface
2. Port assignments
3. An IP interface and address

Of course, the actual running configuration of the switch includes several items not shown here, including commands relating to device security, but they do not directly affect VLAN configuration and were omitted for brevity.

Deleting a VLAN

In order to delete a VLAN from the 8100fl, you must consider several dependencies among VLAN-related items. Although VLAN 2 is user defined, you cannot delete it as long as it is associated with an IP interface and ports. Consequently, to delete the VLAN, you must take three separate steps:

1. Delete the VLAN interface
8100f1(config)# no interface vlan 2
2. Reverse the switchport command from within the port configuration context for all ports associated with VLAN 2
8100f1(config)# int e 6/1
8100f1(config-interface-gig6/1)# no switchport
3. Delete the VLAN
8100f1(config)# no vlan 2

The first two steps in this procedure can be performed in any order.

VLANs and LAGs: Comparison with other ProCurve switches

VLANs and LAGs: Comparison with other ProCurve switches



Terminology differences	
8100fl	Other ProCurve
Access mode/trunk mode	Untagged/tagged port
Aggregator or LAG	Trunk group

Command differences		
Configuration item	8100fl	Other ProCurve
Untagged VLAN traffic	switchport mode access vlan <id> switchport trunk-native-vlan <id>	untagged <port-list>
Tagged VLAN traffic	switchport mode trunk switchport trunk-vlan <id-list>	tagged <port-list>
Aggregated links	aggregator <id> interface <port> LAG<id>	trunk <port-list> trk<id>

Rev 6.11

Student Guide: 3-19

15

Although the VLAN and link-aggregation features of the 8100fl are completely compatible with other ProCurve switches, the 8100fl uses quite different terminology and commands. The biggest differences concern the use of the “trunk” and the processes for configuring ports for VLAN membership.


The 8100fl uses the term “trunk” to denote a VLAN mode. Configured for trunk mode, a port can carry traffic for more than one VLAN. By contrast, a port configured for access mode can carry traffic for only one VLAN. On other ProCurve switches, “trunk” refers to a group of aggregated links.

The 8100fl refers to an aggregated group of links as a “Link Aggregation Group” or LAG. After ports have been added to a LAG, the LAG can be configured for trunk mode or access mode using the same *switchport* commands applied to individual ports.

The next few pages will describe the process for configuring ports for trunk mode on the 8100fl.

Example of switchport trunk mode

Example of switchport trunk mode



Layer 2 edge switches in this example support multiple VLANs

- Switch 2848 must carry tagged traffic for three VLANs over its aggregated uplink
- Switch 3400cl must carry tagged traffic for three VLANs over a single uplink
- Edge switches are configured to assign untagged traffic to VLAN 1
- Edge switch management addresses are within VLAN 1 IP address range

ProCurve Switch 8108fl:

- Must provide default gateway for hosts in all four VLANs
- Must support VLAN tagging on the links to the edge switches

Rev 6.11

Student Guide: 3-20

16

When the ProCurve Switch 8100fl series is deployed as the interconnect fabric for Layer 2 edge switches, its ports are most likely to be configured for VLAN trunk mode, which enables a link to carry tagged traffic for multiple VLANs.

This example uses the same edge switches and physical port connections as the LAG example presented earlier. However, this example provides more information about the network's logical connections and virtual broadcast domain boundaries.

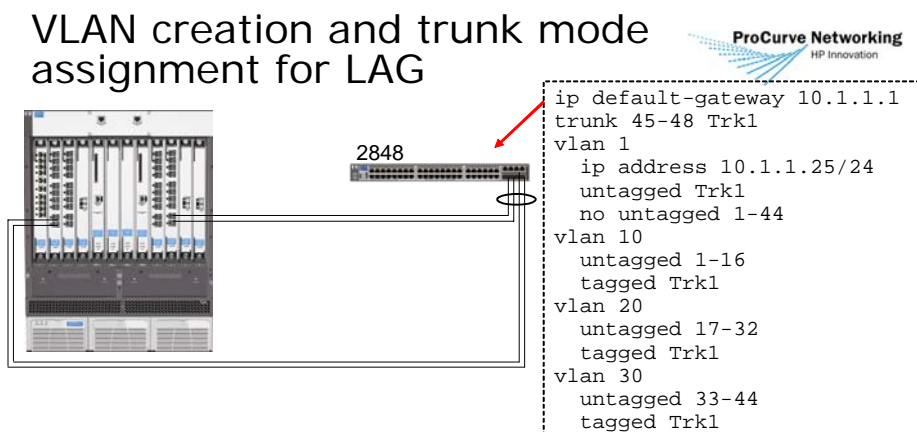
In this topology, hosts in four user VLANs are distributed over two edge switches whose uplinks have already been correctly configured to operate with the 8100fl. The edge switches support hosts in four VLANs (10, 20, 30, 40), with hosts in VLANs 20 and 30 distributed across both switches. Because IP routing is not enabled on the edge switches, the 8100fl must act as default gateway for hosts in all four VLANs.

Furthermore, because each edge switch has an IP address in VLAN 1 that is used for management, the 8100fl ports must carry VLAN 1 over the links to the edge switches and provide default gateway service for the switches.

The VLAN configuration will appropriately associate VLANs with interfaces. This will enable the 8100fl to add the correct tags to traffic it sends to edge switches and to correctly interpret the tags on inbound traffic.

The ProCurve Switch 2848 requires a four-port trunk to provide adequate uplink bandwidth. This trunk and the connected LAG on the 8100fl both have already been configured. The ProCurve Switch 3400cl-48G does not require an aggregated link because it has a 10-GbE uplink.

VLAN creation and trunk mode assignment for LAG



On the 8108fl, create VLANs in the global configuration context:

```

8100fl(config)# vlan 10
8100fl(config-vlan-10)# vlan 20
8100fl(config-vlan-20)# vlan 30
  
```

Define LAG 1 as a trunk mode port and associate VLANs with the LAG:

```

8100fl(config-vlan-30)# int lag 1
8100fl(config-interface-lag1)# switchport mode trunk
8100fl(config-interface-lag1)# switchport trunk-vlans 10,20,30
  
```

Rev 6.11

Student Guide: 3-21

17

Like ports configured for access mode, ports configured for trunk mode cannot be associated with VLANs that are not yet configured on the switch. Consequently, as shown, the first step in configuring trunk-mode ports is the creation of all necessary VLANs. In this case, the 8108fl must be configured for VLANs 10, 20, and 30.

After configuring the VLANs, enter the LAG 1 interface context to configure the LAG for trunk mode. This requires two steps: defining the LAG's switchport mode as "trunk" and associating the trunk with the necessary VLANs. In the example, the VLAN IDs are associated with the VLAN trunk by specifying the IDs separated by commas. If the 2848 supported a range of VLAN IDs, such as VLANs 10, 11, 12, and 13, you could specify them as a range using the command *switchport trunk-vlans 10-13*.

It is not possible to add a port to a LAG after it has been configured for trunk mode. Consequently, you must remove all VLAN associations before adding a port to a LAG.


Notice that Trk1 already has been configured on the 2848 and that it is a tagged member of the user-defined VLANs.

The example also assumes that the LAG has been defined on the 8100fl and that the shutdown state of its ports has been reversed. If the LAG had not been defined, you would need to create the “aggregator” before creating the LAG interface and configuring the VLAN trunk mode. The configuration steps for preparing the LAG are:

```
8100fl(config)# aggregator 1
8100fl(config-lag)# interface e 2/9
8100fl(config-interface-gig2/9)# lag 1
8100fl(config-interface-gig2/9)# no shutdown
8100fl(config-interface-gig2/9)# interface e2/10
8100fl(config-interface-gig2/10)# lag 1
8100fl(config-interface-gig2/10)# no shutdown
8100fl(config-interface-gig2/10)# interface e6/9
8100fl(config-interface-gig6/9)# lag 1
8100fl(config-interface-gig6/9)# no shutdown
8100fl(config-interface-gig6/9)# interface e6/10
8100fl(config-interface-gig6/10)# lag 1
8100fl(config-interface-gig6/10)# no shutdown
```


VLAN creation and trunk mode assignment for port

VLAN creation and trunk mode assignment for port



On the Switch 8108fl, create any additional VLANs required by the edge switch:

```
8100fl(config)# vlan 40
```

Define TenGigabit 8/1 as a "trunk" mode port and associate VLANs:

```

8100fl(config-vlan)# int tengig 8/1
8100fl(config-interface-10gig8/1)# switchport mode trunk
8100fl(config-interface-10gig8/1)# switchport trunk-vlans 20,30,40
8100fl(config-interface-10gig8/1)# no shutdown
  
```

Rev 6.11

Student Guide: 3-23

18

In this example, the 8100fl is configured to support its connections to the 3400cl-48G, which has been configured to support hosts in three user-defined VLANs.

To complete the configuration, the 10-GbE port on the module in Slot 8 of the 8100fl must be configured for trunk mode. The configuration of the 8100fl begins with the creation of VLAN 40, which was not required earlier. Port 8/1 is then configured for trunk mode and VLANs 20, 30, and 40 are defined as *trunk-vlans*. Finally, the default shutdown state of the port is reversed with the *no shutdown* command.

Display VLAN membership for trunk mode ports

Display VLAN membership for trunk mode ports



Trunk mode switch ports associate untagged traffic with the “native” VLAN (Default is VLAN 1)

```
8100fl# show vlan
```

VLAN	Name	Status	Ports	Type
1	Default	active	Gig1/1 Gig1/2 Gig1/3 ... [output omitted] TenGig7/1 TenGig8/1 Lag1	
2	VLAN-2	active	Gig6/1 Gig6/4	
10	VLAN-10	active	Lag1	
20	VLAN-20	active	TenGig8/1 Lag1	
30	VLAN-30	active	TenGig8/1 Lag1	
40	VLAN-40	active	TenGig8/1	

Rev 6.11

Student Guide: 3-24

19

The output of the *show vlan* command indicates that the trunk-mode ports on the 8100fl have retained their membership in VLAN 1. Although the output does not distinguish between tagged and untagged status, the 8100fl’s behavior for trunk mode ports is identical to the behavior of the tagged ports on other ProCurve switches.

For example, on a 3400cl switch, if you assign the uplink (port 49) as a tagged member of VLAN 20 (*vlan 20 tagged 49*), the port remains an untagged member of VLAN 1. Similarly, when a port on the 8100fl is configured for trunk mode, it retains untagged membership in VLAN 1 unless you specifically change its “native VLAN.”

“Native” VLAN on the 8100fl trunk

“Native” VLAN on the 8100fl trunk



Trunk mode port assigns untagged traffic to its “native” VLAN

Native VLAN for all trunk mode ports defaults to VLAN 1

- May be configured independently for each trunk port
- Native VLAN configuration is not necessary if edge switches associate untagged traffic with VLAN 1
- VLAN 1 can be assigned as “native-vlan” (untagged), not as a “trunk-vlan” (tagged)
- If edge switches associate untagged traffic with some other VLAN, the native VLAN may be changed to any existing VLAN ID

```
8100fl(config)# vlan 100
8100fl(config-vlan-100)# interface lag 1
8100fl(config-interface-lag1)# switchport mode trunk
8100fl(config-interface-lag1)# switchport trunk-vlans 10,20,30
8100fl(config-interface-lag1)# switchport trunk-native-vlan 100
```

To cause a trunk mode port to reject all untagged traffic, remove support for native VLAN:

```
8100fl(config-interface-<id>)# switchport trunk-native-vlan disallow
```

Rev 6.11

Student Guide: 3–25

20

When a LAG or port is defined as a trunk-mode port, the port can carry tagged traffic for the VLANs explicitly associated with it. The trunk-mode port also can carry untagged traffic. The 8100fl automatically assigns all untagged traffic to the VLAN that is defined as “native” to the trunk. By default, the “native VLAN” for trunk-mode ports is VLAN 1, the Default VLAN.

To determine which VLAN should be designated the native VLAN, examine the requirements of the switch on the other side of the trunk-mode port. Typically, the device’s management address is in the address range associated with the native VLAN. Many designers dedicate VLAN 1 to management because it can simplify initial configuration and management of edge switches. Since every link comes up as an untagged member of VLAN 1, it is often practical to assign a management address to this VLAN. However, many designers prefer to dedicate a VLAN ID other than 1 to device management.

If a VLAN trunk will be an untagged member of VLAN 1, the trunk-mode port configuration is complete as soon as you associate the *trunk-vlans*.

If untagged traffic should be associated with a different VLAN, you must define that VLAN ID as the trunk mode port’s native VLAN. To override the default native VLAN, enter *switchport trunk-native-vlan <id>* in the port or LAG interface context.

In cases where a trunk port link should not carry any untagged traffic, you can “disallow” untagged traffic. Note, however, that a port configured in this way will reject all untagged traffic, including BPDUs and LLDP messages.

Define VLAN interfaces and display status

Define VLAN interfaces and display status



To enable the 8100fl to forward IP traffic to and from these VLANs, define IP addresses within VLAN interface configuration context:

```
8100fl(config)# interface vlan 10
8100fl(config-interface-vlan10)# ip address 10.1.10.1/24
8100fl(config-interface-vlan10)# interface vlan 20
8100fl(config-interface-vlan20)# ip address 10.1.20.1/24
8100fl(config-interface-vlan20)# interface vlan 30
8100fl(config-interface-vlan30)# ip address 10.1.30.1/24
8100fl(config-interface-vlan30)# interface vlan 40
8100fl(config-interface-vlan40)# ip address 10.1.40.1/24
```

The IP interface table shows the status of local IP interfaces:

```
8100fl# show ip interface brief
```

Interface	IP-Address	Status	Protocol
Vlan1	10.1.1.1/24	Up	Up
Vlan10	10.1.10.1/24	Up	Up
Vlan20	10.1.20.1/24	Up	Up
Vlan30	10.1.30.1/24	Up	Up
Vlan40	10.1.40.1/24	Up	Up

Rev 6.11

Student Guide: 3-26

21

To complete the configuration, define IP interfaces for all user-defined VLANs, as shown above. Note that the process for assigning addresses to user-defined VLANs is the same as for VLAN 1.

The output of *show ip interface brief* indicates that interfaces for all VLANs are “up.” The 8100fl is now prepared to act as default gateway for hosts in all VLANs and for the edge switches with addresses in VLAN 1.

Display IP interface and route tables

Display IP interface and route tables



```
8100fl#show ip route
Codes: R - RIP derived, O - OSPF derived, C - connected,
       S - static,
       * - candidate default route, IA - OSPF inter area route,
       E1 - OSPF external type 1 route, E2 - OSPF external type 2 route,
       N1 - OSPF NSSA external type 1 route,
       N2 - OSPF NSSA external type 2 route
       K - Kernel route remnant after rosrd restart
       A - Aggregate route

Gateway of last resort is not set

10.0.0.0/24 is subnetted, 5 subnets
C      10.1.1.0 is directly connected, Vlan1
C      10.1.10.0 is directly connected, Vlan10
C      10.1.20.0 is directly connected, Vlan20
C      10.1.30.0 is directly connected, Vlan30
C      10.1.40.0 is directly connected, Vlan40

Number of Routes: 5
```

Rev 6.11

Student Guide: 3-27

22

The 8100fl's IP route table stores entries for every network the router has learned from all sources. In this example, the 8100fl is performing only local routing because no dynamic routing protocols have been enabled and no static routes defined.

All of the networks shown in the IP route table are associated with VLANs. You can define up to 1,024 VLANs, but only 180 of them are routable. The 8100fl forwards traffic it receives through its VLAN interfaces if the traffic is destined for another network that appears in the route table, up to a maximum of 180 VLAN interfaces. The others are isolated broadcast domains unless you have provided external router(s) to forward traffic on behalf of clients in the VLANs

Module 5 will describe the process for enabling dynamic routing protocols such as OSPF and RIP.

Module 3 summary

Module 3 summary



In this module, you learned how to configure link aggregation and VLANs on the ProCurve Switch 8100fl

Topics included:

- Similarities and differences with other ProCurve switches
- Configuration of LAGs
- Configuration of VLANs and VLAN interfaces
- Definition and configuration of trunk mode and access mode switch ports
- Definition of IP interfaces associated with VLANs

Rev 6.11

Student Guide: 3–28

23

Module 3 described the processes for configuring link aggregation and VLANs on the 8100fl.

While the 8100fl is completely compatible in these areas with other ProCurve switches, it uses different terminology and command instructions. For instance, on the 8100fl, the term “trunk” refers to a mode that enables a port to carry tagged traffic for multiple VLANs. On other ProCurve switches, “trunk” refers to a group of aggregated ports, which is called a “Link Aggregation Group” on the 8100fl.

When configuring ports or LAGs on the 8100fl for VLAN membership, you must designate them for trunk mode or access mode. In access mode, the port or LAG will carry untagged traffic for a single VLAN. In trunk mode, the interface can carry traffic for multiple tagged VLANs and one untagged VLAN. On other ProCurve switches, ports are designated as tagged or untagged VLAN members.

Learning check

Module 3

1. How is the term “trunk” used differently on the ProCurve Switch 8100fl than on other ProCurve switches?

.....
.....

2. You have defined VLAN 10 and VLAN 30 on a ProCurve Switch 8100fl. What step is necessary to make interface 2/2 a tagged member of both VLANs?

.....
.....

3. What steps are necessary to define a VLAN-based IP interface on the 8100fl?

1.
2.

4. You have defined six IP interfaces on a ProCurve Switch 8100fl? What is necessary to enable the switch to route among them?

.....
.....

Module 3 lab overview

Module 3 lab overview



Module 3 lab tasks:

- Configure link aggregation on 8100fl and 5300xl
- Configure user VLANs on 8100fl and 5300xl
- Configure trunk-mode ports on 8100fl
- Configure VLAN IP interfaces on 8100fl

Rev 6.11

Student Guide: 3–31

25

In this lab activity, you will implement link aggregation and configure VLAN-based IP interfaces on the 8100fl. You will also configure 8100fl ports for trunk mode, enabling them to carry tagged traffic for multiple VLANs.

When the activity is complete, the lab will support four user VLANs. Student A in each group will configure VLAN 10 and VLAN 30. Student B will configure VLAN 20 and VLAN 40. The link connecting the 8100fl switches in each group will carry all four VLANs.

Consult the *Lab Activity Guide* for diagrams and instructions.

Provisioning Network Redundancy on the Switch 8100fl

Module 4

Objectives

After completing this module and its accompanying hands-on activities, you will be able to:

- Given a set of customer requirements, design redundancy solutions that combine MSTP and VRRP
- Describe 8100fl support for VRRP
- Describe 8100fl support MSTP
- Given a design and customer requirements, configure MSTP on the 8100fl and ProCurve Switch Intelligent Edge Switches
- Configure VRRP on the 8100fl
- Monitor, confirm, and troubleshoot VRRP and MSTP configuration

8100fl support for Spanning Tree and VRRP

8100fl support for Spanning Tree and VRRP



The 8100fl supports the following IEEE Spanning Tree standards:

- **IEEE 802.1Q 2003 version** – Multiple Spanning Tree Protocol (MSTP)
 - Formerly defined in IEEE 802.1s
- **IEEE 802.1D 2004 version** – Rapid Spanning Tree Protocol (RSTP)
 - Formerly defined in IEEE 802.1w
- The 8100fl default Spanning Tree version is MSTP
 - Automatically detects and interoperates with RSTP and legacy STP, as specified in standards documents

The 8100fl supports the Virtual Router Redundancy Protocol (VRRP)

- Defined in RFC 2338
- Each IP interface can support up to 15 VRRP instances

Rev 6.11

Student Guide: 4–2

3

In keeping with its role as an interconnect fabric, the ProCurve Switch 8100fl offers extensive support for several technologies that enable it to provide high levels of resiliency for user traffic. The 8100fl supports or interoperates with all standard versions of Spanning Tree to provide rapid failover for failed links. It also supports the Virtual Router Redundancy Protocol (VRRP) to provide default gateway redundancy and seamless failover in the event of core router failure.

The default Spanning Tree version on the 8100fl is Multiple Spanning Tree Protocol (MSTP), as defined in IEEE 802.1Q 2003, supporting up to 16 MST instances. Because its support is standards-based, the 8100fl automatically detects and interoperates with devices using RSTP and legacy STP. In the absence of user-defined MST instances, the 8100fl exhibits RSTP behavior.

When the 8100fl must provide default gateway service for connected hosts, it can support up to 15 VRRP instances per IP interface. Any or all of the 15 available Virtual Router IDs can be used on each IP interface. Because the 8100fl supports up to 180 VLAN or port-based IP interfaces, the maximum number of VRIDs supported is 2,700. However, the actual number of VRRP instances to be configured in a given environment will depend on edge switch and host requirements.

The rest of this module will discuss the design and configuration of 8100fl solutions using these technologies.

MSTP/VRRP redundancy solution

MSTP/VRRP redundancy solution



MSTP/VRRP redundancy solution

- **Benefits**
- **Topology example**
- **Active path affects VRRP role**
- **Edge switch capabilities determine MSTP configuration**

Configure and monitor MSTP

Configure and monitor VRRP

Rev 6.11

Student Guide: 4–3

4

To provide high levels of redundancy while maintaining high levels of resource utilization, the 8100fl supports a solution that combines MSTP and VRRP. The first section of Module 4 will present the high-level topology and benefits of this solution. Later sections will describe steps and tools for configuration and monitoring.

Benefits of combining MSTP and VRRP

Benefits of combining MSTP and VRRP



A solution that combines MSTP and VRRP:

- Results in better resource utilization than one that combines RSTP and VRRP
- Allows MSTP-compliant edge switch with redundant uplinks to utilize all uplinks
- Supports definition of a different set of active links and associated VLANs for each user-defined MST instance
- Allows a pair of 8100fl switches to share default gateway responsibilities

To realize full benefit, most edge switches must support MSTP

Rev 6.11

Student Guide: 4-4

5

In a high-availability environment, redundant links are often provisioned along with redundant router interfaces to ensure that users have at least one alternate path to the backup default gateway. These redundant links also ensure the VRRP Backup can continue to receive the Master's VRRP advertisements after a link failure.

While switches using RSTP define a single loop-free path through a switched domain by blocking redundant links, MSTP allows the VLANs in a Spanning Tree domain to be associated with different Spanning Tree instances. Each instance has a different Root Bridge and, consequently, a different active path.

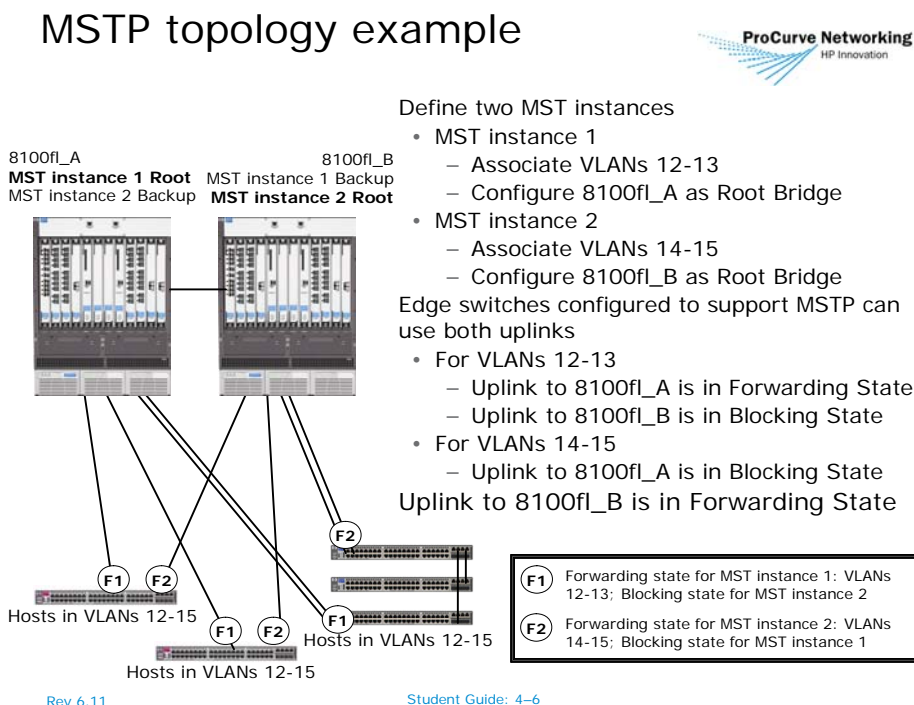
When VRRP is operating in a single Spanning Tree environment based on RSTP, traffic efficiency dictates that the RSTP Root Bridge also be the Owner of all VRRP instances defined within the Spanning Tree domain. This can result in poor resource utilization because many redundant links must be blocked to avoid loops.

By combining MSTP and VRRP, you can ensure better utilization. If configured properly, this solution ensures that all redundant links are used and that the VRRP routers can share default gateway duties.

Under MSTP, a different Root Bridge is elected for each MST instance, which serves an administratively defined set of VLANs. When VRRP is enabled, you can configure each switch to ensure that the MST Root Bridge for a given instance also is the VRRP Owner for all VRIDs associated with the VLANs in the instance.

While this module will focus on the configuration of the 8100fl, the full benefits of this combined solution can be achieved only when edge switches support MSTP. All ProCurve Intelligent Edge Switches support MSTP.

MSTP topology example



To achieve the benefits associated with MSTP, create an MST instance for each redundant path and associate a unique set of VLANs with each instance. In the example, VLANs 12-13 are associated with MST instance 1, and VLANs 14-15 are associated with MST instance 2. Because each edge switch has two uplinks, one for each 8100fl, the administrator has defined two MST instances. The number of VLANs that may be associated with an MST instance is limited only by the number of VLANs supported by the switch.

To ensure optimal redundancy and link utilization, each 8100fl should act as Root Bridge of one MST instance and Backup Root for the other. To achieve this goal, MSTP enables you to define independent Bridge Priority, Port Priority, and Path Cost for each MST instance. In the example, Bridge Priority settings on the 8100fl switches cause 8100fl_A to be the Root Bridge of MST instance 1 and 8100fl_B to be the Root Bridge of MST instance 2. Additionally, each 8100fl switch is configured to be the Backup Root of the instance for which it is not the Root Bridge.

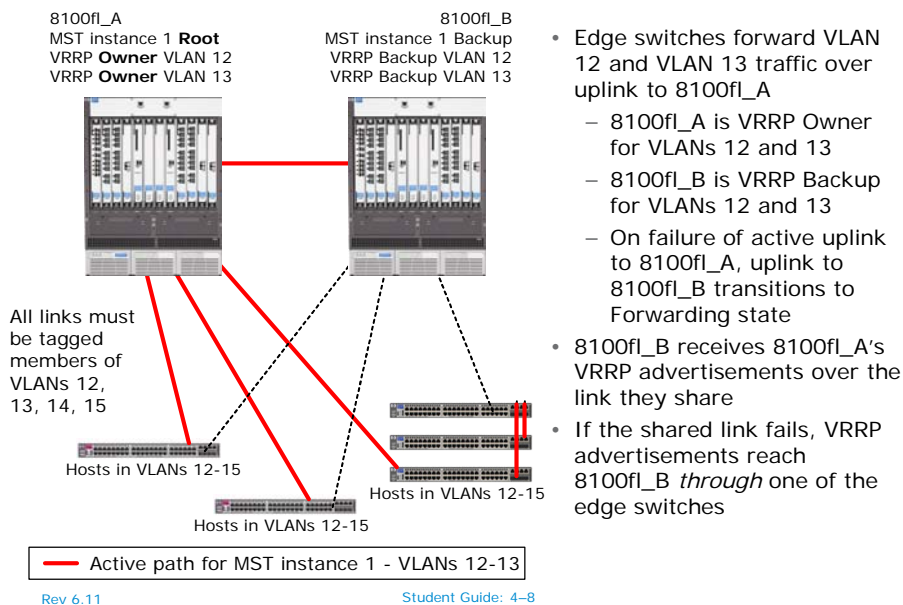
On the edge switches, Bridge Priority is left at default settings for both MST instances, ensuring that the 8100fl switches have the highest priorities in the domain. If either of the 8100fl switches should fail, the remaining switch would become Root of both MST instances.

The edge switches in this example are crucial to this design because the primary benefits of MSTP are realized at the edge. The establishment of a different Root Bridge for each MST instance produces a unique set of links that are placed in Forwarding state. Each link from the edge to the core is placed in Forwarding state for one MST instance and in Blocking state for the other. This enables the edge switches to use both of their uplinks. From the perspective of the edge switches, one uplink carries traffic associated with VLANs 12-13 and the other uplink carries traffic relating to VLANs 14-15.

The next two pages will describe how the active paths for both MST instances affect the assignment of VRRP roles to the core switches.

MST instance 1: Active path affects VRRP role

MST instance 1: Active path affects VRRP role



The definition of MST instances enables the redundant 8100fl switches to share default gateway responsibilities efficiently. As shown above, the active path in MST instance 1 makes the 8100fl_A the best candidate to be VRRP Owner/Master for VLANs 12 and 13.

If all links are active, the uplink from each edge switch to 8100fl_A will remain in the Forwarding state. Traffic in VLANs 12 and 13 typically will not transit 8100fl_B, which is the Backup Root of instance 1.

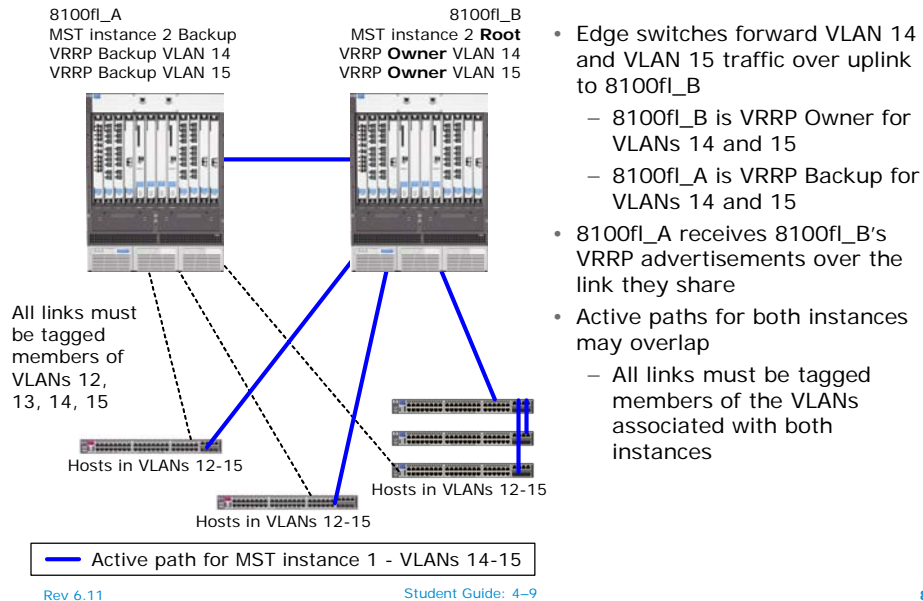
If an uplink between an edge switch and 8100fl_A fails, the uplink to 8100fl_B will transition to the Forwarding state. In that case, 8100fl_B will forward traffic at Layer 2 over its direct link to 8100fl_A, which will continue acting as default gateway for hosts in VLANs 12 and 13.

In the absence of a link failure, 8100fl_B receives 8100fl_A's VRRP advertisements for VLANs 12 and 13 over its Root port, which provides the direct link between the 8100fl switches. **If the link shared by the 8100fl switches should fail, one of the links to an edge switch will become the primary path for VRRP advertisements.** The actual selection of port will depend on the priority settings or MAC addresses on the edge switches.

Because any link could become part of the path from 8100fl_B to 8100fl_A, all switch-to-switch links must be tagged members of all VLANs. This ensures that the Master's VRRP advertisements can reach the Backup.

MST instance 2: Active path affects VRRP role

MST instance 2: Active path affects VRRP role



- Edge switches forward VLAN 14 and VLAN 15 traffic over uplink to 8100fl_B
 - 8100fl_B is VRRP Owner for VLANs 14 and 15
 - 8100fl_A is VRRP Backup for VLANs 14 and 15
- 8100fl_A receives 8100fl_B's VRRP advertisements over the link they share
- Active paths for both instances may overlap
 - All links must be tagged members of the VLANs associated with both instances

The active path for MST instance 2 causes the edge switches to place all uplinks to 8100fl_B into the Forwarding state for VLANs associated with the instance. Consequently, 8100fl_B is the best candidate to be the primary default gateway for hosts in VLANs 14 and 15.


During VRRP configuration, 8100fl_B will be designated as Owner of the VRIDs associated with VLANs 14 and 15. 8100fl_A will be the Backup router. As in MST instance 1, the Bridge Priorities assigned to the 8100fl switches cause the link between them to remain active. This link is the primary path through which 8100fl_A (the Backup) receives VRRP advertisements from 8100fl_B. If the link shared by the 8100fl switches should fail, one of the links to an edge switch will become the primary path for VRRP advertisements.

Because any of the nine links shown in the diagram could become part of the primary path between the 8100fl switches, each link needs to carry all four user VLANs.

Under normal conditions, each edge switch's uplink to 8100fl_B will be its Root port for the MST instance. If the link fails, the edge switch will use its link to 8100fl_A as Root port.

Common Spanning Tree and IST instance

Common Spanning Tree and IST instance



Rev 6.11

Student Guide: 4-10

9

Common Spanning Tree (CST)

- Defines loop-free path that interconnects all STP-enabled switches (MSTP, RSTP, and legacy STP)

- Interconnects MST regions

Internal Spanning Tree (IST) instance

- System-defined default instance within each MST switch
- Associated with VLANs not mapped to MST instances
- Active topology independent of user-defined MST instances

Each region's IST active path is part of the Common Spanning Tree

The Common Spanning Tree (CST) defines a loop-free connection among all STP-enabled devices within a switched domain, including those that support RSTP and legacy STP. In this way, the CST enables MST bridges to interoperate with switches that do not support MSTP.

The Internal Spanning Tree (IST) instance is a default instance within each MSTP bridge. Any VLANs not explicitly associated with user-defined MST instances are mapped to the IST instance. Some MST-enabled ProCurve switches allow VLAN 1 to be mapped to any instance. However, in the 8100fl MSTP implementation, VLAN 1 always belongs to the IST instance.

Each MST Bridge's IST instance connects it to the CST, and the IST instance within an MST region defines the active path of the CST through the region. Consequently, the CST and IST sometimes are referred to collectively as the Common and Internal Spanning Tree (CIST).

The switch labeled 8100fl_A is the Root of the IST instance for this MST region because the Bridge Priority assigned to the IST instance has the lowest numeric value in the region, giving it the highest priority.


Because the entire Spanning Tree domain consists of only the seven switches shown in the diagram, 8100fl_A is also elected the Root of the Common Spanning Tree. If this MST region were part of a larger Spanning Tree domain, another, more centrally located switch might be elected Root of the CST. Such a domain could include RSTP and legacy STP switches, as well as multiple MST regions.

The CST makes it possible for users in the same VLAN to remain in the same broadcast domain even if they are connected to edge switches that support different Spanning Tree versions. If an edge switch that supports RSTP has edge ports assigned to VLANs 12-15, the connected clients would be able to access the same resources as users in VLANs 12-15 that are connected to MST-enabled edge switches.

When designing an MST solution, the careful assignment of VLANs to MST instances enables MSTP edge switches to make full use of their redundant links while also ensuring interoperability with devices that do not support MSTP. The only functional difference between the MSTP and RSTP edge switches will be their utilization of redundant links.

Edge switch capabilities drive MSTP configuration

Edge switch capabilities drive 8100fl MSTP configuration



If edge switches do not support MSTP

- User-defined MST instances unnecessary
- All VLANs remain in IST instance

Edge switch uplinks:

- Forwarding State for uplink to Root Bridge
- Blocking State for uplink to Backup Root

8100fl_A:

- Root of IST instance
- VRRP Owner/Master for all VLAN interfaces

8100fl_B

- Backup root of IST instance
- Becomes VRRP Master only if VRRP Owner becomes unavailable

Rev 6.11

Student Guide: 4-12

10

In the example, stacks of ProCurve 2524 switches have redundant uplinks to two 8100fl switches. Hosts connected to several user VLANs require redundant default gateway support from the core switches. Because the edge switches in this environment support only a single Spanning Tree instance, there would be no benefit in defining MST instances on the 8100fl switches. Without user-defined MST instances, the 8100fl assigns all VLANs to the IST instance.

In this example, 8100fl_A is defined as the Root Bridge of the IST instance. This causes edge switch uplinks to 8100fl_A to be in the Forwarding state, making this switch the best candidate to be Owner of all VRRP instances. If 8100fl_A becomes unavailable, 8100fl_B will become the Root Bridge of the IST instance. The uplinks from edge switches to this switch will transition from Blocking state to Forwarding state. Although 8100fl_B is not defined as the VRRP Owner of any VLANs, it will become the Master if the configured Owner becomes unavailable.

Configure and monitor MSTP

Configure and monitor MSTP



✓ *MSTP/VRRP redundancy solution*

Configure and monitor MSTP

- **Define VLANs and configure ports**
- **Configure IST parameters**
- **Configure MST parameters**
- **Display Spanning Tree statistics**

Configure and monitor VRRP

Rev 6.11

Student Guide: 4–13

11

The next section of Module 4 will describe the procedures for configuring MSTP on the 8100fl.

Spanning Tree configuration overview

Spanning Tree configuration overview



Configuration procedures that apply equally to MSTP and RSTP applications

- Define VLANs, configure switchport modes, and associate VLANs
- Define Internal Spanning Tree instance parameters at global configuration context, for example:
 - IST Bridge Priority
 - IST Port Priority

• ProCurve 8100fl: Enable Spanning Tree per port or LAG

- ProCurve edge switches: Enable Spanning Tree box-wide

Configuration procedures that apply only to MSTP applications

- Define parameters that must match other MSTP switches in the region:
 - Configuration name and configuration revision
 - MST instance IDs and member VLANs
- Define parameters that are typically unique for each MSTP switch
 - Bridge Priority per instance
 - Port Priorities per instance

Rev 6.11

Student Guide: 4–14

12

By default, Spanning Tree is disabled on all ProCurve switches. Because Spanning Tree is most useful in environments where the topology includes loops within VLANs, Spanning Tree cannot be meaningfully configured until after VLANs are defined and associated with ports. Module 3 describes procedures for configuring VLANs on the 8100fl and on ProCurve Intelligent Edge Switches.

Basic Spanning Tree configuration parameters

Some basic Spanning Tree configuration steps must be performed regardless of whether the switches will support multiple MST instances or will support only the IST. On the 8100fl, you must enable Spanning Tree on each port that will be a part of the Spanning Tree domain. On the Intelligent Edge Switches, Spanning Tree is enabled in the global configuration context.

Before enabling Spanning Tree, it is advisable to define Bridge Priority on switches that should be elected Root and Backup Root of the CST. This will avoid unnecessary link-state transitions that are required if the Root and Backup Root are changed after the switches have already converged.

As mentioned earlier, these steps will complete configuration of the 8100fl if it is to be deployed in an RSTP-only environment.

MSTP configuration parameters

A group of switches that have the same MST configuration parameters are part of the same region. Each switch can belong to only one MST Region. While no single configuration parameter identifies a switch as a member of particular region, switches with the following parameters in common will automatically become members of the same region:

- Configuration name
- Configuration revision
- MST instances and mapped VLANs

To achieve the full load-sharing benefit of MSTP, it should be enabled on edge switches as well as core switches such as the 8100fl. The links that interconnect the switches in a region must carry a common set of VLANs. The BPDUs that circulate among the switches in an MST region are generated by the region's IST Root and follow the active path of the IST instance. The switches in the region determine active paths per instance based on the premise that all of the links carry all of the VLANs. If a link that carries a subset of the VLANs should become part of the active path for an instance associated with a VLAN not present on the link, the VLAN's broadcast domain will become fragmented, which will disrupt communication.

As described earlier, MSTP enables you to define distinct active paths for multiple MST instances. To take advantage of this opportunity, identify the best candidate for Root Bridge within each MST instance, and then assign Bridge Priority and/or Port Priority values accordingly. If the topology includes two acceptable candidates for Root Bridge, you can configure each switch to be Root of one instance. If the topology includes only one acceptable Root candidate, you can adjust Port Priority values to cause each instance to have a different active path.

The next several pages will present a process for configuring a group of ProCurve switches to enable a redundant topology similar to the one presented earlier in this module.

Define VLANs and configure ports

Define VLANs and configure ports




Define VLANs

```
8100fl(config)# vlan 12
8100fl(config-vlan-12)# vlan 13
8100fl(config-vlan-13)# vlan 14
8100fl(config-vlan-14)# vlan 15
```

Associate VLANs with ports and LAG

```
8100fl(config-vlan-14)# interface tengig 4/1
8100fl(config-interface-10gig4/1)# switchport mode trunk
8100fl(config-interface-10gig4/1)# switchport trunk-vlans 12-15
8100fl(config-interface-10gig4/1)# interface tengig 8/1
8100fl(config-interface-10gig8/1)# switchport mode trunk
8100fl(config-interface-10gig8/1)# switchport trunk-vlans 12-15
8100fl(config-interface-10gig8/1)# interface lag 1
8100fl(config-interface-lag1)# switchport mode trunk
8100fl(config-interface-lag1)# switchport trunk-vlans 12-15
```



Rev 6.11

Student Guide: 4-16

13

In this example, the administrator of an 8100fl begins by defining four user VLANs that will be carried over all switch-to-switch links. The administrator then associates all relevant switch ports with the VLANs.

A similar procedure is required on the 3400cl and 2848 edge switches. The 3400cl has edge ports in VLANs 12, 13, and 14. Although it does not have edge ports in VLAN 15, its 10-GbE uplinks, which are ports 49 and 50, must be configured as tagged members of all four user VLANs. This is because unexpected link failures elsewhere might cause one or both of its uplinks to become part of the active path for all VLANs.

```
3400cl(config)# vlan 12 untagged 1-16
3400cl(config)# vlan 12 tagged 49-50
3400cl(config)# vlan 13 untagged 17-32
3400cl(config)# vlan 13 tagged 49-50
3400cl(config)# vlan 14 untagged 33-48
3400cl(config)# vlan 14 tagged 49-50
3400cl(config)# vlan 15 tagged 49-50
```

IST configuration example

IST configuration example



Define Bridge Priority for IST instance:


```
8100fl_A(config)#spanning-tree priority 0
8100fl_B(config)#spanning-tree priority 1
```

On 8100fl, enable Spanning Tree within configuration context of each port that should be part of the Spanning Tree domain

```
8100fl(config)#interface lag 1
8100fl(config-interface-lag1)#spanning-tree enable
8100fl(config-interface-lag1)#interface tengig 4/1
8100fl(config-interface-10gig4/1)#spanning-tree enable
8100fl(config-interface-10gig4/1)#interface tengig 8/1
8100fl(config-interface-10gig8/1)#spanning-tree enable
```

On edge switches enable Spanning Tree within global configuration context

```
3400cl(config)#spanning-tree
```



Rev 6.11

Student Guide: 4-17

14

To enable an 8100fl to participate in a single-instance Spanning Tree, it is necessary only to enable Spanning Tree within the configuration context of each interface, as shown above. For a LAG, Spanning Tree is enabled in the LAG interface configuration context and applies to all port members of the LAG. No special RSTP configuration is required on the 8100fl because the switch automatically will use RSTP until and unless MST instances are configured.

On the 8100fl, the IST instance is created as soon as Spanning Tree is enabled on any port. The switch begins participating in the CST immediately, communicating with other switches that have any version of Spanning Tree enabled. All existing VLANs are associated with the IST instance and remain there unless they are explicitly associated with an MST instance. If the switches are configured with a dedicated management VLAN, it is usually advisable to leave the VLAN in the IST instance.

On the ProCurve Intelligent Edge Switches, Spanning Tree is enabled with a single command (*spanning-tree*) entered within the global configuration context.

In this example, the 8100fl switches will be the Root and Backup Root bridges in the IST instance. On the 8100fl and all ProCurve Intelligent Edge Switches, Bridge Priority is configured as a step value between 0 and 15, where the actual priority is the configured value multiplied by 4096.

In the example, the administrator configures 8100fl_A to have the highest priority by assigning it the lowest numeric value in the domain. Another switch should be given the second highest priority, which will cause it to assume the role of Root in the event 8100fl_A fails. To configure 8100fl_B as the Backup Root, assign a Bridge Priority value that is higher than 0, but lower than 8, which is the default Spanning Tree Bridge Priority.

Other parameters on the 8100fl that can affect the active path within the IST instance include Port Priority and Path Cost.

View Spanning Tree statistics on CIST Root

View Spanning Tree statistics on CIST Root




MSTP is the Spanning Tree version even if MST instances have not been defined

```
8100fl_A# show spanning-tree
```

```
Force Version      : 802.1s (MSTP)
Bridge ID         : 0:001321fefaff
Ports In Bridge    : 3
Max Age           : 20 secs
Hello Time        : 2 secs
Forward Delay      : 15 secs
Topology Changes   : 64
Last Topology Chg: 0 days 0 hours 2 min 41 secs ago

Max Hops          : 20
CST Root          : 0:001321fefaff
CST Root Port     : This switch is root
CST Root Path Cost : 0
IST Regional Root : 0:001321fefaff
IST Regional Root Path Cost : 0
IST Remaining Hops : 20
```

Port	Priority	Cost	State	Role	Edge	LinkType
TenGig4/1	128	2000	Forwarding			
TenGig8/1	128	2000	Forwarding			
Lag1	128	20000	Forwarding			



Rev 6.11

Student Guide: 4-19

15

The output of *show spanning-tree* indicates that 8100fl_A is the Root of the Common Spanning Tree. Instead of providing a port ID in the CST Root Port field, the output states that “This switch is root.” The output also indicates that the switch’s Bridge ID matches the ID in the “CST Root” field.

This switch also is the regional root of the IST instance. In this example, the IST is identical to the CST because all switches support MSTP and belong to the same MST region.

To display Spanning Tree information on the 8100fl, enter the *show spanning-tree* command, as shown, which is identical to the command used on Intelligent Edge Switches. On the 8100fl, the output of this command does not display the state of inactive ports or ports for which Spanning Tree is not enabled. On the Intelligent Edge Switches, the output shows all ports because Spanning Tree is enabled box-wide.

Even though no MST instances have been defined on the 8100fl, the *show spanning-tree* output indicates that MSTP is the enabled Spanning Tree version. The 8100fl supports a *force-version* command, as specified in standards documents, but the command is not required to enable the 8100fl to function as an RSTP switch. The *force-version* command is useful if hosts in the domain support applications that are sensitive to the RSTP/MSTP rapid transitions. In such a case, you may need to force the switch to exhibit the slower convergence behavior that is a characteristic of legacy Spanning Tree.

View Spanning Tree statistics on CIST Backup Root

View Spanning Tree statistics on CIST Backup Root




Root Port is tengig 8/1 and priority is 4096

```
8100fl_B# show spanning-tree

Force Version      : 802.1s (MSTP)
Bridge ID         : 4096:001321fe5aff
Ports In Bridge   : 3
Max Age           : 20 secs
Hello Time        : 2 secs
Forward Delay     : 15 secs
Topology Changes  : 6
Last Topology Chg: 0 days 0 hours 3 min 19 secs ago

Max Hops          : 20
CST Root         : 0:001321fefaff
CST Root Port    : TenGig8/1
CST Root Path Cost : 20000
IST Regional Root : 4096:001321fe5aff
IST Regional Root Path Cost : 0
IST Remaining Hops : 20
```

Port	Priority	Cost	State	Role	Edge	LinkType
TenGig4/1	128	2000	Forwarding			
TenGig8/1	128	2000	Forwarding			
Lag1	128	20000	Forwarding			



Rev 6.11

Student Guide: 4-20

16

The output from the *show spanning-tree* command on 8100fl_B indicates that the actual priority is 4096, which is the step value configured with the *spanning-tree priority 1* command entered earlier. The Bridge ID is the actual priority combined with the MAC address. Because the Bridge Priority for other switches in this example has been left at the default value, 8100fl_B has the second highest priority.

As shown above, the Backup Root Bridge does not block any ports, and its Root Port is TenGig8/1.

Define parameters common to switches in MST region

Define parameters common to switches in MST region




Define common MST parameters on 8100fl:

```
8100fl(config)#spanning-tree config-name lab
8100fl(config)#spanning-tree config-revision 1
8100fl(config)#spanning-tree instance 1
8100fl(config-mst)#member-vlans 12-13
8100fl(config-mst)#spanning-tree instance 2
8100fl(config-mst)#member-vlans 14-15
```

Enable MSTP and configure common parameters on 3400cl:

```
3400cl(config)#spanning-tree protocol-version mstp
3400cl(config)#write memory
3400cl(config)#reload
...
3400cl#configure
3400cl(config)#spanning-tree config-name lab
3400cl(config)#spanning-tree config-revision 1
3400cl(config)#spanning-tree instance 1 vlan 12 13
3400cl(config)#spanning-tree instance 2 vlan 14 15
```



Rev 6.11

Student Guide: 4-21

17

By definition, an MST region is a group of switches with three identical MST parameters:

1. A common string assigned as “config-name.” By default, each switch uses its own MAC address as its config-name.
2. A common numeric value assigned as the “config-revision.” By default, the config-revision is 0.
3. A common set of VLAN-to-MST instance mappings

These parameters must be configured identically on all switches that your network design specifies as part of same region. If any value is different on a particular switch, it will not join the MST region you intended but instead its unique parameters will define a separate region. If this happens, the MST instances defined on the switch will be logically disconnected from the instances defined on other switches, eliminating the MSTP load-sharing benefits described earlier.

8100fl configuration

On the 8100fl, the *spanning-tree instance <id>* command leads to a separate configuration context, where you associate VLANs with the instance using the *member-vlans* command. In the example, VLANs are assigned as a range but you can also specify a list of VLANs separated by commas. For example, to assign the VLANs 10, 20, and 30 to an instance, you would issue the command *member-vlans 10,20,30* within the correct instance configuration context.

Intelligent Edge switch configuration

Two groups of ProCurve Intelligent Edge Switches approach MSTP configuration differently. The default Spanning Tree version on 3500zl, 5400zl, and 6200yl switches is MSTP. Like the 8100fl, these switches operate as RSTP switches unless MST instances are defined. RSTP is the default Spanning Tree version on the 2600, 2800, 3400cl, 5300xl, and 6400cl. To enable MSTP on these switches, you must change the Spanning Tree protocol version and reboot the switch.

All of the Intelligent Edge Switches use the same command to create instances and map VLANs to them. In the example, the administrator assigns VLANs 12 and 13 to MST instance 1 on the 3400cl using the command *spanning-tree instance 1 vlan 12 13*. The Intelligent Edge Switches do not have a separate configuration context for MST instances.

Common configuration issues

On the 8100fl and edge switches, all VLANs initially belong to the IST when Spanning Tree is enabled. When a VLAN is assigned to an MST instance, the VLAN is removed from the IST, but all unassigned VLANs remain in the IST indefinitely. In the 8100fl MSTP implementation, the Default VLAN (VLAN 1) always belongs to the IST instance.


On both switch platforms, VLANs must currently be associated with the IST instance in order to be mapped to a user-defined MST instance. You cannot move a VLAN from one user-defined instance directly to another. For example, to move VLAN 12 from MST instance 1 to MST instance 2, you would first remove VLAN 12 from MST instance 1 by entering the *no* version of the command that associated the VLAN with the instance. On both platforms, the command would return VLAN 12 to the IST. The VLAN could then be associated with MST instance 2.

Configure Bridge Priority for each instance

Configure Bridge Priority for each instance



```
8100fl_A(config)#spanning-tree instance 1
8100fl_A(config-mst)#spanning-tree priority 0
8100fl_A(config-mst)#spanning-tree instance 2
8100fl_A(config-mst)#spanning-tree priority 1
```



```
8100fl_B(config)#spanning-tree instance 1
8100fl_B(config-mst)#spanning-tree priority 1
8100fl_B(config-mst)#spanning-tree instance 2
8100fl_B(config-mst)#spanning-tree priority 0
```

	IST instance	MST instance 1	MST instance 2
8100fl_A	Root	Root	Backup Root
8100fl_B	Backup Root	Backup Root	Root

Rev 6.11

Student Guide: 4-23

18

To realize the load-sharing benefits of MSTP, you must specify a different Bridge Priority for each instance. If the Bridge Priority values are left at default settings within the MST instances, the Root of all instances will be the switch with the lowest MAC address, resulting in an identical active path for all instances.

In the example, 8100fl_A is configured to be the Root of MST instance 1 and the Backup Root of MST instance 2. Conversely, 8100fl_B is configured to be Root of MST instance 2 and Backup of MST instance 1. 8100fl_A is also Root Bridge for the CIST.

Defining the Root Bridge role is the surest way to affect the active path. However, the roles of the switches in each MST instance could be reversed without diminishing the benefits of MSTP.

Parameters at port or LAG interface configuration context

Parameters at port or LAG interface configuration context



Port-level parameters:

```
8100fl_A(config-interface-lag1)#spanning-tree ?
edge-port      - Specify port to transition to STP forwarding state
                  immediately(RSTP/MSTP only)
enable         - Start Default Spanning Tree on this interface
hello-time     - Hello interval to use when this switch is MSTP CIST
                  root
instance       - MSTP instance
mcheck        - Force port to send RSTP BPDUs to detect dot1d switches
path-cost      - Default Spanning Tree cost on this (logical) port
point-to-point-mac - Specify a link type for STP use(RSTP/MSTP only)
priority       - Set port priority (the value is in range of 0-240
                  divided into steps of 16 that are numbered from 0 to
                  15, default is step 8).
```

Example: Define any ports connected to end nodes as *edge-ports*

```
8100fl_A(config-interface-gig1/10)#spanning-tree edge-port
```

Rev 6.11

Student Guide: 4-24

19

Many Spanning Tree parameters can be defined at the port configuration context.

For instance, as shown, you can apply the *edge-port* setting to any port that is connected to an end node, such as a server. This setting will apply to all MST instances and will force the port to transition to Forwarding state immediately.

Other parameters, such as port priority and path cost can be configured for specific instances.

View Spanning Tree statistics: Root of MST instance 1


View Spanning Tree statistics: Root of MST instance 1

```
8100fl_A#show spanning-tree instance 1
```

```
Force Version      : 802.1s (MSTP)
Bridge ID          : 0:001321fefaff
Ports In Bridge    : 3
Max Age            : 20 secs
Hello Time         : 2 secs
Forward Delay      : 15 secs
Topology Changes   : 29
Last Topology Chg: 0 days 0 hours 1 min 26 secs ago
```

```
Mapped VLANs      : 12 13
Regional Root      : 0:001321fefaff
Regional Root Port : This switch is root
Regional Root Path Cost: 0
Remaining Hops     : 20
```

Port	Priority	Cost	State	Role
TenGig1/1	128	2000	Forwarding	
TenGig1/2	128	2000	Forwarding	
Lag1	128	20000	Forwarding	



Rev 6.11

Student Guide: 4-25

20

To view details for a user-defined MST instance, include the instance ID in the *show spanning-tree* command, as shown. The example shows statistics for MST instance 1 on 8100fl_A, which is the instance's Root Bridge because of priorities configured earlier. The output also indicates that VLANs 12 and 13 are mapped to the instance.

The switch's Bridge ID field indicates a priority of 0, which is the Bridge Priority in the IST instance. As Root of the IST instance, this switch generates BPDUs for all switches in the MST region. These BPDUs contain information about all instances in the region and their VLAN mappings and are transmitted along the active path of the IST instance.

This behavior is in contrast to many PVST implementations, where the root of each VLAN/instance generates tagged BPDUs.

View Spanning Tree statistics: Root of MST instance 2

View Spanning Tree statistics:
Root of MST instance 2

```

8100fl_B# show spanning-tree instance 2


Force Version      : 802.1s (MSTP)
Bridge ID          : 4096:001321fe5aff
Ports In Bridge    : 3
Max Age            : 20 secs
Hello Time         : 2 secs
Forward Delay      : 15 secs
Topology Changes   : 3
Last Topology Chg: 0 days 0 hours 1 min 39 secs ago

Mapped VLANs       : 14 15
Regional Root      : 0:001321fe5aff
Regional Root Port : This switch is root
Regional Root Path Cost: 0
Remaining Hops     : 20

Port      Priority  Cost      State      Role
Tengig4/1 128       2000      Forwarding
Tengig8/1 128       2000      Forwarding
Lag1      128       20000     Forwarding

```

8100fl_B's
CIST Bridge Priority



Rev 6.11

Student Guide: 4-26

21

The output of *show spanning-tree* for MST instance 2 on the 8100fl_B indicates that the switch is the Root of Instance 2 and displays the VLAN IDs that are mapped to this instance. 8100fl_B is the “Regional Root” of MST instance 2 because it was configured with a priority of 0 for this instance. However, its Bridge ID indicates that its priority is 4096, which is the step value for the switch’s configured IST Bridge Priority of 1.

Troubleshoot MST configuration

Troubleshoot MST configuration



The following must be identical for all MST switches in a region:

- Instance-to-VLAN mappings
- Configuration name
- Configuration revision

MST configuration digest

- Hash product based on configuration name and configuration revision
- Must be identical or each switch will be in its own region

Symptom of mismatched digest

- Multiple switches claim to be Root of the same MST instance

Verify MST configuration digest for 8100fl and Intelligent Edge switches:

```
8100fl1#show spanning-tree mst-config
```

```
MST Configuration Name      : lab
MST Configuration Revision  : 1
MST Configuration Digest    : 0xCA136A235706B316C8DB8F921067A68F
```

Instance ID	Mapped VLANs
1	12 13
2	14 15

Rev 6.11

Student Guide: 4-27

22

Many difficulties with MST configurations arise because of mismatches among the MST configuration parameters on the switches that were expected to form a single region. As mentioned earlier, all switches in a region must be configured with identical VLAN-to-MST instance mappings, configuration names, and configuration revisions. If all of these items do not match, the switches will not become members of the same region.

Each BPDU generated by the Root of the IST instance includes a hashed value that is a digest of the common MST configuration parameters. If more than one switch is reporting itself to be Root of the same instance, examine the configuration digests on all switches to ensure all these parameters match.

As shown, the *show spanning-tree mst-config* command provides a convenient view of all MST configuration parameters. The identical command is used on ProCurve Intelligent Edge Switches.

Configure and monitor VRRP

Configure and monitor VRRP



✓ *MSTP/VRRP redundancy solution*

✓ *Configure and monitor MSTP*

Configure and monitor VRRP

- **Plan for VRRP load sharing**
- **Assign group address**
- **Monitor VRRP roles**

Rev 6.11


Student Guide: 4–28

23

The final section of Module 4 will describe the planning, configuration, and monitoring required to implement VRRP in the combined MSTP/VRRP solution.

Automatic default gateway failover with VRRP

Automatic default gateway failover with VRRP



8100fl switches in the same MST region can share default gateway responsibilities

- 8100fl_A
 - Root Bridge of MST instance mapped to VLANs 12 and 13
 - Root Bridge of IST instance mapped to VLAN 1
 - VRRP Owner for VRIDs mapped to VLANs 1, 12, and 13
- 8100fl_B
 - Root Bridge of MST instance mapped to VLANs 14 and 15
 - VRRP Owner for VRIDs mapped to VLANs 14 and 15

Rev 6.11

Student Guide: 4–29

24

By combining VRRP and MSTP, you can increase network availability by providing redundancy at Layer 2 with redundant links and at Layer 3 with redundant default gateways. To receive the maximum benefit from this configuration, the switch that is the Root Bridge of an MST instance also should be the VRRP Owner of the virtual IP addresses associated with the VLANs mapped to the instance.

If you implement RSTP and VRRP, the division of VRRP Owner duties between the routers will lead to inefficient forwarding patterns. Because some ports will be blocked by Spanning Tree, hosts who use the Backup Root as their default gateway will only be able to communicate with the router indirectly, through the Root bridge. An analogous pattern emerges if you implement MSTP while configuring a single router to be Owner of all VRIDs.

In the example above, 8100fl_A is Root of MST instance 1 and VRRP Owner of the virtual IP addresses associated with VLANs 12 and 13. 8100fl_B is Root for MST instance 2 and VRRP Owner of the virtual IP addresses associated with VLANs 14 and 15. As the diagram shows, the active path in this topology ensures that hosts in all VLANs can communicate with their default gateways without going through the other 8100fl.


In the event of link failure, the combined MSTP and VRRP configuration will ensure that network services remain available. If one of the 8100fl switches fails, the remaining switch will become the Root Bridge role for the IST instance and for both MST instances. The same 8100fl will also become VRRP Master for all VLANs.

Plan for VRRP load sharing

Plan for VRRP load sharing



VLAN ID	Host IP address range	Default Gateway address	MST Instance / Root Bridge
1	10.1.1.0/24	10.1.1.1	IST / 8100fl_A
12	10.1.12.0/24	10.1.12.1	1 / 8100fl_A
13	10.1.13.0/24	10.1.13.1	1 / 8100fl_A
14	10.1.14.0/24	10.1.14.1	2 / 8100fl_B
15	10.1.15.0/24	10.1.15.1	2 / 8100fl_B



As shown above, each router must have an interface in every VLAN in order for the routers to share default gateway responsibilities for the user VLANs. On the 8100fl, the IP address used as default gateway for the hosts in a given VLAN is referred to as the “group” address. Some routers refer to this as the “virtual” IP address. In VRRP terms, the “Owner” of the virtual or group address is the router interface whose actual IP address matches that of the group address.

In a topology that combines MSTP and VRRP, you can achieve the most efficient traffic forwarding patterns if the switch that performs the Root Bridge role for a given VLAN is also defined as the VRRP Owner for that VLAN.

The table above describes the information required for VRRP configuration:

- The address range associated with each VLAN
- The default gateway address that is used by hosts in each VLAN
- Which switch functions as the Root Bridge for each VLAN

By coordinating this information, you can readily determine which router should be the VRRP Owner—and therefore the primary default gateway—for each VRID.

The example illustrates one possibility for assigning router interfaces to meet the needs of this topology. All IP addresses are shown along with the role, either Master or Backup, that each 8100fl interface will fill as long as its state is up. Owner addresses are shown in bold.

Because 8100fl_A is configured as the Root of the IST instance and MST instance 1, it is configured as the Owner of the virtual or group address used by hosts in VLANs 1, 12, and 13. Because 8100fl_B is configured as the Root of MST instance 2, it is configured as the Owner of the virtual or group address used by hosts in VLANs 14 and 15.

8100fl VRRP configuration overview

8100fl VRRP configuration overview



VRRP is enabled within IP interface configuration context

- Command syntax:

```
8100fl(config-interface-<id>)# vrrp <vrid> ip <group-address>
```

- Valid values for *<vrid>* are between 1 and 15
- *<group-address>* is IP hosts' default gateway address
- VRRP interface state is not defined with explicit Master/Backup command:
 - Owner/Master (priority 255) if actual IP interface address is the same as *<group-address>*
 - Non-owner/Backup (default priority: 100) if actual IP interface address is different from *<group-address>*

VRRP is not enabled globally

- No 'router vrrp' command

Rev 6.11

Student Guide: 4–32

26

On the 8100fl, VRRP configuration requires only that you enable VRRP within the configuration context of interfaces that will participate in a VRID. This task is performed with a single command, *vrrp <vrid> ip <group-address>*, where *<group-address>* is the IP address that hosts will use as their default gateway.

It is not necessary to enable VRRP in the global configuration context. It is also not necessary to specify the Owner or Backup status of the router for each VRID. The 8100fl will become Owner of the VRID if the address of its IP interface is the same as the *group-address* assigned to the VRID.

The next few pages will illustrate VRRP configuration for VLAN-based interfaces. However, VRRP can also be configured for port-based interfaces.

VRRP configuration example: Configure VRIDs on 8100fl_A

VRRP configuration example: Configure VRIDs on 8100fl_A



Verify existing IP interface addresses

```
8100fl_A# show ip interface brief
```

Interface	IP-Address	Status	Protocol
Vlan1	10.1.1.1/24	Up	Up
Vlan12	10.1.12.1/24	Up	Up
Vlan13	10.1.13.1/24	Up	Up
Vlan14	10.1.14.2/24	Up	Up
Vlan15	10.1.15.2/24	Up	Up

Define VRIDs within each VLAN configuration context

```
8100fl_A(config)# interface vlan 1
8100fl_A(config-interface-vlan1)# vrrp 1 ip 10.1.1.1
8100fl_A(config-interface-vlan1)# interface vlan 12
8100fl_A(config-interface-vlan12)# vrrp 2 ip 10.1.12.1
8100fl_A(config-interface-vlan12)# interface vlan 13
8100fl_A(config-interface-vlan13)# vrrp 3 ip 10.1.13.1
8100fl_A(config-interface-vlan13)# interface vlan 14
8100fl_A(config-interface-vlan14)# vrrp 4 ip 10.1.14.1
8100fl_A(config-interface-vlan14)# interface vlan 15
8100fl_A(config-interface-vlan15)# vrrp 5 ip 10.1.15.1
```

Group-address same
as corresponding
IP interface address

Rev 6.11

Student Guide: 4-33

27

As noted earlier, configuring an actual IP address that is identical to the group address causes 8100fl_A to be the Owner of VRIDs 1, 2, and 3, which are associated with VLANs 1, 12, and 13 respectively. The Owner status gives the highest possible priority setting (255) to these interfaces. As long as the Owner interfaces are up, they will be Master of these VRIDs.

This router's actual IP address does not match the group address for VRIDs 4 and 5, which are mapped to VLANs 14 and 15 respectively. The default priority for non-Owner VRRP interfaces is 100. 8100fl_A could become the Master of VRIDs 4 and 5 only if no other VRRP interface with a higher priority (such as an Owner at priority level 255) is present on the networks associated with VLANs 14 and 15.

If this is the first switch you are configuring, and you haven't yet configured 8100fl_B, this switch will have the Master role of all five VRIDs. If the other switch has been configured as Owner of VRIDs 4 and 5, 8100fl_A will have the Backup role for VRIDs 4 and 5.

The example uses a unique VRID for each VLAN in order to make them easily distinguishable in diagrams and text. On the 8100fl, however, it is not necessary to associate a unique VRID with each IP interface. If desired, you can use any of the 15 available VRIDs for every IP interface.

VRRP configuration example: Configure VRIDs on 8100fl_B

VRRP configuration example: Configure VRIDs on 8100fl_B



Verify existing IP interface addresses

```
8100fl_B# show ip interface brief
```

Interface	IP-Address	Status	Protocol
Vlan1	10.1.1.2/24	Up	Up
Vlan12	10.1.12.2/24	Up	Up
Vlan13	10.1.13.2/24	Up	Up
Vlan14	10.1.14.1/24	Up	Up
Vlan15	10.1.15.1/24	Up	Up

Define VRIDs within each VLAN configuration context

```
8100fl_B(config)# interface vlan 1
8100fl_B(config-interface-vlan1)# vrrp 1 ip 10.1.1.1
8100fl_B(config-interface-vlan1)# interface vlan 12
8100fl_B(config-interface-vlan12)# vrrp 2 ip 10.1.12.1
8100fl_B(config-interface-vlan12)# interface vlan 13
8100fl_B(config-interface-vlan13)# vrrp 3 ip 10.1.13.1
8100fl_B(config-interface-vlan13)# interface vlan 14
8100fl_B(config-interface-vlan14)# vrrp 4 ip 10.1.14.1
8100fl_B(config-interface-vlan14)# interface vlan 15
8100fl_B(config-interface-vlan15)# vrrp 5 ip 10.1.15.1
```

Group-address same
as corresponding
IP interface address

VRID configuration uses the **same** commands for both switches

- Router derives its VRRP role by comparing actual and group addresses

Rev 6.11

Student Guide: 4-34

28

In this example, 8100fl_B is configured to participate in the same VRIDs earlier configured on 8100fl_A. However, 8100fl_B is configured as the Owner of VRIDs 4 and 5, which are associated with VLANs 14 and 15. If these VRRP interfaces are up, this router will perform the Master role for these VRIDs.

The table below shows the final VRRP configurations for both switches for all VRIDs.

VRID and VLAN assignments

VRID	VLAN	Group Address	Master	Backup
1	1	10.1.1.1	8100fl_A	8100fl_B
2	12	10.1.12.1	8100fl_A	8100fl_B
3	13	10.1.13.1	8100fl_A	8100fl_B
4	14	10.1.14.1	8100fl_B	8100fl_A
5	15	10.1.15.1	8100fl_B	8100fl_A

Note that VLANs associated with VRID 2 and VRID 3 are also associated with MST instance 1. The VLANs associated with VRID 4 and VRID 5 are associated with MST instance 2. VLAN 1 is associated with VRID 1 and with the IST instance.

View VRRP status

View VRRP status



```
8100fl_A# show vrrp
...
Interface Vlan12 - Group 2
-----
Uptime                0 days, 2 hours, 11 minutes, 2 seconds.
State                  Master
Priority                255 (default value)
Virtual MAC address    00005E:000102
Advertise Interval     1000 msec(s) (default value)
Preempt Mode           enabled delay = 0 sec
Master Down Interval   3000
Authentication         None (default value)
Primary Address        10.1.12.1
Associated Addresses   10.1.12.1
...
Interface Vlan14 - Group 4
-----
Uptime                0 days, 2 hours, 3 minutes, 47 seconds.
State                  Backup
Priority                100 (default value)
Virtual MAC address    00005E:000104
Advertise Interval     1000 msec(s) (default value)
Preempt Mode           enabled delay = 0 sec
Master Down Interval   3000
Authentication         None (default value)
Primary Address        10.1.14.2
Associated Addresses   10.1.14.1
...
```

Rev 6.11

Student Guide: 4–35

29

The *show vrrp* command indicates whether the routers in each VRID are correctly performing their Master and Backup roles.

The State for every router in each VRID should be either Master or Backup, except for a brief initialization period, when the State will be “Init.” If the “Init” state persists on either Master or Backup, check the configuration and verify interface state. If an interface is down, the VRRP state cannot move past “Init.”

If the output for two routers displays a state of “Master” for the same VRID, check to see if the IP interface and group address are correctly configured on both routers.

Module 4 summary

Module 4 summary



In this module, you learned how to design and configure a redundancy solution that combines MSTP and VRRP on the 8100fl

Topics included:

- Design considerations to ensure VRRP and MSTP interoperate efficiently
- How to configure RSTP and MSTP on the 8100fl and Intelligent Edge switches
- How to configure VRRP on the 8100fl and Intelligent Edge switches
- How to verify and troubleshoot MSTP and VRRP configurations

Rev 6.11

Student Guide: 4–36

30

To perform its role as interconnect fabric, the 8100fl supports industry standards for providing redundant links and default gateways. Module 4 presented a design that used VRRP and MSTP to provide for optimal usage of network resources while also providing redundancy. The module also described steps for configuring the design parameters on the 8100fl and on the ProCurve Switch 3400cl.

Learning check

Module 4

1. What is the benefit of combining MSTP and VRRP in a network that features redundant 8100fl switches?
.....
.....
2. What configuration items must match among all switches in an MSTP region?
 - a.
 - b.
 - c.
3. In what context is VRRP enabled on the 8100fl?
.....
.....
4. What is necessary to configure an 8100fl to be Master of a VRID?
.....
.....

Module 4 lab overview

Module 4 lab overview



Module 4 lab tasks:

- Enable Spanning Tree on 8100fl and 5300xl
- Configure Multiple Spanning Tree on 8100fl and 5300xl
- Configure VRRP on 8100fl
- Test VRRP failover

Rev 6.11

Student Guide: 4–39

32

In the Module 4 lab activity, you and your partner will implement the combined MSTP/VRRP solution described in the module.

You will begin by enabling and verifying single Spanning Tree and then will configure Multiple Spanning Tree on the 5300xl and the 8100fl.

After verifying MST operation, you will configure VRRP on the 8100fl. Along with your partner, you will test VRRP failover.

Configuring IP Routing on the Switch 8100fl

Module 5

Objectives

After completing this module and its accompanying hands-on activities, you will be able to:

- Describe the IP routing features and capabilities of the ProCurve Switch 8100fl
- Define IP interfaces, including loopback, port-based, and VLAN interfaces, on the 8100fl
- Define static routes on the 8100fl
- Configure RIP on the 8100fl
- Configure OSPF on the 8100fl
- Configure route redistribution on the 8100fl
- Interpret the 8100fl route table
- Given a set of customer requirements, design a routing solution using the 8100fl as an interconnect fabric
- Monitor and troubleshoot 8100fl routing configurations

IP routing on the 8100fl

IP routing on the 8100fl



The 8100fl automatically routes among locally connected networks defined by any of the following interface types:

- VLAN interfaces
 - Associate multiple switch ports with a single broadcast domain
 - 8100fl routes traffic to and from hosts in the networks defined by the VLAN interface
- Port-based interfaces
 - Associate an IP address directly with a switch port
 - Each port is in a separate IP network and broadcast domain
- Loopback interface
 - IP interface whose state is not bound to any physical port state

To route traffic toward remote networks, the 8100fl must have one or more of the following enabled:

- RIP
- OSPF
- Static routes

Rev 6.11

Student Guide: 5–2

3

The 8100fl automatically routes among the local networks that are defined by its IP interfaces, including VLAN interfaces, port-based interfaces, and the loopback interface. In keeping with its role in the network core, the 8100fl supports RIP, OSPF, and static routes.

As with other 8100fl features, the routing capabilities are designed specifically to enable the 8100fl to fulfill its role as an enterprise interconnect fabric. The actual design and configuration of an 8100fl routing solution will depend largely upon the capabilities of the edge switches that the 8100fl interconnects.

Port-based and loopback interfaces

A port-based interface associates an IP address directly with a switch port. In this configuration, a broadcast domain may include only one port. The port does not carry any tagged traffic and consequently can be a member of only one broadcast domain.

Like other routers, the 8100fl supports the configuration of a loopback interface, which is an interface that is not bound to any port state. Because a loopback interface is not connected to any network media, it is often configured with a 32-bit mask. The 8100fl derives an address range based on the loopback interface's IP address and mask and places that range in its IP route table.

Routing toward remote networks

As well as routing traffic among its local networks, the 8100fl can forward traffic toward remote networks that appear in its route table. When configured to support dynamic routing protocols, the 8100fl can exchange route information with other routers using RIP and OSPF. The 8100fl can also learn about remote networks through user-defined static routes.

Define port-based IP interfaces

Define port-based interfaces



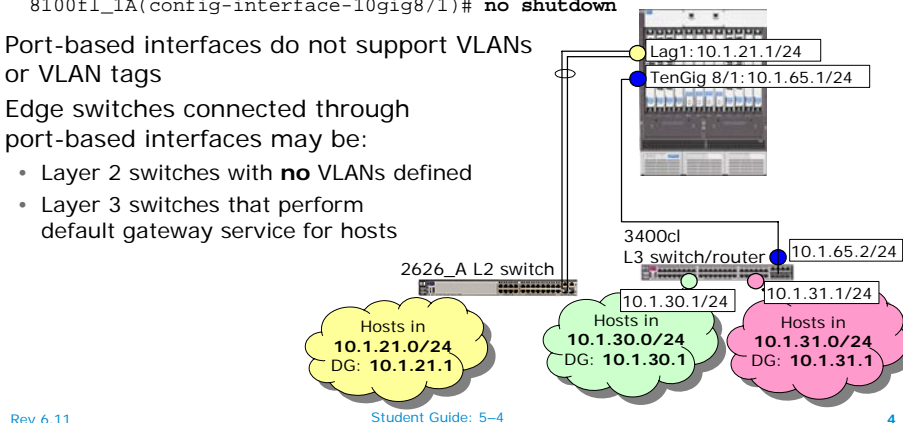
The 8100fl can be configured as a router that assigns IP addresses directly to a LAG or physical port

```
8100fl_1A(config)# interface lag 1
8100fl_1A(config-interface-lag1)# ip address 10.1.21.1/24
8100fl_1A(config-interface-lag1)# int tengig 8/1
8100fl_1A(config-interface-10gig8/1)# ip address 10.1.65.1/24
8100fl_1A(config-interface-10gig8/1)# no shutdown
```

Port-based interfaces do not support VLANs or VLAN tags

Edge switches connected through port-based interfaces may be:

- Layer 2 switches with **no** VLANs defined
- Layer 3 switches that perform default gateway service for hosts



Rev 6.11

Student Guide: 5-4

4

In this example, an administrator assigns port-based IP interfaces to a LAG and to a gigabit port. In one case, the 8100fl will provide default gateway services to hosts connected to a ProCurve Switch 2626. In the other case, the 8100fl supports a 10-GbE uplink from a 3400cl that provides default gateway services to connected hosts.

Supporting a Layer 2 switch

LAG1 includes GigabitEthernet 1/1 and GigabitEthernet 1/2 and connects directly to a Layer 2 edge switch. All of the connected hosts are in the network 10.1.21.0/24. The aggregated uplink that leads to the 8100fl is in the same broadcast domain as the edge ports that lead to the end stations. In this topology, the 8100fl will act as default gateway for all hosts in the network 10.1.21.0/24. Accordingly, the LAG is configured with the IP address 10.1.21.1/24, which will be the address of the hosts' default gateway.

When a port or LAG on the 8100fl is associated directly with an IP address, no other ports on the 8100fl can be in the same IP address range. If you attempt to assign an IP address to a port and the address is within the range associated with another interface, the CLI returns an error indicating that it cannot validate the address and mask. In the example, this means edge switches connected to other 8100fl switch ports cannot have hosts in the address range 10.1.21.0/24. If this address range were required to support hosts connected to multiple ports on the 8100fl, the administrator would configure a VLAN interface, using the process described earlier.

Supporting an edge router

In the second example, port TenGig 8/1 is configured with the IP address 10.1.65.1/24. The connected edge switch, a 3400cl, routes IP traffic on behalf of hosts in networks 10.1.30.0/24 and 10.1.31.0/24. To perform this duty, the 3400cl forwards traffic to the 8100fl over a routed uplink, a network that includes only the router interfaces on the 8100fl and 3400cl.

No VLAN tags

Port-based interfaces are only appropriate in situations where the links will never carry tagged traffic. For instance, the topologies shown above could not support Layer 2 prioritization using the IEEE 802.1p priority marker included in the IEEE 802.1Q VLAN tag. If the edge switches needed to forward prioritized traffic, they could be configured for Layer 3 QoS.

Impact of port-based interfaces on Default VLAN

Impact of port-based interfaces on Default VLAN



When IP addresses are assigned directly to ports, those ports are removed from the VLAN 1 and placed in a separate Layer 2 broadcast domain

```
8100fl_1A# show vlan
VLAN      Name           Status      Ports           Type
1          Default        active      Gig1/3
           Gig1/4
           Gig1/5
           Gig1/6
           Gig1/7
           Gig1/8
           Gig1/9
           Gig1/10

...[output omitted]

4096       VLAN-Lag1      active      Lag1
4097       VLAN-TenGig8/1 active      TenGig8/1

8100fl_1A#
```

Rev 6.11

Student Guide: 5-6

5

As with the 5300xl and other ProCurve switches, every port on the 8100fl must belong to a VLAN. Consequently, when you define a port-based interface on the 8100fl, the port is removed from the Default VLAN (VLAN 1) and placed into a separate VLAN with a system-defined ID.

The VLAN ID is greater than 4095, which is the greatest value that can be represented in the 12-bit VLAN ID field in a VLAN tag. The non-standard VLAN ID does not raise interoperability issues because these VLAN IDs are locally significant. The “invalid” VLAN ID values never appear in tags because port-based interfaces do not use tags.

This behavior is quite different from other ProCurve switches. For instance, when you define a port-based interface on the 9300m and 9400sl routing switches, the port retains its membership in the Default VLAN. This can create mismatched Layer 2 and Layer 3 broadcast domains, which is why ProCurve typically recommends that interfaces on those switches be placed within VLANs.

The ProCurve Intelligent Edge Switches do not support port-based interfaces, but the same effect can be accomplished by defining a VLAN and VLAN interface with only one port member.

View IP interfaces and route table

View IP interfaces and route table



To view status of configured IP interfaces:

```
8100fl_1A# show ip interface brief
Interface      IP-Address      Status  Protocol
Lag1           10.1.21.1/24    Up      Up
TenGig1/10     10.1.65.1/24    Up      Up
```


Without routing protocols, only local networks appear in route table

```
8100fl_1A# show ip route
Codes: R - RIP derived, O - OSPF derived, C - connected,
       S - static,
       * - candidate default route, IA - OSPF inter area route,
       E1 - OSPF external type 1 route, E2 - OSPF external type 2 route,
       N1 - OSPF NSSA external type 1 route,
       N2 - OSPF NSSA external type 2 route
       K - Kernel route remnant after rosrd restart
       A - Aggregate route
```

Gateway of last resort is not set

```
10.0.0.0/24 is subnetted, 2 subnets
C    10.1.21.0 is directly connected, Lag1
C    10.1.65.0 is directly connected, TenGig1/10

Number of Routes: 2
```



Rev 6.11

Student Guide: 5-7

6

To view the status of configured IP interfaces, issue the *show ip interface* command with the *brief* argument, as shown above. To see more detail about each interface, omit the *brief* argument. You can also limit the scope of the output by specifying an interface ID.

View the route table

As with other ProCurve products, you can view the route table of the 8100fl by issuing the *show ip route* command. In the output shown above, the route table does not include networks 10.1.30.0/24 and 10.1.31.0/24 because those networks are connected to the 3400cl. To enable the 8100fl to learn about those networks, the administrator could enable RIP or OSPF on both routers. The administrator could alternatively define a static route to the networks that shows the neighboring interface on the 3400cl as the next hop.

Steps for RIP configuration

Steps for RIP configuration



Enable RIP within global configuration context:

```
8100fl_1A(config)# router rip
```

Use *network* statement to specify interfaces that must send and receive RIP updates


```
8100fl_1A(config-router)# network 10.1.65.0 255.255.255.0
```

Define RIP version to be used for all interfaces:

```
8100fl_1A(config-router)# version 2
```

Define RIP version for a single interface:

```
8100fl_1A(config-router)# interface tengig 8/1
8100fl_1A(config-interface-10gig8/1)# ip rip version 2
```



Rev 6.11

Student Guide: 5-8

7

On the 8100fl, as on other ProCurve switches, you enable RIP by issuing the *router rip* command in the global configuration context. To enable RIP per interface, use a *network* statement within the RIP router configuration context to specify one or more address ranges that include all interface IP addresses that must send and receive RIP updates.

In the example, TenGig 8/1 is the only interface with a neighbor and is therefore the only interface that must participate in RIP communication. Use the RIP router configuration context command *network 10.1.65.0 255.255.255.0* to enable RIP on this interface.

The default RIP version on the 8100fl is version 1. Because RIP version 2 is the default on all ProCurve Intelligent Edge Switches, you must either configure all edge switches to use RIP version 1 or configure the core switch(es) to use RIP version 2. The configuration of the core switches is recommended because RIP version 2 is superior to version 1. To change the RIP version on all RIP interfaces, specify *version 2* within the RIP router configuration context.

To change the RIP version on a single interface, issue the *ip rip version 2* command in the interface configuration context. This will override the RIP version that is set in the RIP router configuration context. Note, however, that this command is not sufficient to enable RIP on the interface, as it is on ProCurve Intelligent Edge switches.

The *network* statement

The *network* statement



The *network* statement specifies a network address range

- IP interfaces whose addresses fall within *network* statement ranges will send and receive RIP updates
- Network addresses of IP interfaces that fall within *network* statement ranges are included in outbound RIP updates
- Split Horizon rules apply
 - RIP updates sent through each IP interface contain a unique set of advertisements
 - Networks learned through a given IP interface are not included in updates sent through that interface
- Use address ranges where possible to minimize the number of network statements, for example:
 - Assign to RIP all interfaces between 10.0.0.0 and 10.255.255.255:

```
8100f1(config-router)# network 10.0.0.0 255.0.0.0
```

Rev 6.11

Student Guide: 5–9


8

Unlike other ProCurve switches, the 8100f1 uses a *network* statement to determine which interfaces will participate in RIP updates. You may supply a separate *network* statement for each interface that should participate in RIP communication or you can supply one or more address ranges that include multiple IP interfaces. In the example, the statement *network 10.0.0.0 255.0.0.0* causes RIP to be enabled on every IP interface whose address falls within the specified range. The next page will describe a scenario with more detail about this range.

On the 8100f1, the syntax for the *network* statement requires that you enter the mask in dotted-decimal form. The CLI will not accept a *network* statement mask entered as a bit count, as in 10.1.64.0/22.

Network numbering strategies

Network numbering strategies



Rev 6.11

Student Guide: 5–10

9

“Non-edge” interfaces specified in *network* statement

- TenGig 4/1: ip address 10.1.64.1/24
- TenGig 7/1: ip address 10.1.65.1/24
- TenGig 8/1: ip address 10.1.66.1/24
- Gig 5/1: ip address 10.1.67.1/24

“Edge” interfaces **not** specified in network statements

- Redistribute connected, non-RIP interfaces
 - VLAN 12: ip address 10.1.12.1/24
 - VLAN 13: ip address 10.1.13.1/24
 - VLAN 14: ip address 10.1.14.1/24
 - VLAN 22: ip address 10.1.22.1/24
 - VLAN 23: ip address 10.1.23.1/24
 - VLAN 24: ip address 10.1.24.1/24

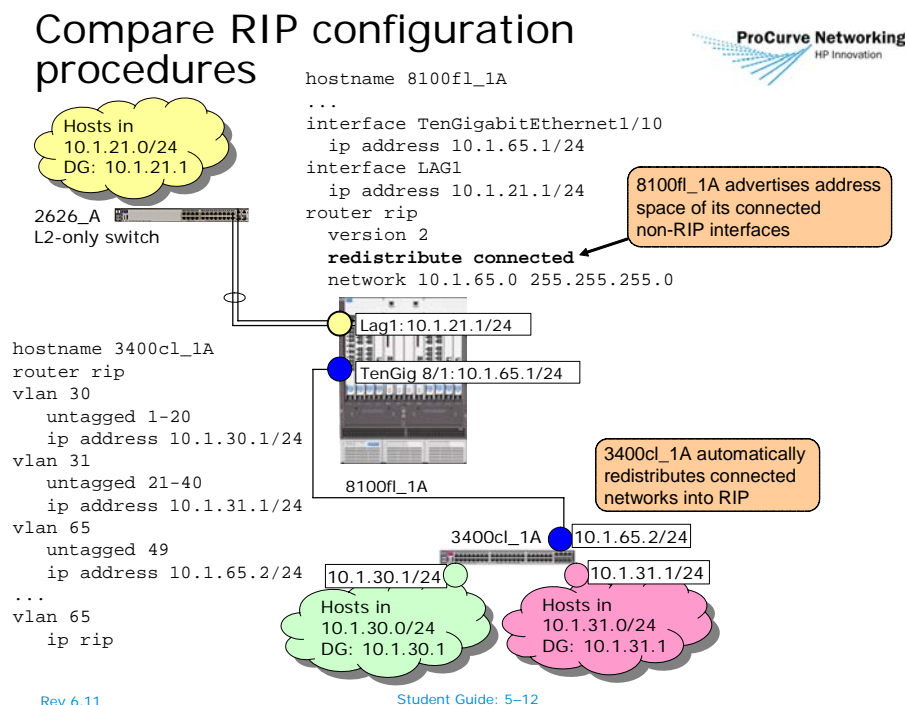
This example uses a larger number of interfaces than earlier examples to illustrate a network numbering scheme that can minimize the number of network statements. The router will have RIP neighbors on four interfaces whose associated network address space is contiguous. Because the four edge interfaces have the same value in the first 22 bits of their addresses, their address space can be summarized with the following statement: *network 10.1.64.0 255.255.252.0*. The table below illustrates how to derive the starting address and mask that includes these four networks. The third octet has been omitted from the binary value column because the last eight bits are host bits, as indicated by the 24-bit mask associated with each interface.

Dotted-decimal value	Binary value of first 24 bits		
10.1.64.0	00001010	00000001	010000 00
10.1.65.0	00001010	00000001	010000 01
10.1.66.0	00001010	00000001	010000 10
10.1.67.0	00001010	00000001	010000 11
Common bits set to “1”	11111111	11111111	111111 00

In addition to specifying which IP interfaces will participate in RIP communication, the *network* statement defines which networks will be advertised in RIP updates.

An administrator has also defined six VLAN interfaces on the 8100fl as edge ports. Because the router has no neighbors on the edge interfaces their address space does not need to be specified in *network* statements.

Compare RIP configuration procedures



This example continues the scenario started on the previous page. The administrator has configured the 8100fl and the 3400cl so that they can exchange information about their connected networks over the network 10.1.65.0/24.

The network statement specifies 10.1.65.0 255.255.255.0 because only one of 8100fl_1A's interfaces must participate in RIP. The RIP version is set at "2" for compatibility with 3400cl_1A.

Redistribution of connected networks

The 3400cl and other ProCurve Intelligent Edge Switches automatically redistribute connected, non-RIP networks into RIP. In the 3400cl configuration example, an administrator has associated RIP with VLAN 65, which is the VLAN associated with the network 10.1.65.0/24. RIP is not applied to VLANs 30 and 31. The address space associated these VLANs will be automatically redistributed into RIP. The 3400cl will include the address spaces associated with these VLANs in the RIP updates it sends to its neighbor, 8100fl_1A.

The 8100fl must be explicitly configured to redistribute connected, non-RIP networks. In this example, an administrator has entered the command *redistribute connected* within the RIP router configuration context. Consequently, the 8100fl will advertise the network 10.1.21.0/24 in updates sent to its neighbor, 3400cl_1A.

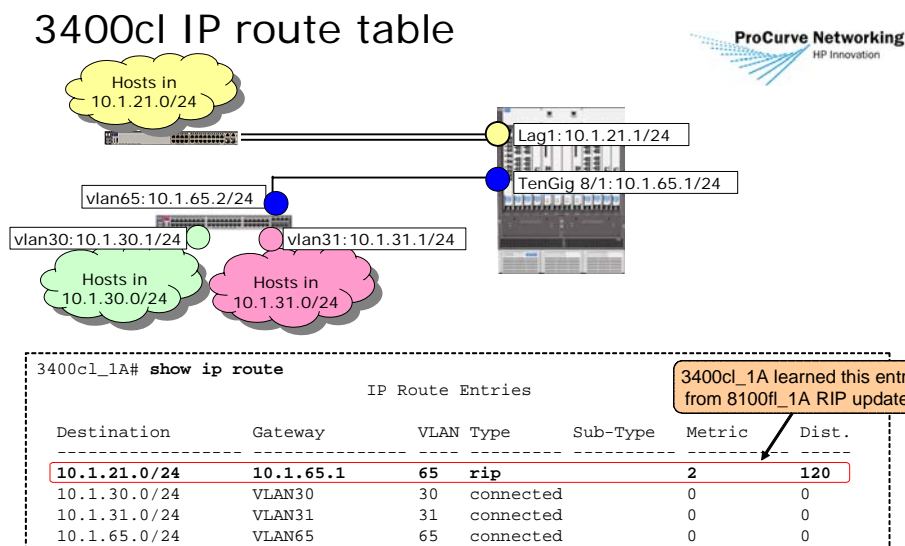
You can configure the 8100fl to redistribute some of its connected networks by specifying a *route-map* after the *redistribute connected* command. See the *8100fl Management and Configuration Guide* for more information on restricting redistributed RIP routes. To restrict the connected networks redistributed by the 3400cl, use the *restrict* command within the RIP router configuration context. See the *3400cl/5300xl Advanced Traffic Management Guide* for more information on restricting redistributed networks.

Redistribution of edge network space is the preferred method for both the 3400cl and the 8100fl. In addition to improving efficiency, redistribution is somewhat more secure than defining the edge networks as native RIP interfaces. If RIP is enabled on the edge networks, an unauthorized router connected to an edge network could inject information into the routed domain that disrupts connectivity.

On both switch platforms, an administrator could define the connected edge networks as native RIP interfaces to cause the associated address space to be included in RIP updates. However, this is not the most efficient approach because the router would continually send RIP updates through these networks.

Redistribution of the edge networks accomplishes the goal of address space inclusion without incurring unnecessary overhead.

3400cl IP route table



Rev 6.11

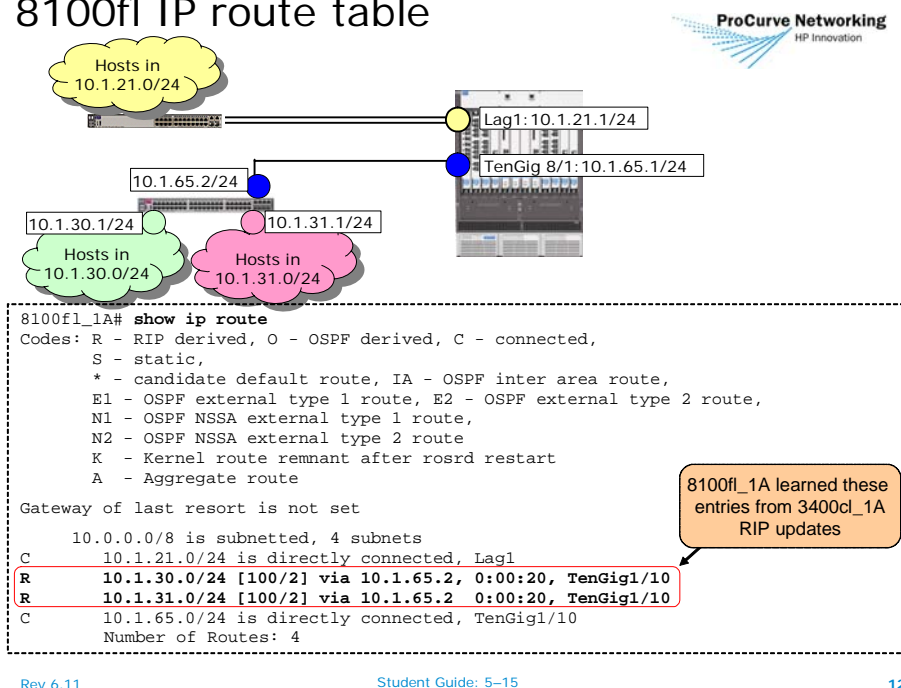
Student Guide: 5-14

11

After a connected network is redistributed into RIP on the 8100fl and the network is included in RIP updates, RIP neighbors recognize the network as a RIP route, even though it is not a native RIP interface. In the example above, the IP route table on the 3400cl indicates that it has learned of network 10.1.21.0/24 from a RIP update it received through its VLAN 65 interface. The networks associated with VLANs 30, 31, and 65 are reported as connected networks.

8100fl IP route table

8100fl IP route table



Rev 6.11

Student Guide: 5-15

12

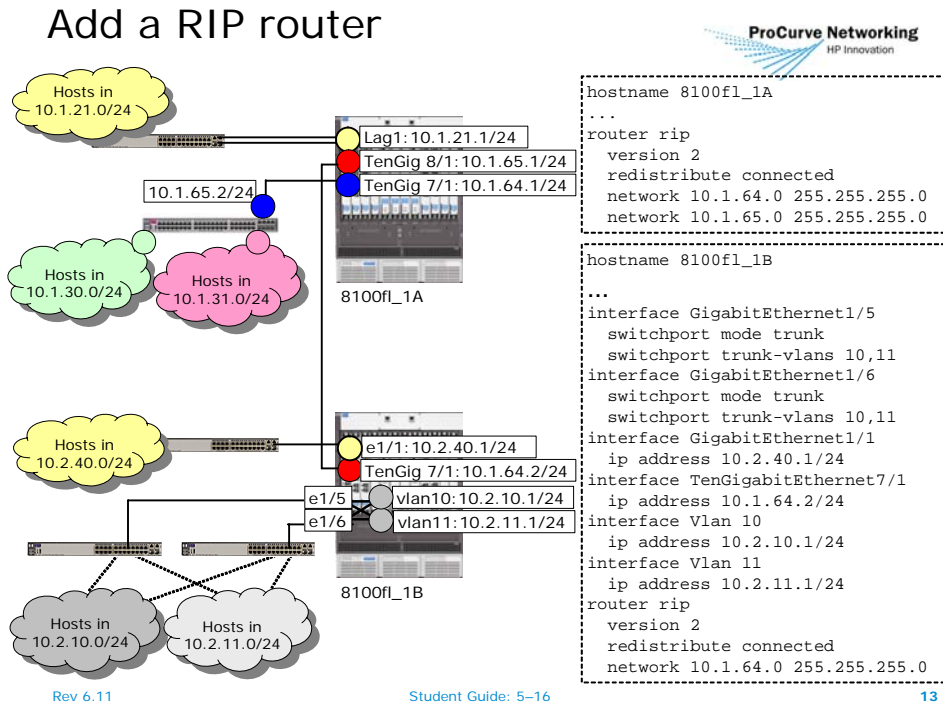
In this configuration, the IP route table of the 8100fl includes four entries, including two connected networks and two networks learned through RIP updates received from the 3400cl.

The bracketed values in the RIP entries—such as [100/2]—indicate the route's administrative distance and cost. The 8100fl assigns an administrative distance of 100 to all RIP routes. Because the 8100fl learned of networks 10.1.30.0/24 and 10.1.31.0/24 from the same neighbor, they have the same cost (2), and the same next hop (10.1.65.2).

Like other ProCurve switches, the 8100fl uses administrative distance to determine route preference. A lower administrative distance for a type of route, such as RIP or static, means the type has a higher route preference. When the router discovers multiple paths to the same destination using different routing protocols or methods, the route with the highest preference is placed in the IP route table.

The 8100fl's default administrative distance for RIP routes is 100. You can change the administrative distance for RIP routes within the RIP router configuration context using the *distance* command. This is the same command and configuration context you would use to change administrative distance on ProCurve Intelligent Edge Switches.

Add a RIP router

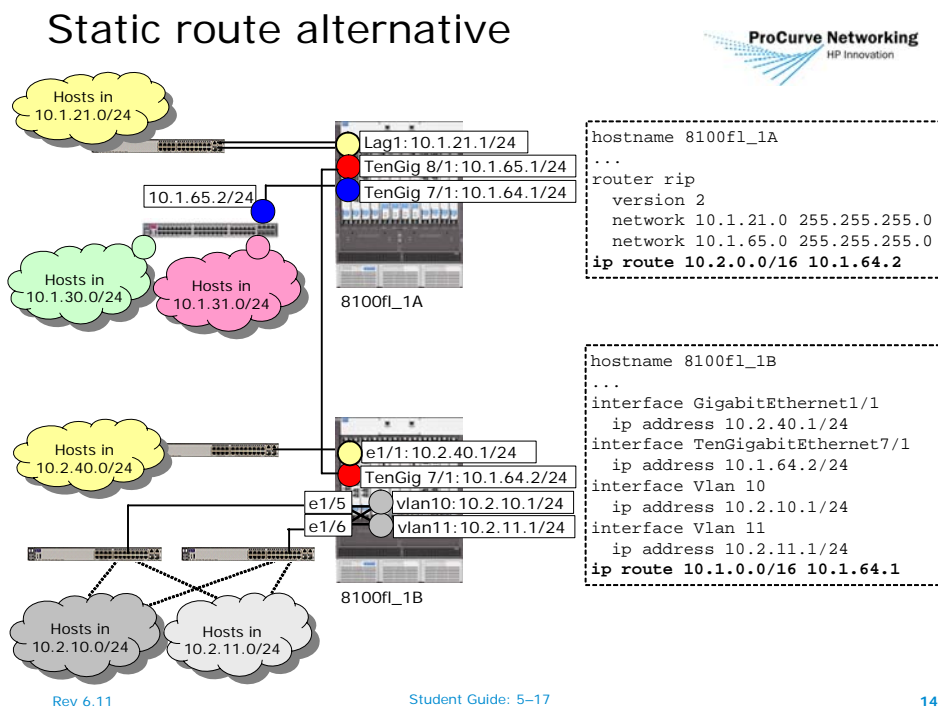


In this example, a new RIP router, labeled “8100fl_1B,” has been added to the topology and is connected to 8100fl_1A through TenGig 7/1. To enable the two 8100fl switches to exchange RIP updates, the administrator has entered *network 10.1.64.0 255.255.255.0* within the RIP router configuration context of both 8100fl routers.

Some networks on 8100fl_B are associated with port-based interfaces (Gig 1/1 and TenGig 7/1) and others are associated with VLAN interfaces. Gig 1/5 and Gig 1/6 are configured as trunk mode switchports associated with VLAN 10 and VLAN 11. From a routing perspective, there is no functional difference between port-based and VLAN interfaces.

In the configuration output for 8100fl_1B, each port must be enabled using the *no shutdown* command. However, for the sake of brevity, these commands are not shown in the example.

Static route alternative



Instead of enabling RIP on the interface between the 8100fl switches, the administrator could configure a static route on each switch that defines the entire address range on the other switch, including the neighbor's interface.

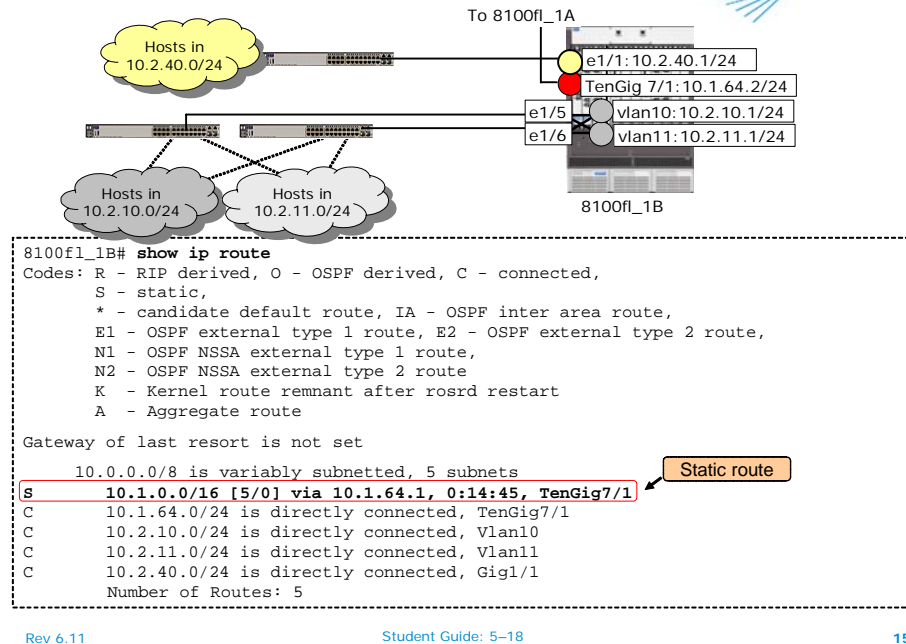
In the example, the administrator defined a static route to the network range 10.1.0.0/16 on the 8100fl_1B and defined a route to 10.2.0.0/16 on the 8100fl_1A. The static route defined on each 8100fl specifies an interface on the network that the routers share.

As well as eliminating routing protocol traffic over a given link, the definition of a static route provides an opportunity to summarize address space. In the example, the static routes will minimize the number of route table entries by substituting a summarized route for a longer list of individual networks.

Like many other ProCurve switches, the 8100fl automatically summarizes routes on classful boundaries. However, auto-summary would not be helpful in the example because the ranges 10.1.0.0/16 and 10.2.0.0/16 are in the same classful network 10.0.0.0/8.

Static entries in IP route table

Static entries in IP route table



Rev 6.11

Student Guide: 5-18

15

When you define a static route, the address space and next hop appear in the IP route table as a type “S” route. The default administrative distance for static routes is 5, and the default cost is 0. You can change the administrative distance of a static route by adding the value to the statement that defines the route. For instance, the command *ip route 10.1.0.0/16 10.1.64.2 1* would place the static route into the table with an administrative distance of 1, giving it a higher preference than other static routes to the same destination that use the default administrative distance of 5.

In the display, the classful network 10.0.0.0/8 is considered “variably subnetted” because a 16-bit mask is applied to the static route and 24-bit masks to all other subnets in that range.

Choose the best path

Choose the best path



The 8100fl switch may learn multiple paths to the same remote address range

- For multiple paths learned from *the same* routing method:
 - The path with the lowest cost is placed in the IP route table
 - Traffic may be distributed over equal cost paths
- For multiple paths learned from *different* routing methods:
 - The route learned from the routing method with the highest preference (lowest numerical *administrative distance* value) is placed in the IP route table
 - Default administrative distance for RIP routes: 100
 - Default administrative distance for static routes: 5

Rev 6.11

Student Guide: 5–19

16

The 8100fl supports up to four equal-cost paths for static routes and 16 for OSPF.

Like other routers, the 8100fl uses two different methods for choosing among multiple paths to the same network. When the 8100fl learns about an address range from multiple neighbors using the same routing protocol or method, the router places the route with the lowest metric or cost into the IP route table. The 8100fl supports up to 16 equal-cost paths for static routes and up to 16 equal cost paths for OSPF.

The 8100fl also supports up to four equal-cost paths for RIP; however, due to RIP's path calculation algorithm, equal-cost paths in RIP may not support equal bandwidth and may actually impair rather than improve performance. OSPF is often recommended over RIP in environments with redundant routed links because it calculates path cost based on bandwidth. Other ProCurve Intelligent Edge Switches do not support equal-cost paths for RIP routes, and this support may not continue on the 8100fl platform indefinitely.

8100fl Administrative Distance

Type	Default Administrative Distance
Connected	0
Static	5
OSPF Inter-area	10
OSPF Intra-area	10
OSPF External (Type 2)	150
RIP	100

8100fl OSPF support

8100fl OSPF support



Supports equal-cost multi-path

- Up to 16 paths learned from OSPF and static routes
- Traffic flows distributed using round-robin algorithm

Each router can have up to 55 adjacent neighbors

- Supports simple password and MD5 authentication methods
- Authentication configurable per interface and per area

Loopback interface provides Router ID

- Alternate 32-bit value may be statically defined

Supports all standard area types

- Summarize address space at area border router
- Summarization filters

Supports *virtual link* to backbone through another area

- Ensure continued connectivity on failure of area's direct backbone link(s)

Supports redistribution of RIP, static, and connected routes

- Redistribution list
- External route summarization filter

Rev 6.11

Student Guide: 5-21

17

The 8100fl features robust support for OSPF, including support for:

- Up to 16 equal-cost paths
- Up to 55 adjacent neighbors
- All standard area types
- Virtual links

By default, like many routers, the 8100fl derives its Router ID from the address of its loopback interface. However, you can statically define a 32-bit Router ID.

As well as supporting the standard area types, the 8100fl supports the summarization of address space at area border boundaries. The common area types include:

Normal

- Intra-area, inter-area, and external routes appear in route table as individual networks.
- ASBR permitted

Stub

- Intra-area and inter-area routes appear in route table as individual networks.
- External (non-OSPF) routes are summarized as default route.
- ASBR not permitted

Stub No-summary

- Intra-area routes appear in route table as individual networks.
- Inter-area and external routes are summarized as default route.
- ASBR not permitted

Not-So-Stubby (NSSA)

- Intra-area and inter-area routes appear in route table as individual networks.
- External routes that originate outside the area are summarized as default route.
- ASBR permitted

Virtual link support

Occasionally, it is not cost-effective or practical to provision a physical connection from every area to the backbone area. If the single link in the non-backbone area fails, the area will be completely isolated. Virtual links overcome this limitation by providing indirect backbone links that transit other non-backbone areas.

For more information, see the “OSPF Configuration” section of the *Management and Configuration Guide*.

OSPF configuration steps

OSPF configuration steps



Define loopback interface address or statically define Router ID:

```
8100fl_1A(config)# interface loopback 1
8100fl_1A(config-interface-loop1)# ip address 10.1.208.1/32
8100fl_1A(config)# ip router-id <ip-address>
```

Enable OSPF and assign *process ID* at global configuration context:

```
8100fl_1A(config)# router ospf 1
```

Create and configure areas to which this router will connect:

```
8100fl_1A(config-router)# area 0
8100fl_1A(config-ospf-area)# area 1
8100fl_1A(config-ospf-area)# stub no-summary
```

Assign network ranges to areas:

```
8100fl_1A(config-ospf-area)# network 10.1.0.0 255.255.0.0 area 1
8100fl_1A(config-router)# network 10.0.0.0 255.255.0.0 area 0
```

- Redistribution of “connected” non-OSPF networks will prohibit assignment of non-backbone areas as “stub” or “stub no-summary”

Rev 6.11

Student Guide: 5–23

18

The steps for OSPF configuration on the 8100fl are similar to the steps for RIP configuration. The high-level steps are:

1. Define loopback interface address or statically define Router ID
2. Enable OSPF and assign process ID at global configuration context
3. Create and configure areas to which this router will connect
4. Assign network ranges to areas

Define Router ID

In general, it is recommended that you configure a Router ID before enabling OSPF so that you can avoid changing the ID after the OSPF interfaces are up and the router has formed adjacencies. To define a Router ID on the 8100fl, either configure a loopback interface address or configure a static ID using the *ip router-id* command, as shown above. This command is used in the same way on all ProCurve switches that support OSPF.

Enable OSPF

To enable OSPF, enter the *router ospf* command in the global configuration context, followed by a process ID, which is “1” in the example. When OSPF is enabled, the CLI will display a message indicating that no areas have been defined. This is normal operation.

Create and configure OSPF areas

As shown, you create an OSPF area by entering the *area* command, followed by an area ID, in decimal or dotted-decimal notation, in the OSPF configuration context. The context will change to the OSPF area configuration context, where you can define area types and define range summaries. When you define an area, the CLI will display a message indicating that no interfaces are associated with the area.

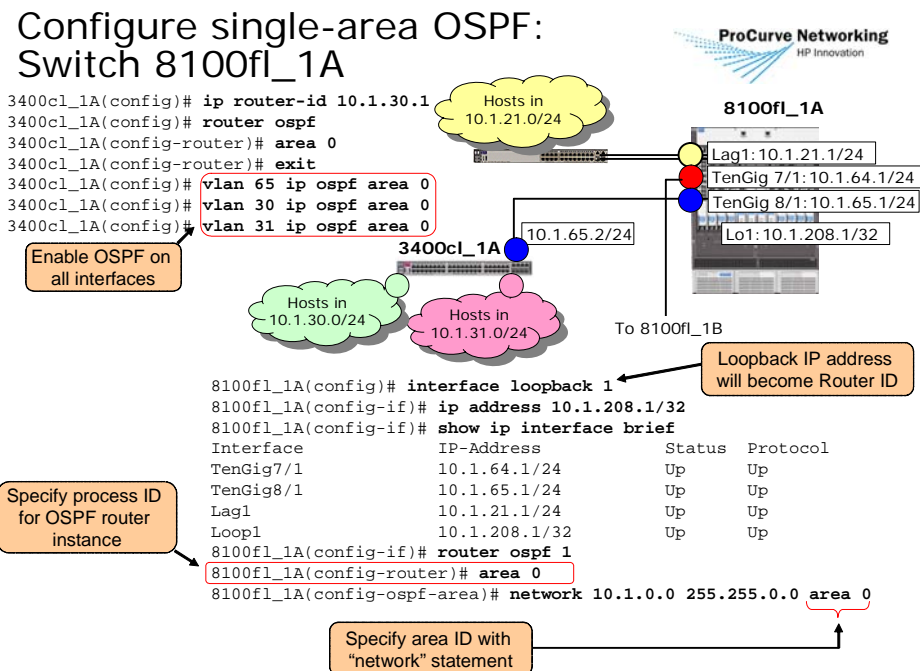
Assign network ranges to areas

Network statements within the OSPF router configuration context should specify ranges that include all IP interfaces whose address space should be included in the router's Link State Advertisements. As shown, the statements must include an area with which the address range will be associated. This is true even if you have only defined area 0 and all OSPF interfaces are to be assigned to area 0. The CLI will return an error if the specified area does not exist.

OSPF redistribution

Connected networks not defined within network statements may be redistributed; however, if you choose to redistribute networks into OSPF, it can limit your choices for assignment of OSPF area types. OSPF Link State Advertisements indicate whether each route is a native OSPF interface or was redistributed into OSPF. If you choose to redistribute a connected network instead of defining it as a native OSPF interface, the area to which the router belongs cannot be defined as a Stub or Stub No-summary area.

Configure single-area OSPF: Switch 8100fl_1A



Rev 6.11

Student Guide: 5–25

19

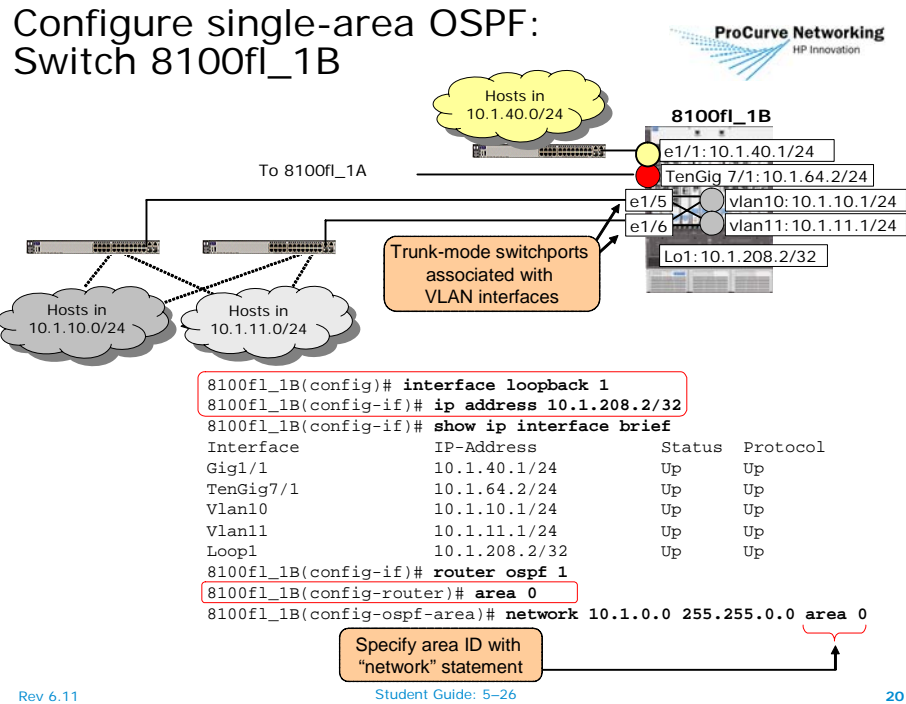
In this example, an administrator configures 8100fl_1A and 3400cl_1A to participate in a single OSPF area.

On the 8100fl, the administrator follows the four steps described on the previous page: define a loopback interface, enable OSPF, define an area, and assign network ranges to areas. Note that the *router ospf* command includes a process ID; the process ID will be required anytime the administrator wants to enter the OSPF router configuration context in the future, even if only one process is defined. Furthermore, the administrator must define an area, even if the router will only be a member of the backbone area. The *show ip interface brief* command is not required, but information about defined IP interfaces is useful when developing *network* statements.

In the example, a Router ID was statically defined on the 3400cl before OSPF was enabled. If you want to subsequently change the Router ID, you must first disable OSPF. OSPF may be associated with VLAN interfaces from the global configuration context, by following the VLAN ID with the *ip ospf* command. On the 3400cl and other ProCurve Intelligent Edge Switches that support OSPF, the *ip ospf* command associates the interface with area 0 if this area is configured on the switch. To associate the interface with an area other than 0, you must specify the area and its ID.

You may, alternatively, enter the *ip ospf area <id>* command in the configuration context of each VLAN. When the configuration is displayed, *area 0* is expressed as “area backbone” within the VLAN context.

Configure single-area OSPF: Switch 8100fl_1B




In this example, the administrator configures OSPF on 8100fl_1B. Like 8100fl_1A, this router associates its edge and non-edge interfaces with OSPF using a network statement that specifies the address space 10.1.0.0/16.

The routers must have other matching configuration items, such as Hello Interval and Dead Interval. The default values for these parameters are the values recommended in OSPF standard specifications. They are considered optimal for broadcast-type networks such as Ethernet. The intervals are sometimes adjusted for types of WAN links.

View OSPF neighbor table

View OSPF neighbor table



```
8100fl_1A# show ip ospf neighbor
```

OSPF Router with ID (10.1.208.1) (Process ID 1)

ID	Pri	State	Dead Time	Address	Interface
10.1.208.2	1	FULL/BDR	0:00:38	10.1.64.2	TenGig7/1
10.1.30.1	1	FULL/DR	0:00:21	10.1.65.2	TenGig8/1

```
8100fl_1B# show ip ospf neighbor
```

OSPF Router with ID (10.1.208.2) (Process ID 1)

ID	Pri	State	Dead Time	Address	Interface
10.1.208.1	1	FULL/DR	0:00:35	10.1.64.1	TenGig7/1

Rev 6.11

Student Guide: 5-27

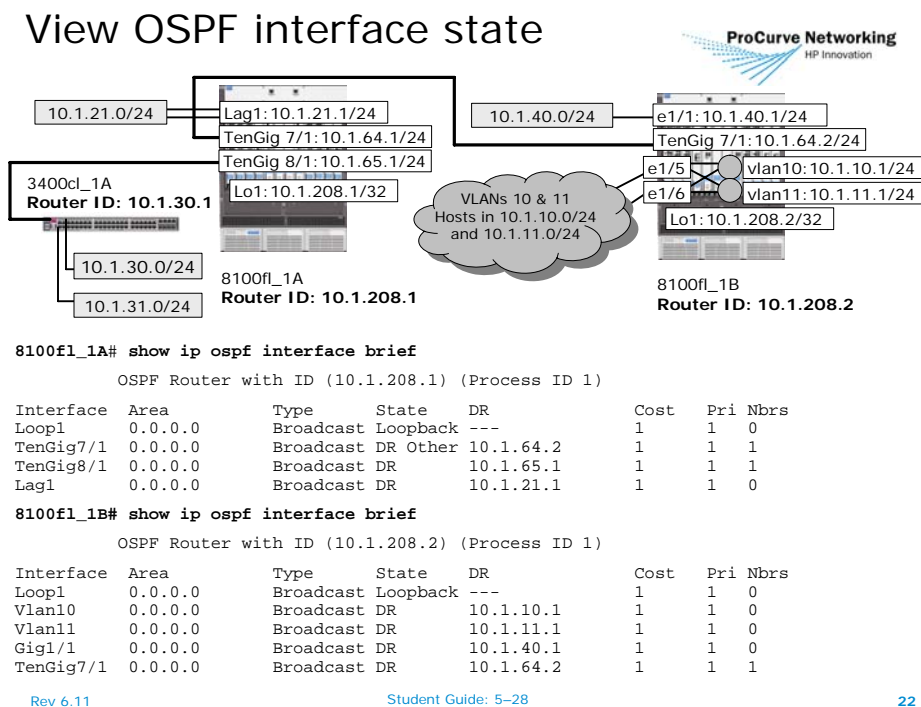
21

The two 8100fl switches have now become OSPF neighbors. On 8100fl_A, the *show ip ospf neighbor* output lists 8100fl_B and 3400cl_1A, which also has been configured for OSPF. 8100fl_1A has established a full adjacency with both neighbors because each network has only two routers.

In the diagram, all router-to-router connections are shown in heavy black lines. None of the routers finds a neighbor on the edge interfaces, which are shown in lighter lines.

8100fl_1A is the backup DR for the network it shares with 8100fl_1B because its Router ID (10.1.208.1) is lower than the Router ID of 8100fl_1B (10.1.208.2). However, 8100fl_1A is the DR of the network it shares with 3400cl_1A because its Router ID is higher than the Router ID of 3400cl_1A (10.1.30.1). By default, a router becomes the DR if its Router ID is the highest on a given network. You can affect the DR selection by configuring a higher OSPF priority for an interface.

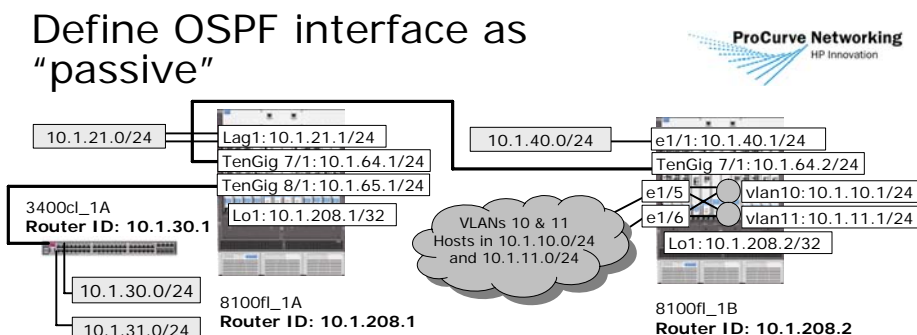
View OSPF interface state



The *show ip ospf interface* command provides details on OSPF parameters for all interfaces. You can make the output less verbose by using the *brief* keyword as shown. However, the detailed output often is useful for troubleshooting adjacency failures.

Each interface sends Hello messages every 10 seconds in an effort to find neighbors. Note that each edge interface considers itself to be the DR.

Define OSPF interface as “passive”



```
8100fl_1A(config)# router ospf 1
8100fl_1A(config-router)# passive-interface Lag1
8100fl_1A(config-router)# show ip ospf interface brief
```

OSPF Router with ID (10.1.208.1) (Process ID 1)

Interface	Area	Type	State	DR	Cost	Pri	Nbrs
Loop1	0.0.0.0	Broadcast	Loopback	---	1	1	0
Tengig7/1	0.0.0.0	Broadcast	DR other	10.1.64.2	1	1	1
Tengig8/1	0.0.0.0	Broadcast	DR	10.1.65.1	1	1	1
Lag1	0.0.0.0	Broadcast	Wait	0.0.0.0	1	1	0

Rev 6.11

Student Guide: 5-29

23

An OSPF interface defined as “passive” does not send Hello messages and cannot form adjacencies with neighbors. However, because the interface is defined as an OSPF interface, its associated address space is included in outbound advertisements. Only edge interfaces should be defined as passive.

On the 8100fl, the *passive-interface* command is entered at the OSPF router configuration context and takes the interface name as an argument.

In the diagram, LAG1 on 8100fl_1A is defined as a passive interface.

Passive interfaces on 8100fl_B

8100fl_1B is configured with three edge networks that have no OSPF neighbors. The port-based and VLAN interfaces may be defined as passive interfaces with the following commands:

```
8100fl_1B(config)# router ospf 1
8100fl_1B(config-router)# passive-interface e1/1
8100fl_1B(config-router)# passive-interface vlan 10
8100fl_1B(config-router)# passive-interface vlan 11
```

OSPF interface parameters

OSPF interface parameters



Items that may be defined within interface configuration context

- Define cost per OSPF interface
 - `8100fl_1A(config-interface-lag1)# ip ospf cost 10`
 - Calculated automatically using 1 Gbps as reference bandwidth
 - Change reference bandwidth to 10 Gbps:
 - `8100fl_1A(config-router)# auto-cost reference-bandwidth 10000`
- Router priority level (used in electing DR)
 - `8100fl_1A(config-interface-10gig7/1)# ip ospf priority 100`
- Authentication key
- Authentication method
 - May also be defined in area configuration context
- Hello Interval
- Dead Interval
- SPF Delay Timer

Rev 6.11

Student Guide: 5–30

24

Several important OSPF parameters can be defined within the interface context. All of the interface-level parameters shown above can be applied to VLAN or port-based interfaces. To see a list of OSPF parameters available in an interface context, enter *ip ospf ?*.

Cost assignment

On the 8100fl, you can change the cost of any interface by issuing the *ip ospf cost* command in the interface configuration context. Often, however, it is more useful to change the reference bandwidth on which all automatic costs are based. By default, on the 8100fl, automatic assignment of costs is enabled using a reference bandwidth of 1000 Mbps. As a result, gigabit interfaces are assigned a cost of 1.

On an 8100fl with 10-GbE interfaces, or when the 8100fl is put into an existing environment that uses 10-GbE elsewhere, it may be advisable to change the reference-bandwidth to 10000 Mbps. With this configuration, the cost of 10-GbE interfaces will be 1. Gigabit interfaces will be assigned a cost of 10. If any 100 Mbps interfaces are installed, they will have a cost of 100. As shown, the reference bandwidth is changed in the OSPF configuration context.

Router priority

You can affect the selection of the DR on a given network by assigning a particular router with a higher priority by issuing the *ip ospf priority* command.

This configuration does not necessarily guarantee that the router will be elected DR, especially on a multi-access network with more than two routers. If the router is one of the first two to discover each other, its higher priority will cause it to be elected DR. If the priority is configured on an interface where a DR already exists, the DR will not relinquish its role because a new router advertises a higher priority.

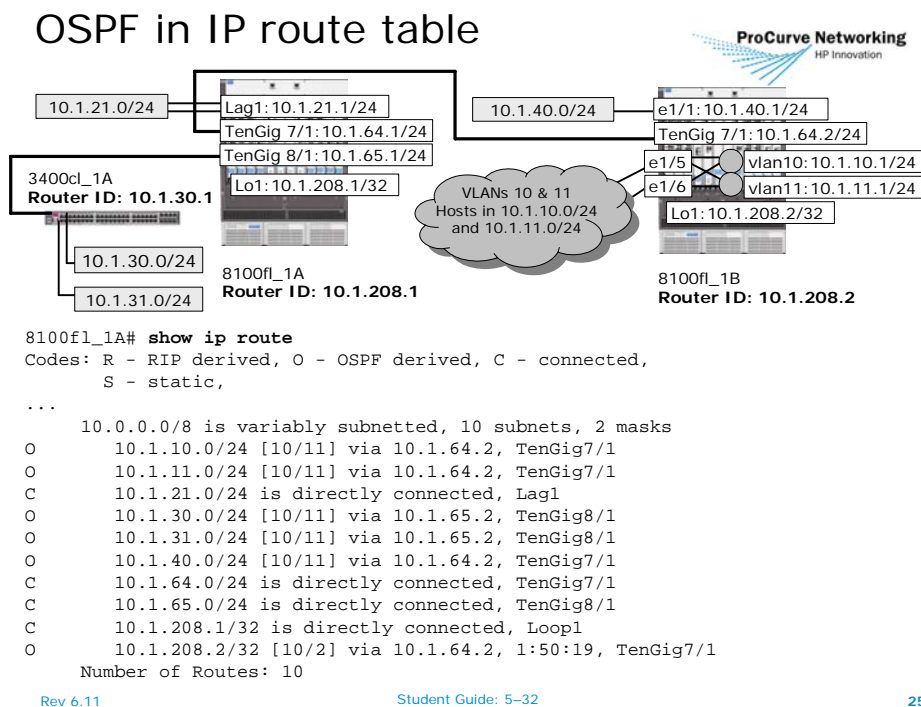
Authentication key

OSPF authentication is designed to ensure that unauthorized routers do not introduce inaccurate or disruptive route information into the OSPF domain. When enabling OSPF authentication on the 8100fl, you can specify a simple key or message digest key per area or per interface.

Other parameters

Other parameters include the Hello and Dead Intervals and the SPF Delay Timer. The SPF Delay Timer sets a minimum period of time between iterations of the shortest path first algorithm. This prevents the router from continuously running the algorithm in a dynamic environment, which would degrade performance because the switch cannot route IP packets while the algorithm is running.

OSPF in IP route table

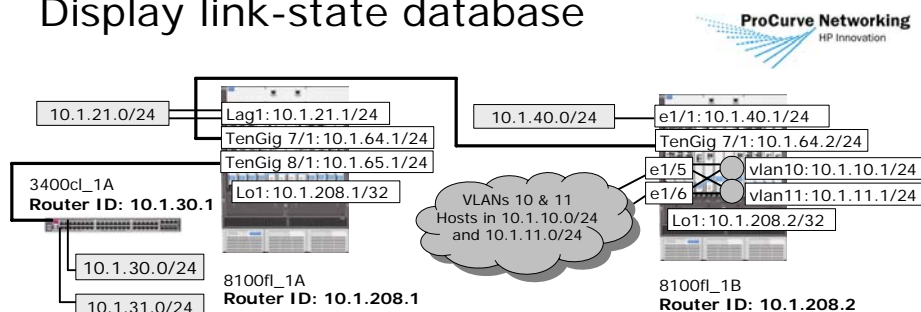


The IP route table of 8100fl_1A shows the impact of the OSPF configuration. The routes whose entries begin with an “O” have been derived from Link State Advertisements received from other routers. The table also shows several directly connected networks, whose entries begin with a “C.” The classful network 10.0.0.0/8 is described as “variably subnetted” because some of the interfaces in this range have a 24-bit mask and the loopback interfaces have a 32-bit mask.

Each OSPF route entry includes the interface used to reach the destination address range and the next hop, which is the router that leads toward the address range. Each entry also includes two values surrounded by brackets. The first number, “10,” is the administrative distance the 8100fl associates with OSPF routes. The second number, “11,” is the cumulative cost between this router and the address range. Due to the change in the “auto-cost reference-bandwidth,” each of the 10-GbE links that interconnect the routers has a cost of 1. The gigabit links have a cost of 10. As described earlier, costs for individual links may be altered at the interface configuration context.

Display link-state database

Display link-state database



```
8100fl_1A# show ip ospf database
          OSPF Router with ID (10.1.208.1) (Process ID 1)

          Router Link States (Area 0.0.0.0)

Link ID        ADV Router    Age      Seq#       Checksum    Link Count
10.1.30.1      10.1.30.1      648      0x8000000F 0x9F04      3
10.1.208.1     10.1.208.1     634      0x80000014 0xDF5A      4
10.1.208.2     10.1.208.2     1253     0x80000012 0xA358      5

          Net Link States (Area 0.0.0.0)

Link ID        ADV Router    Age      Seq#       Checksum    Router Count
10.1.64.2      10.1.208.2    394      0x80000007 0xDC93      2
10.1.65.1      10.1.208.1    394      0x80000004 0x39E2      2
```

Rev 6.11

Student Guide: 5-33

26

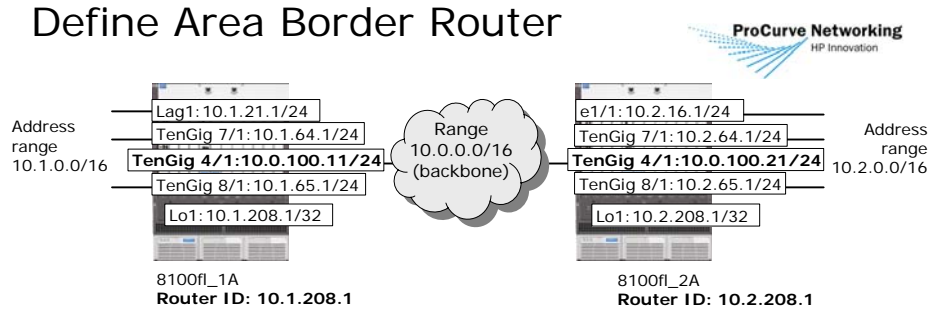
Like all OSPF routers, the 8100fl uses the information in its link-state database to determine the shortest path to each destination network and places the next hop on the shortest path into its IP route table.

In the example, the link-state database on 8100fl_1A contains one entry for each LSA the router has received. The advertisements are separated by type.

Because the example includes three routers, the database contains three Router LSAs. The DR of each multi-access network sends a Network LSA.

All routers in the area should have identical link-state database entries.

Define Area Border Router



Create backbone interface:

```
8100fl_1A(config)# interface tengig 4/1
8100fl_1A(config-interface-10gig4/1)# ip address 10.0.100.11/24
```

Enable OSPF, assign a process ID, and create areas:

```
8100fl_1A(config-interface-10gig4/1)#router ospf 1
8100fl_1A(config-router)# area 0
8100fl_1A(config-ospf-area)# area 1
```

Define range for address summarization:

```
8100fl_1A(config-ospf-area)# range 10.1.0.0/16
```

Assign interface address ranges to existing OSPF areas:

```
8100fl_1A(config-ospf-area)# network 10.0.0.0 255.255.0.0 area 0
8100fl_1A(config-router)# network 10.1.0.0 255.255.0.0 area 1
```

Rev 6.11

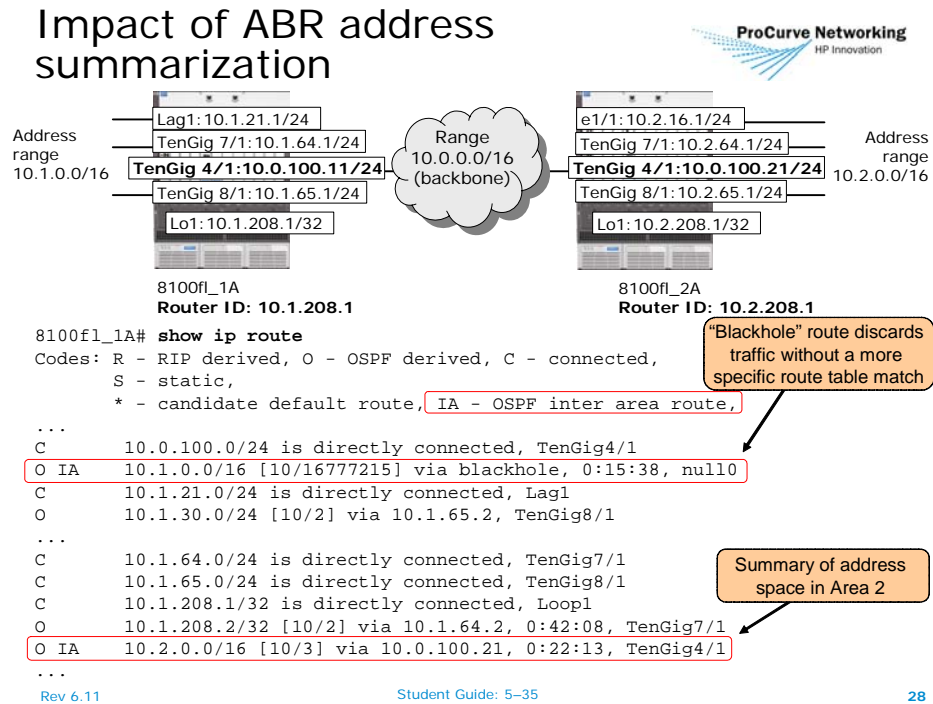
Student Guide: 5-34

27

Dividing an OSPF domain into multiple areas can improve overall efficiency, especially when the domain encompasses more than 50 routers or 500 networks. In a multiple-area topology, a router with interfaces in more than one area is called an Area Border Router (ABR). You can configure a router to function as an ABR by defining more than one area and assigning networks to the areas. One of the areas must be the backbone area, 0.0.0.0.

In the example, an administrator has created two normal OSPF areas. The division of an OSPF domain into multiple areas enables efficient summarization of address space at the ABRs. In the example, the administrator has defined the range 10.1.0.0/16 as the network range that 8100fl_1A will advertise into the backbone instead of the individual networks within the range.

Impact of ABR address summarization



The definition of a network range and its association with an area immediately affects the behavior of an 8100fl in two ways.

Like all routers, the 8100fl advertises the summary route rather than the individual networks that exist in the area. This benefits routers in other areas because they are not required to store all of the networks in the area 1 in their route tables. Instead, they store only the summary route.

By default, the 8100fl places a “blackhole” route for the summarized range in its route table. The “null0” keyword causes the router to discard matching traffic. However, the router does not discard all packets with a destination address in the specified range, only those that do not have a more specific match in the route table.

In the example, if a packet arrives at the router with a destination address of 10.1.30.42, the router will forward the packet through TenGig 8/1 because the route table has a more specific match for that packet. The “discard” action associated with the less specific match (10.1.0.0/16) is not invoked or followed if there is a more specific match. On the other hand, if a packet is destined for the address 10.1.250.250 and there isn’t any match other than 10.1.0.0/16, the packet is discarded.

To override the default null route, add the *no-discard* option when defining the range.

Autonomous System Boundary Router

Autonomous System Boundary Router



An ASBR redistributes information about external (non-OSPF) networks into the OSPF domain

- Examples of external information include:
 - Connected networks not defined as OSPF interfaces
 - Networks in route table learned from RIP neighbors
 - Static routes that appear in route table
- Redistribute external information into OSPF and optionally define within 'redistribute' configuration context:
 - Default metric
 - Metric type (External 1 or External 2):

```
8100fl_1A(config)# router ospf 1
8100fl_1A(config-router)# redistribute connected
8100fl_1A(config-ospf-redistribute-connected)# redistribute rip
8100fl_1A(config-ospf-redistribute-rip)# redistribute static
8100fl_1A(config-ospf-redistribute-static)#
```

Rev 6.11

Student Guide: 5–36

29

An ASBR is an OSPF router with non-OSPF interfaces whose address space it must advertise into the OSPF domain. In the example, 8100fl_1A is an ASBR because the administrator has configured redistribution of non-OSPF interfaces instead of defining them as passive interfaces.

When redistributing non-OSPF information into the OSPF domain, you can define filters to restrict the information that is redistributed. For example, you can configure the router to advertise only some of its RIP-learned networks into OSPF.

Define additional OSPF area types

Define additional OSPF area types



OSPF area type is defined within OSPF area configuration context

- Define area as stub type
 - External information summarized as default route
 - AS boundary routers not permitted within area

```
8100fl_1A(config)# router ospf 1
8100fl_1A(config-router)# area 1
8100fl_1A(config-ospf-area)# stub
```
- Define area as 'totally stubby'
 - External and inter-area routes summarized as default route
 - AS boundary routers not permitted within area

```
8100fl_1A(config-router)# area 1
8100fl_1A(config-ospf-area)# stub no-summary
```
- Define area as not-so-stubby
 - External information originating outside the area is summarized as default route
 - AS boundary routers permitted within the area

```
8100fl_1A(config-router)# area 1
8100fl_1A(config-ospf-area)# nssa
```

Rev 6.11

Student Guide: 5–37

30

While other ProCurve switches require that you define an area's type with the same command that creates the area, the 8100fl creates a configuration context for each defined area. "Normal" is the default area type and does not need to be specified within the area configuration context.

However, to define a Stub, Stub No-summary, or Not-So-Stubby Area, you must specify the area type within the area's configuration context as shown. The differences among area types primarily concern how the ABRs advertise information about non-OSPF networks to other routers in the area.

If you change an area type from normal to Stub, Stub No-summary or NSSA, any existing adjacencies the router has with other routers in the area will be terminated. The adjacencies will not be reestablished until the neighbor routers have been configured to support the new area type. Area type must match if two routers are to establish and maintain an adjacency.

The only exception to this rule is the Stub No-summary area type, which is a variation of the Stub area type that is relevant only for ABRs. An ABR configured as Stub No-summary withholds Type 3 Summary LSAs from its neighbors within a Stub area. Internal routers within the Stub area use the default route in place of inter-area OSPF routes as well as external routes.

When an area contains AS boundary routers, which are OSPF routers that must redistribute non-OSPF information into the OSPF domain, the area cannot be defined as a Stub or Stub No-summary area type. Some examples of non-OSPF networks include connected networks, RIP-derived networks, and static routes. If you want the routers in the area to have the benefits of a Stub area—that is, to use the default route to reach external networks that originate outside the area—you can define it as a Not-So-Stubby Area. ABRs and internal routers within a Not-So-Stubby Area must agree on the area type in order to form adjacencies.

For more information on area types, see the *IP Routing Foundations* course.

Module 5 summary

Module 5 summary



In this module, you learned how to configure IP routing on the 8100fl

Topics included:

- Configuring RIP
- Configuring OSPF
- Defining static routes
- Interpreting the IP route table

Module 5 described the processes for configuring IP routing on the ProCurve Switch 8100fl.

Learning check

Module 5

1. Describe a situation where the definition of a port-based interface on an 8100fl would be appropriate.

.....
.....

2. How is the *network* statement used in RIP configuration on the 8100fl?

.....
.....
.....

3. Describe the 8100fl support for Equal Cost Multi-Path routing.

.....
.....

4. How does an 8100fl determine its router ID?

.....
.....

Module 5 lab overview

Module 5 lab overview



Module 5 lab tasks:

- Enable IP routing on the 5300xl
- Enable OSPF on the 5300xl
- Configure redundant routed links between 5300xl switches and 8100fl switches
- Configure OSPF on 8100fl
- Confirm connectivity with other groups
- Define a range summary on the 8100fl

Rev 6.11

Student Guide: 5–43

33

In the final lab activity, you will implement OSPF routing on the 5300xl and the 8100fl. You will observe the effects of connectivity with other groups on the 8100fl route table.

ProCurve Switch 8100fI

Command Reference

Appendix A


System-level tasks	8100fl	9300m	5300xl
Reset to factory defaults	8100fl#erase startup 8100fl#reload	9300m#erase startup 9300m#reload	5300xl#erase startup
Access Privileged EXEC or Manager level	8100fl>enable	9300m>enable	5300xl>enable (Required only if Operator level is password-protected)
Define password for Privileged EXEC/Manager level	8100fl(config)#enable secret <password>	9300m(config)#enable super-user <string>	5300xl(config)#password manager <string>
Enter configuration context	8100fl#configure [terminal]	9300m#configure terminal	5300xl#configure [terminal]
Enable Telnet access	8100fl(config)#ip telnet 8100fl(config-telnet)#no shutdown 8100fl(config)#line vty 0	Enabled by default (when an IP interface exists)	Enabled by default (when an IP interface exists)
Define hostname	8100fl(config)#hostname <name>	9300m(config)#hostname <name>	5300xl(config)#hostname <name>
Set system clock	8100fl#clock set <HH:MM:SS> <DD mo yr>	9300m#clock set <HH:MM:SS> <DD mo yr>	5300xl(config)#clock set <HH:MM:SS> <DD mo yr>
Save running configuration to startup	8100fl(config)#write memory	9300m(config)#write memory	5300xl#write memory
Copy configuration to TFTP server	8100fl#copy [start run] tftp://<ip add>/<filename>	9300m#copy [start run] tftp <ip-add><path/filename>	5300xl#copy [start run] tftp <ip-add> <path/filename>

Interface-level tasks	8100fl	9300m	5300xl
Assign IP address to a physical port	8100fl(config)#int <port-id> 8100fl(config-interface-gig<id>)#ip add <address/mask>	9300m(config)#int <port-id> 9300m(config-if-e1000-<port-id>)# ip add <address/mask>	N/A – supports VLAN-based interfaces only
Enable ports	8100fl(config-interface-gig<id>)# no shutdown (Ports will not forward traffic or illuminate LED without this command)	N/A	N/A
Create a VLAN	8100fl(config)#vlan <id>	9300m(config)#vlan <id>	5300xl(config)#vlan <id>
Associate ports that will carry untagged traffic for a single VLAN	8100fl(config)#int <id> 8100fl(config-interface- gig<id>)#switchport mode access vlan <id>	9300m(config)#vlan <id> 9300m(config-vlan-<id>)#untag <port-list>	5300xl(config)#vlan <id> untag <port-list> OR 5300xl(config)#vlan <id> 5300xl(vlan-<id>)#untag <port-list>
Define port as a tagged member of a VLAN	8100fl(config)#int <id> 8100fl(config-interface- gig<id>)#switchport mode trunk 8100fl(config-interface- gig<id>)#switchport trunk-vlans <vlan-id> 8100fl(config-interface- gig<id>)#switchport trunk-vlans <vlan-id> (Repeat to assign mult. VLANs to port)	9300m(config-vlan-<id>)#tag <port-list>	5300xl(config)#vlan <id> tag <port-list> OR 5300xl(config)#vlan <id> 5300xl(vlan-<id>)# tag <port-list>
Define the VLAN to which a port's untagged traffic should be assigned	8100fl(config-interface- gig<id>)#switchport trunk-native-vlan <id> (Native VLAN on a trunk port defaults to 1 – use above command to specify a different VLAN ID for untagged traffic on a trunk port)	9300m(config-if-e1000-<id>)# dual-mode [<vlan-id>] (By default, 9300m dual-mode port assigns untagged traffic to VLAN 1) [Note: 9408sl does not use dual- mode, allows explicit assignment of ports to default VLAN]	5300xl(config)#vlan <id> untag <port-list> OR 5300xl(config)#vlan <id> 5300xl(vlan-<id>)#untag <port-list>
Associate IP address with a VLAN	8100fl(config)#int vlan <id> 8100fl(config-interface-vlan<id>)#ip add <address/mask>	9300m(config)#vlan <id> 9300m(config-vlan-<id>)# router-int ve <id> 9300m(config-vlan-<id>)#int ve <id> 9300m(config-vif-<id>)#ip add <address/mask>	5300xl(config)#vlan <id> ip add <address/mask> OR 5300xl(config)#vlan <id> 5300xl(vlan-<id>)# ip add <address/mask>

Routing configuration	8100fl	9300m	5300xl
Enable RIP globally	8100fl(config)#router rip	9300m(config)#router rip	5300xl(config)#router rip
Enable RIP version globally	8100fl(config)#router rip 8100fl(config-router)#version [1 2] (You can override RIP version within interface configuration context)	N/A – version is assigned per interface (v1 is default)	N/A – version is assigned per interface (v2 is default)
Enable RIP on an interface	N/A. Specify network address and mask within RIP router configuration context. For example, to apply RIP to all router interfaces in the network 172.16.0.0/16: 8100fl(config)#router rip 8100fl(config-router)#network 172.16.0.0 255.255.0.0	9300m(config-[v]if-<id>)#ip rip [v1 v2-only v1-compatible-v2]	5300xl(vlan-<id>)# ip rip
Enable OSPF globally and define area 0	8100fl(config)#router ospf <process-id> 8100fl(config-router)#area 0	9300m(config)#router ospf 9300m(config-ospf-router)#area 0	5300xl(config)#router ospf 5300xl(ospf)#area 0
Enable OSPF on an interface	N/A – specify network address and mask within OSPF router configuration context. For example, to apply OSPF to all router interfaces in the network 10.0.0.0/8: 8100fl(config)#router rip 8100fl(config-router)#network 10.0.0.0 255.0.0.0	9300m(config-[v]if-<id>)#ip ospf area <id>	5300xl(vlan-<id>)#ip rip

ProCurve Switch 8100fl Series

Reviewer's Guide V2.0



- 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 10GbE X2 Media-Flexible Module (J8736A) 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	15
Software & 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	16
Minimizing System Downtime for Software Upgrades.....	17
Features and Benefits	17
L2 & L3 Performance and Latency.....	17
10 Gigabit Performance: Full Mesh, 8 & 16 10GbE ports.....	18
1 Gigabit Performance: Full Mesh, 80 & 160 GbE ports	18
Notes on Latency Measurements	18
High Availability	18
Optional Redundant Switch Fabric Modules	18
Optional Redundant Management Modules.....	18
Optional Redundant Power Supplies	19
Link Aggregation Groups (LAGs).....	19
IP Routing.....	19
Spanning Tree	19
Virtual Router Redundancy Protocol (VRRP)	19
Scalability	20
Passive Backplane Chassis.....	20
Scaleable 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	24

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, wire-speed, 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 of which 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, wire-speed non-blocking performance for up to eight 10GbE ports 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, wire-speed non-blocking performance for up to 16 10GbE 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 multi-stage replication for multicast traffic ensure unicast traffic is minimally affected 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 non-stop 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 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)

8-slot Chassis
Management Module (J8731A)
8-slot Fabric Module (J8729A)
Two (2) Fan Trays (J8989-61001)
1 Power Supply (J8732A)
Two (2) Power Supply Blanks
Nine (9) Blank slot covers

Switch 8116fl (J8728A)

16-slot Chassis
Management Module (J8731A)
16-slot Fabric Module (J8730A)
Two (2) Fan Trays (J8989-61001)
Two (2) Power Supplies (J8732A)
1 Power Supply Blank
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 bi-color LED to indicate booting, running or fault conditions) and one to indicate whether it is the active or non-active 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 uncompressed operating system and lookup tables. In addition, a 512 MB Compact Flash device is used for non-volatile 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 bi-color LED to indicate link state and speed (green for Gigabit, orange for 100 bps) and another to indicate port activity. The module has a bi-color 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 bi-color green/orange LED on the bottom to indicate booting, running or fault conditions.

Switch fl 10GbE X2 Media-Flexible Module (J8736A)

This single-port, X2 form factor 10 Gigabit module accepts X2 10Gbps transceivers, capable of full duplex, 10 Gigabit line rate. The port has Link and Activity LEDs. The module has a bi-color green/orange LED on the bottom to indicate system status (booting, running or fault conditions).

This J8736A 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 bi-color 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 8-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 event log entry are generated. Each fan tray is designed to cool its respective half of the chassis (containing fabric, management and interface modules).

¹ The earlier J8733A fixed-port LR 10G module was discontinued in early 2006. The functional equivalent is a J8736A 10G X2 Media Flexible module with a J8437A LR X2 transceiver.

² 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.

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 on-board 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. Simply click on the Online Shopping link, then the HP Parts Page link, to order part number J8989-61001.

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 8-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: 8100fl 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 is always 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 mostly located 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.

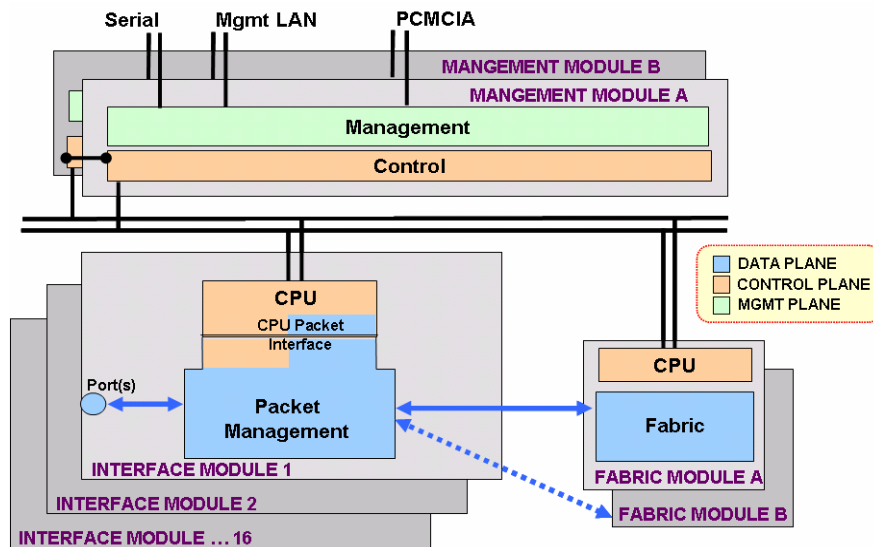


Figure 5: 8100fl Series architecture: Data/Control/Management Plane

Data Plane

The 8100fl Series data plane provides all of the processing and forwarding of packets throughout the system, which are primarily implemented 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).

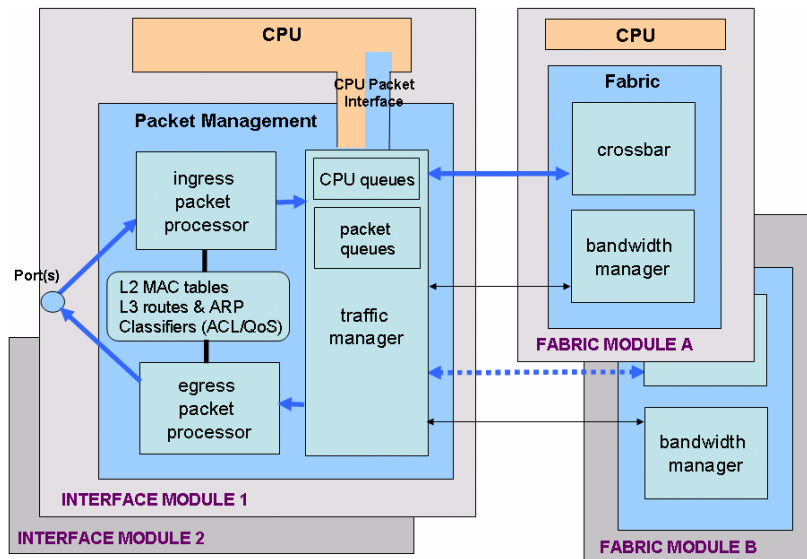


Figure 6: Components of the 8100fl Series data plane

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 1000 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, 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 provides efficient replication of multicast.

Backplane Interface Between Traffic Manager and Switching Fabric

Interface modules are physically connected 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 8-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.

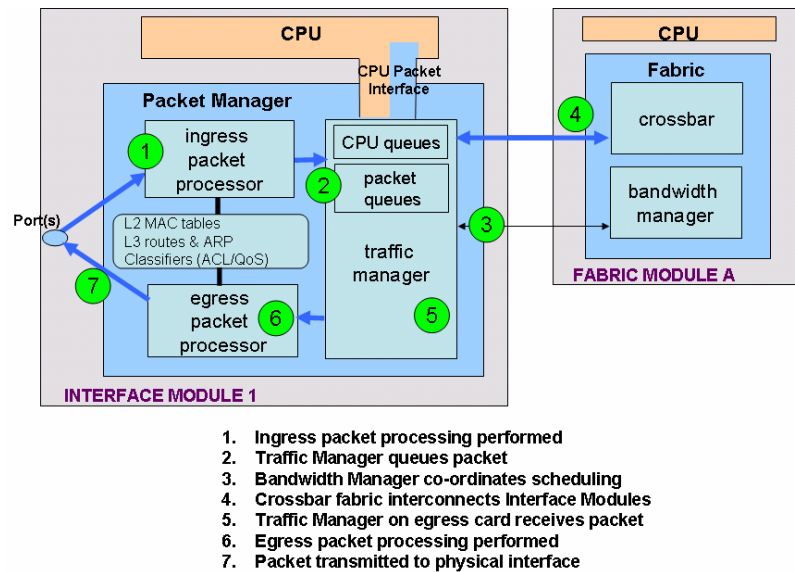


Figure 8: Data Plane Packet Walkthrough Overview

Data Plane Review: Packet Walkthrough

Figure 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 Class of Service (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 bi-directional 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, RIP, etc. 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 to ensure line-rate, hardware-based forwarding.

In addition, the control plane software on the management module translates user-specified policies and actions (ACLs, QoS policies, etc.) into data plane specific structures, 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 & 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, 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 are also 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 of which 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 non-volatile 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.

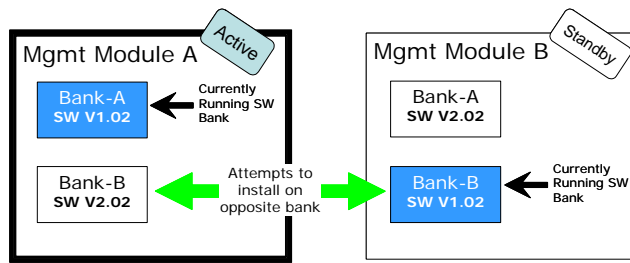


Figure 8: Management Module Software Bank Synchronization

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 & L3 Performance and Latency

- 10 Gbps full duplex bandwidth between modules
- Line rate non-blocking performance
- L2 & 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, 8 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 & 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

1 Gigabit Performance: Full Mesh, 80 & 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 – 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 8-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 wire speed 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 multi-path
 - Support for NSSA (NotSoStubbyAreas) 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:

- Disabling 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.

Scaleable Hardware Architecture

- 100K IP routes & MAC addresses, scaled up to 200K ³
- Up to 4000 Multicast Routes ⁴, scaled up to 8000 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 wire speeds 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
 - 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

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

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

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

⁶ See Hardware Architecture on page 9

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.
J8989-61001	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.

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.

⁷ In the USA, go to www.hp.com, click on Online Shopping, then HP Parts Store

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.

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.

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.

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 on Resellers.

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

www.procurve.com



© 2006 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

4AA0-2624ENW Rev 2, 02/2006