



# **BCNP in a Nutshell Study Guide for Exam 150-220**



Global Education Services  
Revision 0611

**Corporate Headquarters - San Jose, CA USA**

T: (408) 333-8000

info@brocade.com

European Headquarters - Geneva, Switzerland

T: +41 22 799 56 40

emea-info@brocade.com

Asia Pacific Headquarters - Singapore

T: +65-6538-4700

apac-info@brocade.com

© 2011 Brocade Communications Systems, Inc. All Rights Reserved.


Brocade, the Brocade B-weave logo, Fabric OS, File Lifecycle Manager, MyView, Secure Fabric OS, SilkWorm, and StorageX are registered trademarks and the Brocade B-wing symbol and Tapestry are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. FICON is a registered trademark of IBM Corporation in the U.S. and other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

Revision 0611

## BCNP in a Nutshell First Edition

---



**Objective:** The BCNP Nutshell guide is designed to help you prepare for the BCNP Certification, exam number 150-220.

**Audience:** The BCNP Nutshell self-study guide is intended for those who have successfully completed the CNP 300 Brocade Certified Network Professional (BCNP) Training course, and who wish to undertake self-study or review activities before taking the actual BCNP exam. The BCNP guide is not intended as a substitute for classroom training or hands-on time with Brocade products.

**How to make the most of the BCNP guide:** The BCNP guide summarizes the key topics on the BCNP exam for you in an easy to use format. It is organized closely around the exam objectives. We suggest this guide be used in conjunction with our free online knowledge assessment test. To benefit from the BCNP guide, we strongly recommend you have successfully completed the CNP 300 Brocade Certified Network Professional (BCNP) Training course.

We hope you find this useful in your journey towards BCNP Certification, and we welcome your feedback by sending an email to [jcannata@brocade.com](mailto:jcannata@brocade.com).

Joe Cannata  
Certification Manager

A handwritten signature in blue ink that reads "Joe Cannata".



# Table of Contents

## 1 - QoS and Voice Deployment

Rate Limiting .....	1
Rate Shaping .....	1
Traffic Classification and QoS .....	1
Voice over IP (VoIP) .....	5

## 2 - Security Concepts

MAC Port Security .....	6
802.1X Port Security .....	6
802.1X Message Exchange During Authentication .....	6
General Security Concepts .....	7

## 3 - OSPF Concepts

OSPF Hello Packet .....	10
Adjacency .....	10
Neighbor and Adjacency Establishment Process .....	10
OSPF AS, Areas, and Router Types .....	11
OSPF LSA Types .....	11
Link State Cost .....	12
OSPF Virtual Link .....	12
Redistribution .....	13
External Route Summarization .....	13

## 4 - BGP Concepts

EBGP Multihop .....	15
Using Loopback Interfaces for IBGP Peering .....	15
Multipath EBGP .....	15
Peer Group .....	15
Confederation .....	16
The BGP Table .....	17

## 5 - Advanced Layer 3 Concepts

Default Route Origination .....	20
IPv6 .....	21
PIM Sparse Mode .....	22
IGMP Snooping Overview .....	22
VRRP vs. VRRP-E .....	23
Layer 2 and Layer 3 Multicast Address Mapping .....	24
PBR .....	25

## 6 - Layer 2 Protocols

MRP .....	26
RSTP Bridges and Bridge Port Roles .....	28
DHCP Snooping .....	30
Spanning Tree Protocol (STP) .....	31
BPDU Guard .....	31
LLDP .....	32
Dual-mode VLAN Ports .....	32
Q-in-Q .....	34
VLAN .....	34

## 7 - Monitoring, Maintenance, and Troubleshooting

OSPF External Route Summarization .....	38
OSPF Internal Route Summarization .....	38
BGP Route Flap Dampening .....	39
OSPF Network Types and DR and BDR Election .....	40
OSPF Interface: Passive and Ignore .....	40
Dynamically Refreshing BGP Routes and Placing BGP Policy Changes Into Effect .....	41
Allocating Memory for More VLANs or Virtual Routing Interfaces .....	42
Specify Types of OSPF Syslog Messages to Log .....	43
Port Mirroring and Monitoring .....	43
SNMP .....	44
Recovering from a Lost Password .....	45

# List of Figures

Packet Trust Levels .....	3
802.1X Message Exchange .....	7
IPv6 Address Format .....	21
Multicast Address Layout .....	24
Metro Ring Example .....	26
Multiple Metro Rings .....	27
Multiple MRP Rings Sharing an Interface .....	28
Dual-mode VLAN Ports .....	32
Dual-mode Port .....	33
Q-in-Q Tagging .....	34
Private VLAN .....	36



# List of Tables

QoS Queues .....	4
Power Classes .....	5
Default Administrative Distances .....	21
Private and Standard Port-based VLANs Comparison .....	37
VLAN and Virtual Routing Interface Maximums .....	42



# 1 - QoS and Voice Deployment

## Rate Limiting

Inbound rate limiting allows you to specify the maximum number of Kbps a given port can receive. To configure inbound rate limiting on a port, enter the following commands:

```
FastIron(config) #interface ethernet 0/2/1
FastIron(config-if-e10000-0/2/1) #rate-limit input fixed 1000000
Rate Limiting on Port 0/2/1 - Config: 1000000 Kbps, Actual: 1000000 Kbps
```

The above commands configure a fixed rate limiting policy that allows port 0/2/1, a 10-GbE port, to receive a maximum of 1,000,000 kilobits per second. If the port receives additional bits during a given one-second interval, the port drops all inbound packets on the port until the next one-second interval starts.

Outbound rate limiting allows you to specify the maximum number of kilobits a given port can transmit.

- Port-based: Limits the rate of outbound traffic on an individual physical port or trunk port, to a specified rate. Traffic that exceeds the maximum rate is dropped. Only one port-based outbound rate limiting policy can be applied to a port.
- Port- and priority-based: Limits the rate on an individual 802.1p priority queue on an individual physical port or trunk port. Traffic that exceeds the rate is dropped. Only one priority-based rate limiting policy can be specified per priority queue for a port. This means that a maximum of eight port- and priority-based policies can be configured on a port.

Here is an example of port-based rate limiting:

```
FastIron(config) #interface ethernet 0/1/34
FastIron(config-if-e1000-0/1/34) #rate-limit output fixed 65
Outbound Rate Limiting on Port 0/1/34 Config: 65 Kbps, Actual: 65 Kbps
```

The above commands configure a fixed rate limiting policy that allows port 0/1/34 to transmit 65Kbps. If the port transmits additional bits during a given one-second interval, the port will drop all outbound packets on the port until the next one-second interval starts.

## Rate Shaping

Outbound Rate Shaping is a port-level feature that is used to shape the rate and control the bandwidth of outbound traffic on a port. This feature smoothes out excess and bursty traffic to the configured maximum limit before it is sent out on a port. Packets are stored in available buffers and then forwarded at a rate no greater than the configured limit. This process provides for better control over the inbound traffic of neighboring devices. The device has one global rate shaper for a port and one rate shaper for each port priority queue. Rate shaping is done on a single-token basis, where each token is defined to be 1 byte.

## Traffic Classification and QoS

Quality of Service (QoS) features are used to prioritize the use of bandwidth in a switch. When QoS features are enabled, traffic is classified as it arrives at the switch, and processed through on the basis of configured priorities. Traffic can be dropped, prioritized for guaranteed delivery, or subject to limited delivery options as configured by a number of different mechanisms.

Classification is the process of selecting packets on which to perform QoS, reading the QoS information and assigning a priority to the packets. The classification process assigns a priority to packets as they enter the switch. These priorities can be determined on the basis of information contained within the packet or assigned to the packet as it arrives at the switch. Once a packet or traffic flow is classified, it is mapped to a forwarding priority queue. Packets on Brocade devices are classified in up to eight traffic classes with values between 0 and 7. Packets with higher priority classifications are given precedence for forwarding.

### *Processing of Classified Traffic*

The trust level in effect on an interface determines the type of QoS information the device uses for performing QoS. The Brocade device establishes the trust level based on the configuration of various features and if the traffic is switched or routed. The trust level can be one of the following:

- Ingress port default priority
- Static MAC address
- Layer 2 Class of Service (CoS) value – This is the 802.1p priority value in the Ethernet frame. It can be a value from 0 – 7. The 802.1p priority is also called the Class of Service.
- Layer 3 Differentiated Service Code Point (DSCP) – This is the value in the six most significant bits of the IP packet header's 8-bit DSCP field. It can be a value from 0 – 63. These values are described in RFCs 2472 and 2475. The DSCP value is sometimes called the DiffServ value. The device automatically maps a packet's DSCP value to a hardware forwarding queue.

ACL keyword is an ACL that can also prioritize traffic and mark it before sending it along to the next hop.

[Figure 1 on page 3](#) illustrates how a Brocade device determines a packet's trust level:

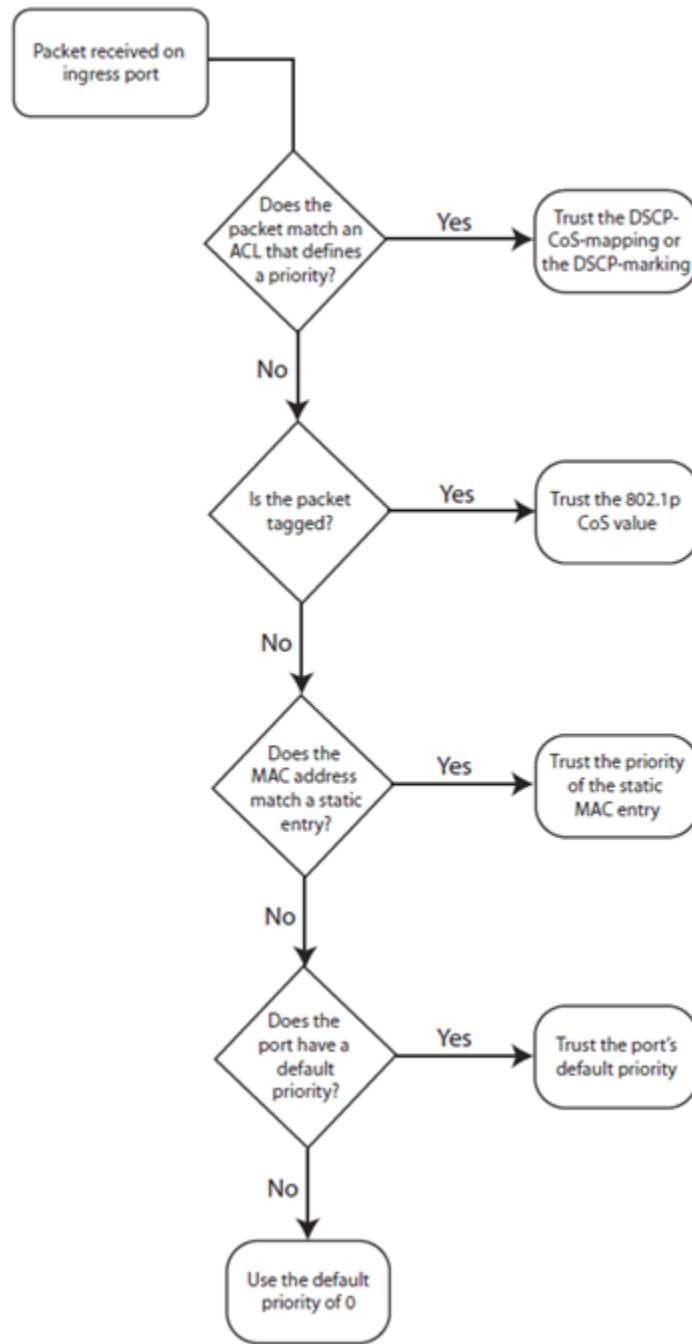


Figure 1: Packet Trust Levels

The first criteria considered is whether the packet matches on an ACL that defines a priority. If this is not the case and the packet is tagged, the packet is classified with the 802.1p CoS value. If neither of these are true, the packet is next classified based on the static MAC address, ingress port default priority, or the default priority of zero (0).

## QoS Queues

Brocade devices support eight QoS queues (qosp0 – qosp7) listed below:

TABLE 1 QoS Queues

QoS Priority Level	QoS Queue
0	qosp0 (lowest priority queue)
1	qosp1
2	qosp2
3	qosp3
4	qosp4
5	qosp5
6	qosp6
7	qosp7

## Traffic Policy Verification

To view traffic policies that are currently defined on the Brocade device, enter the `show traffic-policy` command. An example display output is shown below:

```
FastIron#show traffic-policy
Traffic Policy - t_voip:
Metering Enabled, Parameters:
  Mode: Adaptive Rate-Limiting
  cir: 100 kbps, cbs: 2000 bytes, pir: 200 kbps, pbs: 4000 bytes
Counting Not Enabled
Number of References/Bindings:1
```

The fields in the above output are explained as follows:

**Traffic Policy:** The name of the traffic policy.

**Metering:** Shows whether or not rate limiting was configured as part of the traffic policy:

- **Enabled:** The traffic policy includes a rate limiting configuration.
- **Disabled:** The traffic policy does not include a rate limiting configuration.

**Mode:** If rate limiting is enabled, this field shows the type of metering enabled on the port:

- **Fixed Rate-Limiting**
- **Adaptive Rate-Limiting**
- **cbs:** The committed burst size, in bytes per second, for the adaptive rate-limiting policy.
- **pir:** The peak information rate, in kbps, for the adaptive rate-limiting policy.
- **pbs:** The peak burst size, in bytes per second, for the adaptive rate-limiting policy.

**Counting:** Shows whether or not ACL counting was configured as part of the traffic policy:

- **Enabled** – Traffic policy includes an ACL counting configuration.
- **Disabled** – Traffic policy does not include an ACL traffic counting configuration.

**Number of References/Bindings:** The number of port regions to which this traffic policy applies. For example, if the traffic policy is applied to a trunk group that includes ports e 9/9, 9/10, 9/11, and 9/12, the value in this field would be 2, because these four trunk ports are in two different port regions.

## Voice over IP (VoIP)

### *Power Over Ethernet*

When Power over Ethernet (PoE) is enabled on a port to which a power consuming device is attached, by default, the Brocade PoE device supplies 15.4 watts of power at the RJ45 jack, minus any power loss through the cables.

To configure the maximum power level for a power consuming device, enter the following commands:

```
FastIron#config t
FastIron(config)#interface e 1/1
FastIron(config-if-e1000-1/1)#inline power power-limit 14000
```

These commands enable in-line power on interface e 1 in slot 1 and set the PoE power level to 14,000 milliwatts (14 watts). The default is 15400.

A power class specifies the maximum amount of power that a Brocade PoE device supplies to a power consuming device. The table below shows the different power classes and their respective maximum power allocations:

**TABLE 2 Power Classes**

Class	Maximum Power (Watts)
0	15.4 (default)
1	4
2	7
3	15.4
4	29 (FCS devices only)

By default, the power class for all power consuming devices is zero (0). A power consuming device with a class of 0 receives 15.4 watts of power. To configure the power class for a PoE power consuming device, enter the following commands:

```
FastIron#config terminal
FastIron(config)#interface e 1/1
FastIron(config-if-e1000-1/1)#inline power power-by-class 3
```

## 2 - Security Concepts

### MAC Port Security

You can configure the Brocade device to learn *secure* MAC addresses on an interface. The interface forwards only packets with source MAC addresses that match these learned secure addresses. The secure MAC addresses can be specified manually, or the Brocade device can learn them automatically. After the device reaches the limit for the number of secure MAC addresses it can learn on the interface, if the interface then receives a packet with a source MAC address that does not match the learned addresses, it is considered a security violation.

When a security violation occurs, a Syslog entry and an SNMP trap are generated. In addition, the device takes one of two actions either drops packets from the violating address (and allows packets from the secure addresses), or disables the port for a specified amount of time. You specify which of these actions takes place.

The secure MAC addresses are not flushed when an interface is disabled and re-enabled. The secure addresses can be kept secure permanently (the default), or can be configured to age out, at which time they are no longer secure. You can configure the device to automatically save the secure MAC address list to the startup-config file at specified intervals, allowing addresses to be kept secure across system restarts.

The port security feature applies only to Ethernet interfaces.

### 802.1X Port Security

The 802.1X standard defines the roles of *supplicant*, *authenticator*, and *authentication server* in a network. The *supplicant*, or client, provides username/password information to the authenticator. The authenticator sends this information to the authentication server. Based on the supplicant's information, the authentication server determines whether the supplicant can use the services provided by the authenticator. The authentication server passes this information to the authenticator, which then provides services to the supplicant, based on the authentication result.

The authenticator is the device that controls access to the network. In an 802.1X configuration, the Brocade device serves as the authenticator. The authenticator passes messages between the supplicant and the authentication server. Based on the identity information supplied by the supplicant, and the authentication information supplied by the authentication server, the authenticator either grants or denies network access to the supplicant.

The supplicant is the device that seeks to gain access to the network. Supplicants must be running software that supports the 802.1X standard (for example, the Windows XP operating system). Supplicants can either be directly connected to a port on the authenticator, or can be connected through a hub.

The authentication server is the device that validates the supplicant and specifies whether the supplicant may access services on the device. Brocade supports authentication servers running RADIUS.

### 802.1X Message Exchange During Authentication

[Figure 2 on page 7](#) illustrates a sample exchange of messages between an 802.1X-enabled supplicant, a FastIron switch acting as an authenticator, and a RADIUS server acting as an authentication server.

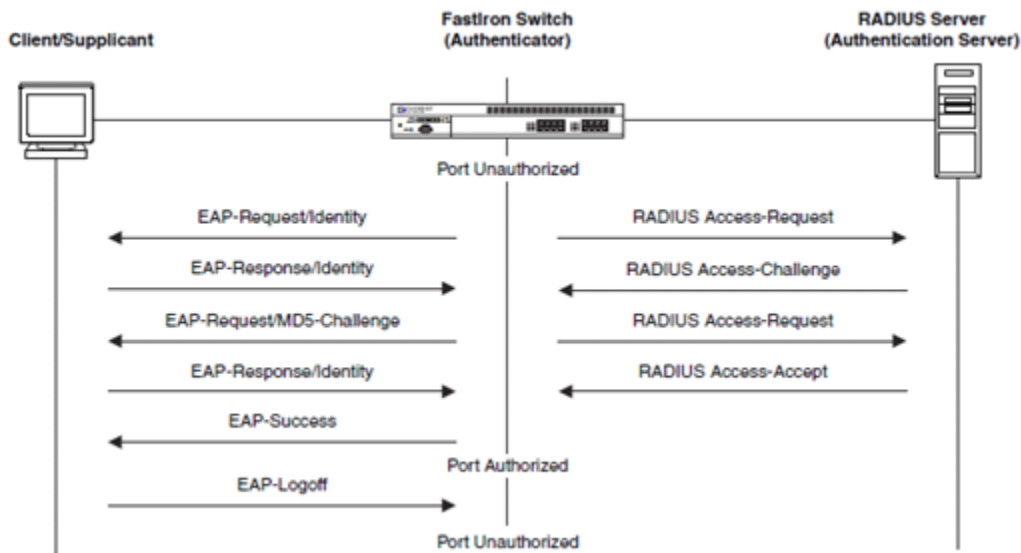


Figure 2: 802.1X Message Exchange

In this example, the authenticator (the FastIron switch) initiates communication with an 802.1X-enabled supplicant. When the supplicant responds, it is prompted for a username (255 characters maximum) and password. The authenticator passes this information to the authentication Server, which determines whether the supplicant can access services provided by the authenticator. When the supplicant is successfully authenticated by the RADIUS server, the port is authorized. When the supplicant logs off, the port becomes unauthorized again.

The Brocade 802.1X implementation supports dynamic VLAN assignment. If one of the attributes in the Access-Accept message sent by the RADIUS server specifies a VLAN identifier, and this VLAN is available on the Brocade device, the supplicant's port is moved from its default VLAN to the specified VLAN. When the supplicant disconnects from the network, the port is placed back in its default VLAN.

If a supplicant does not support 802.1X, authentication cannot take place. The Brocade device sends EAP-Request/Identity frames to the supplicant, but the supplicant does not respond to them. When a supplicant that supports 802.1X attempts to gain access through a non-802.1X-enabled port, it sends an EAP start frame to the Brocade device. When the device does not respond, the supplicant considers the port to be authorized, and starts sending normal traffic.

## General Security Concepts

### *Controlling User Access Into Devices*

In order to validate users requesting access into switches and routers, you need to specify what type of access and the authentication methods.

```
aaa authentication (what type of access) default (how to validate)
```

```
Syntax: aaa authentication <snmp-server|web-server|enable|login> default
<method1>[<method2> <method3> <method4> <method5> <method6> <method7>]
```

The access type can be Web, SNMP, Telnet, and Console. The validation methods can be Line, Enable, Local, RADIUS, TACACS, and TACACS+.

The following command example causes TACACS/TACACS+ to be the primary authentication method for securing Telnet/SSH access to the CLI. If the TACACS/TACACS+ authentication fails due to an error with the server, authentication is performed using local user accounts instead:

```
FastIron(config) #aaa authentication login default tacacs local
```

Here is another example:

```
FastIron(config) #aaa authentication web default radius line enable
```

To gain access to the switch or router using a web browser, first use: RADIUS usernames; if username/password is not configured or RADIUS server not available, then use Telnet password; if the Telnet password is not configured, then use the “enable” super-user, port-config, and read-only passwords.

## *Access Control List*

Brocade devices support rule-based Access Control Lists (ACLs; sometimes called hardware-based ACLs), where the decisions to permit or deny packets are processed in hardware and all permitted packets are switched or routed in hardware. All denied packets are also dropped in hardware.

Rule-based ACLs program the ACL entries you assign to an interface into Content Addressable Memory (CAM) space allocated for the ports. The ACLs are programmed into hardware at startup (or as new ACLs are entered and bound to ports). Devices that use rule-based ACLs program the ACLs into the CAM entries and use these entries to permit or deny packets in the hardware, without sending the packets to the CPU for processing.

ACLs consist of ACL IDs and ACL entries:

- *ACL ID* is a number from 1 – 99 (for a standard ACL) or 100 – 199 (for an extended ACL) or a character string. The ACL ID identifies a collection of individual ACL entries. When you apply ACL entries to an interface, you do so by applying the ACL ID that contains the ACL entries to the interface, instead of applying the individual entries to the interface. This makes applying large groups of access filters (ACL entries) to interfaces simple.
- *ACL entry*, also called an *ACL rule*, is a filter command associated with an ACL ID. The maximum number of ACL rules you can configure is a system-wide parameter and depends on the device you are configuring. You can configure up to the maximum number of entries in any combination in different ACLs.

You configure ACLs on a global basis, then apply them to the incoming traffic on specific ports. The software applies the entries within an ACL in the order they appear in the ACLs configuration. As soon as a match is found, the software takes the action specified in the ACL entry (permit or deny the packet) and stops further comparison for that packet.

## *Numbered and Named ACLs*

When you configure an ACL, you can refer to the ACL by a numeric ID or by an alphanumeric name. The commands to configure numbered ACLs are different from the commands for named ACLs.

- *Numbered ACL* – If you refer to the ACL by a numeric ID, you can use 1 – 99 for a standard ACL or 100 – 199 for an extended ACL.
- *Named ACL* – If you refer to the ACL by a name, you specify whether the ACL is a standard ACL or an extended ACL, then specify the name.

You can configure up to 99 standard numbered IP ACLs and 99 extended numbered IP ACLs. You also can configure up to 99 standard named ACLs and 99 extended named ACLs by number. If you can, try to apply ACLs “Inbound” rather than “Outbound”. On some platforms, outbound ACL is not supported.

The default ACL action when no ACLs are configured on a device is to permit all traffic. However, once you configure an ACL and apply it to a port, the default action for that port is to deny all traffic that is not explicitly permitted on the port:

- If you want to tightly control access, configure ACLs consisting of permit entries for the access you want to permit. The ACLs implicitly deny all other access.
- If you want to secure access in environments with many users, you might want to configure ACLs that consist of explicit deny entries, then add an entry to permit all access to the end of each ACL. The software permits packets that are not denied by the deny entries.

The following extended ACL command sequence example denies ping and allows other ICMP packets through. It blocks video streams from a 10.10.10.10 IP address to the 192.168.1.0/24 subnet, and allows all traffic not specifically denied to pass through.

```
FESX(config)# access-list 102 deny icmp any any echo log
FESX(config)# access-list 102 deny icmp any any echo-reply log
FESX(config)# access-list 102 permit icmp any any
FESX(config)# access-list 102 deny igmp host 10.10.10.10 192.168.1.0
0.0.0.255
FESX(config)# access-list 102 permit ip any any
FESX(config)# int e 1
FESX(config-if-1/1)# ip access-group 102 in
```

## AAA

Access control is the way you control who is allowed access to the network server and what services they are allowed to use once they have access. Authentication, authorization, and accounting (AAA) network security services provide the primary framework through which you set up access control on your router or access server.

- **Authentication** refers to the process of identifying users, including login and password dialog, challenge and response, messaging support, and encryption (depending on the security protocol you select). Authentication is the process to identify the user before he/she is allowed access to the network and network services.
- **Authorization** provides the method for remote access control, including one-time authorization or authorization for each service, per-user account list and profile, user group support, and support of IP, IPX, ARA, and Telnet. AAA authorization works by assembling a set of attributes that describe what the user is authorized to perform.
- **Accounting** provides the method for collecting and sending security server information used for billing, auditing, and reporting, such as user identities, start and stop times, executed commands (such as PPP), number of packets, and number of bytes.

## 3 - OSPF Concepts

### OSPF Hello Packet

All OSPF routers send Hellos to 224.0.0.5 (all OSPF routers on this subnet) to find neighbors. Certain items must match on both routers for them to become OSPF neighbors. These items are: subnet mask, area ID, Hello/Dead intervals, authentication password, and stub flag.

### Adjacency

Adjacency occurs when a relationship is formed between neighboring routers for the purpose of exchanging routing information. Adjacent OSPF neighbor routers go beyond the simple Hello packet exchange; they exchange database information. In order to minimize the amount of information exchanged on a particular segment, one of the first steps in creating adjacency is to assign a Designated Router (DR) and a Backup Designated Router (BDR). The Designated Router ensures that there is a central point of contact, thereby improving convergence time within a multi-access segment. In an OSPF point-to-point network, where a direct Layer 3 connection exists between a single pair of OSPF routers, there is no need for Designated and Backup Designated Routers, as is the case in OSPF multi-access networks. Without the need for Designated and Backup Designated routers, a point-to-point network establishes adjacency and converges faster. The neighboring routers become adjacent whenever they can communicate directly. In contrast, in broadcast and non-broadcast multi-access (NBMA) networks, the Designated Router and Backup Designated Router become adjacent to all other routers attached to the network.

Adjacency is the next step after the neighboring process (which is the simple Hello exchange during the Down, Init and 2-Way states). Adjacent routers are routers who go beyond the hello exchange and proceed into the database exchange process. Each router forms adjacency with the DR/BDR. When two routers' LSDBs become identical, they are said to be "adjacent" and reach the "full" neighbor state. OSPF routing updates are only sent across adjacencies to 224.0.0.6 (DR/BDR routers).

### Neighbor and Adjacency Establishment Process

An OSPF interface passes through the following steps before becoming adjacent to another router:

1. Down: No OSPF information has been received from any other router on the segment.
2. Attempt: On NBMA (non-broadcast multi-access) networks such as Frame Relay, this state indicates that no recent information has been received from the neighbor. An effort should be made to contact the neighbor by sending Hellos at the reduced rate Poll Interval.
3. Init: The interface has sent a Hello packet but bi-directional communication has not yet been established.
4. Two-way: The bi-directional communication has been established with a neighbor. The router has seen itself in the Hello packets coming from a neighbor. At the end of this stage the DR and BDR election would have been done. At the end of the 2-way stage, both routers will decide whether to proceed forward to build an adjacency or not. The decision is based on whether one of the routers is a DR or BDR, or the link is a point-to-point link, or a virtual link.
5. Exstart: Both routers try to establish the initial sequence number that is going to be used in the Exchange packets. The sequence numbers ensure that routers always get the most recent information. One router will become the master and the other will become secondary. The master will poll the secondary for information.

6. Exchange: Both routers describe their entire LSDB (Link-State Database) by sending DD (database description) packets.
7. Loading: If both routers' databases are not identical, they would go through this state. Routers finalize the information exchange in this state. Missing, incomplete, or outdated info will be put on a LSR (Link-State Request) list and sent to the neighbor. The neighbor replies with LSU (Link-State Update) packets. Any LSU that has been sent will be put on the retransmission list until it gets acknowledged (Link-State Acknowledgement).
8. Full: At this state, the adjacency is established. The neighboring routers are fully adjacent. Their LSDBs become identical.

## OSPF AS, Areas, and Router Types

An OSPF Autonomous System (AS) is an entire OSPF routing domain under the same technical administration.

An OSPF AS can be divided into multiple areas. The idea of using areas is to put a boundary on the explosion of link state updates. Flooding and SPF calculation on a router is limited to changes within an area. An area can be represented by either a single number or in dotted-decimal notation. All routers within an area have the exact link-state database. Area 0 is also known as the backbone area. All other areas must border the backbone area.

An area is interface-specific. An OSPF router can be a member of multiple areas (among which one area must be area 0). These routers are known as Area Border Routers (ABRs). Each ABR maintains a separate topological database for each area the router is in. Each topological database contains all of the LSA databases for each router within a given area. The routers within the same area have identical topological databases. The ABR is responsible for forwarding routing information or changes between its border areas.

An Autonomous System Boundary Router (ASBR) is a router that is running multiple routing protocols and serves as a gateway to OSPF routers in order to reach networks within other routing domains. The ASBR imports routes from different routing protocols into OSPF through a process known as redistribution. An ASBR may exist in either a normal area or a NSSA.

## OSPF LSA Types

Type 1: Router LSA: It is generated by each router for each area it belongs to. This type of LSA lists all local OSPF interfaces including cost. It is flooded within originating area.

Type 2: Network LSA: It is generated by DRs describing the set of routers attached to a particular network. It is flooded only within the originating area.

Type 3: Network Summary LSA: It is generated by ABRs describing inter-area routes. It is flooded to other areas.

Type 4: ASBR Summary LSA: It is generated by ABRs advertising the interface IP address of the ASBR. It is flooded into areas not connected to the ASBR.

Type 5: External LSA: It is generated by the ASBR describing networks external to the Autonomous System (AS) or Default Routes. It is flooded only to Normal Areas, not to Stub or NSSA.

Type 7: NSSA External LSA: It is generated by ASBR in a NSSA area, and is flooded only within the Not-So-Stubby area. It advertises an external destination or a default route. The ABR converts Type 7 LSAs into type 5 before flooding them into the backbone area and other normal areas.

## Link State Cost

Each interface on which OSPF is enabled has a cost associated with it. The Layer 3 Switch advertises its interfaces and their costs to OSPF neighbors. For example, if an interface has an OSPF cost of ten, the Layer 3 Switch advertises the interface with a cost of ten to other OSPF routers.

By default, an interface's OSPF cost is based on the port speed of the interface. The cost is calculated by dividing the reference bandwidth by the port speed. The default reference bandwidth is 100 Mbps, which results in the following default costs:

- 10 Mbps port – 10
- All other port speeds – 1 (If the resulting cost is less than 1, the software rounds the cost up to 1.)

You can change the reference bandwidth, to change the costs calculated by the software. The software uses the following formula to calculate the cost:

$$\text{Cost} = \text{reference-bandwidth}/\text{interface-speed}$$

For 10 Gbps OSPF interfaces, in order to differentiate the costs between 100 Mbps, 1000 Mbps, and 10,000 Mbps interfaces, you can set the auto-cost reference bandwidth to 10000, whereby each slower link is given a higher cost, as follows:

- 10 Mbps port's cost =  $10000/10 = 1000$
- 100 Mbps port's cost =  $10000/100 = 100$
- 1000 Mbps port's cost =  $10000/1000 = 10$
- 10000 Mbps port's cost =  $10000/10000 = 1$

The bandwidth for interfaces that consist of more than one physical port is calculated as follows:

- Trunk group: The combined bandwidth of all the ports.
- Virtual interface: The combined bandwidth of all the ports in the port-based VLAN that contains the virtual interface.

The default reference bandwidth is 100 Mbps. You can change the reference bandwidth to a value from 1 – 4294967. If a change to the reference bandwidth results in a cost change to an interface, the Layer 3 Switch sends a link-state update to update the costs of interfaces advertised by the Layer 3 Switch.

## OSPF Virtual Link

Virtual Link is used to link an area to the backbone through a transit area.

Rules for setting up Virtual Links:

1. Virtual links must be configured between two ABRs. It must be configured on both routers using router IDs.
2. The area through which the virtual link is configured, known as the transit area, must be a non-backbone normal area (must have full routing information).

All ABRs must have either a direct or indirect link to an OSPF backbone area (0.0.0.0 or 0). If an ABR does not have a physical link to a backbone area, you can configure a virtual link from the ABR to another router within the same area that has a physical connection to the backbone area.

The path for a virtual link is through an area shared by the neighbor ABR (router with a physical backbone connection) and the ABR requiring a logical connection to the backbone. Two parameters must be defined for all virtual links—transit area ID and neighbor router:

- The transit area ID represents the shared area of the two ABRs and serves as the connection point between the two routers. This number should match the area ID value.
- When assigned from the router interface requiring a logical connection, the neighbor router field is the router ID (IPv4 address) of the router that is physically connected to the backbone.

When assigned from the router interface with the physical connection, the neighbor router is the router ID (IPv4 address) of the router requiring a logical connection to the backbone.

When you establish an area virtual link, you must configure it on both of the routers (both ends of the virtual link). For example, imagine that ABR1 in areas 1 and 2 is cut off from the backbone area (area 0). To provide backbone access to ABR1, you can add a virtual link between ABR1 and ABR2 in area 1 using area 1 as a transit area. To configure the virtual link, you define the link on the router that is at each end of the link. No configuration for the virtual link is required on the routers in the transit area.

To define the virtual link on ABR1, enter the following command on ABR1:

```
FastIron(config-ospf6-router) #area 1 virtual-link 209.157.22.1
```

To define the virtual link on ABR2, enter the following command on ABR2:

```
FastIron(config-ospf6-router) #area 1 virtual-link 10.0.0.1
```

To display OSPF virtual link information, enter the following command at any CLI level:

```
FastIron#show ip ospf virtual-link
```

Virtual Links add a layer of complexity and troubleshooting difficulty to any internetwork. When two or more internetworks are merged, sufficient planning should take place beforehand so that no area is left without a direct link to the backbone. If a Virtual Link is configured, it should be used only as a temporary fix to an unavoidable topology problem.

## Redistribution

Redistribution must be enabled on routers configured to operate as ASBRs. For example, to enable redistribution of RIP and static IP routes into OSPF, enter the following commands:

```
FastIron(config) #router ospf
FastIron(config-ospf-router) #redistribution rip
FastIron(config-ospf-router) #redistribution static
```

## External Route Summarization

When the Layer 3 Switch is an OSPF Autonomous System Boundary Router (ASBR), you can configure it to advertise one external route as an aggregate for all redistributed routes that are covered by a specified address range.

When you configure an address range, the range takes effect immediately. All the imported routes are summarized according to the configured address range. Imported routes that have already been advertised and that fall within the range are flushed out of the AS and a single route corresponding to the range is advertised.

If a route that falls within a configured address range is imported by the Layer 3 Switch, no action is taken if the Layer 3 Switch has already advertised the aggregate route; otherwise the Layer 3 Switch advertises the aggregate route. If an imported route that falls within a configured address range is removed by the Layer 3 Switch, no action is taken if there are other imported routes that fall within the same address range; otherwise the aggregate route is flushed.

You can configure up to 32 address ranges. The Layer 3 Switch sets the forwarding address of the aggregate route to zero and sets the tag to zero. If you delete an address range, the advertised aggregate route is flushed and all imported routes that fall within the range are advertised individually.

If an external LSDB overflow condition occurs, all aggregate routes are flushed out of the AS, along with other external routes. When the Layer 3 Switch exits the external LSDB overflow condition, all the imported routes are summarized according to the configured address ranges.

To configure a summary address for OSPF routes, enter commands such as the following.

```
FastIron(config-ospf-router) #summary-address 10.1.0.0 255.255.0.0
```

The command in this example configures summary address 10.1.0.0, which includes addresses 10.1.1.0, 10.1.2.0, 10.1.3.0, and so on. For all of these networks, only the address 10.1.0.0 is advertised in external LSAs.

To display the configured summary addresses, enter the following command at any level of the CLI:

```
FastIron#show ip ospf config  
OSPF Redistribution Address Ranges currently defined:  
Range-Address Subnetmask  
1.0.0.0 255.0.0.0  
1.0.1.0 255.255.255.0  
1.0.2.0 255.255.255.0
```

## 4 - BGP Concepts

### EBGP Multihop

EBGP speakers are usually directly connected (i.e. over a WAN link). Sometimes they cannot be directly connected. In this special case, the `neighbor x.x.x.x ebgp-multihop [<num>]` command is used. `ebgp-multihop [<num>]` specifies that the neighbor is more than one hop away and that the session type with the neighbor is thus EBGP-multihop. This option is disabled by default. The `<num>` parameter specifies the TTL you are adding for the neighbor. You can specify a number from 0 – 255. The default is 0. If you leave the EBGP TTL value set to 0, the software uses the IP TTL value. Multihop is only used for EBGP, not for IBGP.

### Using Loopback Interfaces for IBGP Peering

Using a loopback interface to establish peers is commonly used with IBGP rather than EBGP. The loopback interface is used to ensure that the neighbor relationship stays up as long as there is IP connectivity between the two peers.

IBGP Neighbors may be located anywhere in the AS, even several hops away from one another if reachable via local IGP such as OSPF. There may be multiple physical paths between IBGP peers. By default, BGP will use the IP address of the physical interface as the source IP in the packets sent to the peer. If the physical interface goes down, even though there is another path to the peer, the packets can't be sent. By using the loopback interfaces to establish peers, even if one physical link goes down, the two peers may still be reachable via another physical link.

The neighbor router needs to tell BGP that it is using a loopback interface rather than a physical interface to initiate the BGP neighbor TCP connection. The BGP command `neighbor x.x.x.x update-source <ip-addr> | ethernet [<slotnum>/]<portnum> | loopback <num> | ve <num>` configures the router to communicate with the neighbor through a specified local interface.

### Multipath EBGP

By default, when BGP4 load sharing is enabled, both IBGP and EBGP paths are eligible for load sharing, while paths from different neighboring ASs are not eligible. You can change load sharing to apply only to IBGP or EBGP paths, or to support load sharing among paths from different neighboring ASs.

To enable load sharing of IBGP paths only, enter the following command at the BGP configuration level of the CLI:

```
FastIron(config-bgp-router) #multipath ibgp
```

To enable load sharing of EBGP paths only, enter the following command at the BGP configuration level of the CLI:

```
FastIron(config-bgp-router) #multipath ebgp
```

### Peer Group

A peer group is a set of BGP4 neighbors that share common parameters. Peer groups provide the following benefits:

- Simplified neighbor configuration: You can configure a set of neighbor parameters and then apply them to multiple neighbors. You do not need to configure the common parameters individually on each neighbor.
- Flash memory conservation: Using peer groups instead of individually configuring all the parameters for each neighbor requires fewer configuration commands in the startup-config file.

You can set all neighbor parameters in a peer group. When you add a neighbor to the peer group, the neighbor receives all the parameter settings you set in the group, except parameter values you have explicitly configured for the neighbor. Here is an example of configuring a peer group:

```
FastIron(config-bgp-router) #neighbor PeerGroup1 peer-group
FastIron(config-bgp-router) #neighbor PeerGroup1 description "EastCoast_Peers"
FastIron(config-bgp-router) #neighbor PeerGroup1 remote-as 100
FastIron(config-bgp-router) #neighbor PeerGroup1 distribute-list out 1
```

## *Route Reflector*

Normally, all the BGP routers within an AS are fully meshed. Each of the routers has an IBGP session with each of the other BGP routers in the AS. Each IBGP router thus has a route for each of its IBGP neighbors. For large ASs containing many IBGP routers, the IBGP route information in each of the fully-meshed IBGP routers can introduce too much administrative overhead.

To avoid this problem, you can hierarchically organize your IGP routers into clusters. A cluster is a group of IGP routers organized into route reflectors and route reflector clients. You configure the cluster by assigning a cluster ID on the route reflector and identifying the IGP neighbors that are members of that cluster. All the configuration for route reflection takes place on the route reflectors. The clients are unaware that they are members of a route reflection cluster. All members of the cluster must be in the same AS. The cluster ID can be any number from 0 – 4294967295. The default is the router ID, expressed as a 32-bit number. Note if the cluster contains more than one route reflector, you need to configure the same cluster ID on all the route reflectors in the cluster. The cluster ID helps route reflectors avoid loops within the cluster.

A route reflector is an IGP router configured to send BGP route information to all the clients (other BGP4 routers) within the cluster. Route reflection is enabled on all Brocade BGP4 routers by default but does not take effect unless you add route reflector clients to the router.

A route reflector client is an IGP router identified as a member of a cluster. You identify a router as a route reflector client on the router that is the route reflector, not on the client. The client itself requires no additional configuration. In fact, the client does not know that it is a route reflector client. The client just knows that it receives updates from its neighbors and does not know whether one or more of those neighbors are route reflectors.

## **Confederation**

A confederation is a BGP4 Autonomous System (AS) that has been subdivided into multiple, smaller ASs. Subdividing an AS into smaller ASs simplifies administration and reduces BGP-related traffic, thus reducing the complexity of the Interior Border Gateway Protocol (IBGP) mesh among the BGP routers in the AS.

Normally, all BGP routers within an AS must be fully meshed, so that each BGP router has interfaces to all the other BGP routers within the AS. This is feasible in smaller ASs but becomes unmanageable in ASs containing many BGP routers.

When you configure BGP routers into a confederation, all the routers within a sub-AS (a subdivision of the AS) use IBGP and must be fully meshed. However, routers use EBGP to communicate between different sub-ASs.

To configure a confederation, configure groups of BGP routers into sub-ASs. A sub-AS is simply an AS. The term “sub-AS” distinguishes ASs within a confederation from ASs that are not in a confederation. For the viewpoint of remote ASs, the confederation ID is the AS ID. Remote ASs do not know that the AS represents multiple sub-ASs with unique AS IDs.

You can use any valid AS numbers for the sub-ASs. If your AS is connected to the Internet, Brocade recommends that you use numbers from within the private AS range (64512 – 65535). These are private AS numbers and BGP4 routers do not propagate these AS numbers to the Internet.

## Next-Hop-Self

The BGP command `neighbor x.x.x.x next-hop-self` may be applied to an IBGP neighbor. `next-hop-self` specifies that the router should list itself as the next hop in updates sent to the specified neighbor, rather than letting the protocol choose the next hop. This option is disabled by default.

## The BGP Table

Each BGP router has a BGP table. It includes information such as the destination networks, the next hops, MED (Multi-Exit Discriminator, metric), Local Preference, Weight, and AS Path. The > indicates that the route is the best way to get to the destination network and hence is put into the routing table.

Here is an example output:

```
Total number of BGP Routes: 14
Status codes: s suppressed, d damped, h history, *valid, >best, i internal
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop      Metric LocPrf Weight Path
*>i161.19.7.192/26 161.19.7.5   0      100    25    100 11 i
* 161.19.7.192/26 161.19.8.5   50     200    0     100 100 11 i
*>i161.49.4.0/26   161.49.5.2   0      300    25     10 i
*>i161.49.4.64/26 161.49.5.2   0      300    25     10 i
*>i161.49.4.128/26 161.49.5.2   0      300    25     10 i
*>i161.49.4.192/26 161.49.5.2   0      300    25     10 i
*> 161.49.6.0/24  0.0.0.0      50     200   32768  i
*i 161.49.6.0/24  161.49.7.1   0      100    25     i
```

The command `show ip bgp route` displays similar information:

```
FastIron#show ip bgp route
Total number of BGP Routes: 5
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
H:HISTORY I:IBGP L:LOCAL M:MULTIPATH S:SUPPRESSED
Prefix          Next Hop      Metric      LocPrf      Weight      Status
1 0.0.0.0/0      10.1.0.2     0           100         0           BI
AS_PATH: 65001 4355 701 80
2 102.0.0.0/24  10.0.0.1     1           100         0           BI
AS_PATH: 65001 4355 1
3 104.0.0.0/24  10.1.0.2     0           100         0           BI
AS_PATH: 65001 4355 701 1 189
4 240.0.0.0/24  102.0.0.1    1           100         0           BI
```

```
AS_PATH: 65001 4355 3356 7170 1455
5 250.0.0.0/24      209.157.24.1    1          100         0          I
AS_PATH: 65001 4355 701
```

## *The Routing Table*

Each router has a routing table which contain the best route each destination network. Here is an example:

```
FastIron#show ip route
Total number of IP routes: 50834
B:BGP D:Directly-Connected O:OSPF R:RIP S:Static
Network Address      NetMask             Gateway             Port              Cost   Type
3.0.0.0              255.0.0.0          192.168.13.2       1/1              0     B
4.0.0.0              255.0.0.0          192.168.13.2       1/1              0     S
9.20.0.0             255.255.128.0     192.168.13.2       1/1              0     B
10.1.0.0             255.255.0.0       0.0.0.0            1/1              1     D
10.10.11.0           255.255.255.0     0.0.0.0            2/24             1     D
12.2.97.0            255.255.255.0     192.168.13.2       1/1              0     O
```

The `Type` column indicates from where the network has been learned.

## *AS Path*

A basic concept of BGP is the idea that each BGP packet will keep track of the Autonomous Systems (ASs) it crosses in the AS PATH. If a router sees its own AS number in the AS PATH, it drops the packet as it represents a loop.

- BGP attributes keep track of route-specific information such as AS path information and route origin. Attributes are used in filtering and choosing the best route. Next-HOP and AS-PATH are example of attributes.
- Routing loops are the most important addition to the creation of BGP was its ability to prevent routing loops by checking the AS PATH and dropping packets if a packet is received containing the same AS as packet is entering.

You may prepend the local AS numbers to the front of the route's AS-Path. By adding AS numbers to the AS-Path, you can cause the route to be less preferred when compared to other routes on the basis of the length of the AS-Path.

## *Useful Commands to Display Summary BGP Neighbor and Route Information*

- `show ip bgp summary`
- `show ip bgp config`
- `show ip bgp neighbor`
- `show ip bgp neighbor x.x.x.x`
- `show ip bgp neighbor x.x.x.x routes-summary`
- `show ip bgp neighbor x.x.x.x advertised-routes`
- `show ip bgp`
- `show ip bgp route`
- `show ip bgp route summary`

- `show ip bgp route best`
- `show ip bgp route unreachable`
- `show ip bgp flap-statistics`

## 5 - Advanced Layer 3 Concepts

### Default Route Origination

When the Brocade device is an OSPF Autonomous System Boundary Router (ASBR), you can configure it to automatically generate a default external route into an OSPF routing domain. This feature is called *default route origination* or *default information origination*.

By default, the Brocade device does not advertise the default route into the OSPF V3 domain. If you want the device to advertise the OSPF default route, you must explicitly enable default route origination.

When you enable OSPF default route origination, the device advertises a type 5 default route that is flooded throughout the AS (except stub areas). The device advertises the default route into OSPF even if OSPF route redistribution is not enabled, and even if the default route is learned through an IBGP neighbor. Note that the Brocade device does not advertise the OSPF default route, regardless of other configuration parameters, unless you explicitly enable default route origination.

For example, to create and advertise a default route with a metric of 2 and as a type 1 external route, enter the following command:

```
FastIron(config-ospf6-router) #default-information-originate always metric 2 metric-type type1
```

Syntax:

```
[no] default-information-originate [always] [metric <value>] [metric-type <type>]
```

The `always` keyword originates a default route regardless of whether the device has learned a default route. This option is disabled by default.

The `metric <value>` parameter specifies a metric for the default route. If this option is not used, the value of the `default-metric` command is used for the route. The `metric-type <type>` parameter specifies the external link type associated with the default route advertised into the OSPF routing domain. The `<type>` can be one of the following:

- Type 1 external route
- Type 2 external route

If you do not use this option, the default redistribution metric type is used for the route type.

### IP Route Selection and Administrative Distance

IP routers use a lookup mechanism. IP lookup is an important action in router that is to find the next hop of each incoming packet with a longest-prefix-match address in the routing table. In other words, when a router determines the path to a certain destination, it will first choose the entry with the longest prefix match.

There may be multiple entries learned from different sources for the same destination network and these entries have the same prefix length. These different sources may be BGP4, OSPF, RIP, static routes, and so on. The software compares the routes on the basis of each route's administrative distance. If the administrative distance of the paths is lower than the administrative distance of paths from other sources (such as static IP routes, RIP, or OSPF), the BGP4 paths are installed in the IP route table.

Table 3 lists the default administrative distances which are found on the Brocade router.

**TABLE 3 Default Administrative Distances**

Protocol	Cost
Directly connected	0 (this value is not configurable)
Static	1 (applies to all static routes, including default routes)
External BGP (eBGP)	20
OSPF	110
RIP	120
Internal BGP (iBGP)	200
Local BGP	200
Unknown	255 (the router will not use this route)

Lower administrative distances are preferred over higher distances. For example, if the router receives routes for the same network from OSPF and from RIP, the router will prefer the OSPF route by default.

### Route Redistribution

Redistribution is done on the router that runs multiple routing protocols. In other words, it is run on a border router that borders multiple routing domains (such as OSPF and RIP). Here is an example of redistributing RIP and static routes into OSPF:

```
FastIron(config)#router ospf
FastIron(config-ospf-router)#redistribution rip
FastIron(config-ospf-router)#redistribution static
```

Do not enable redistribution until you have configured the redistribution filters. Otherwise, you might accidentally overload the network with routes you did not intend to redistribute.

### IPv6

An IPv6 address is 128 bits and is composed of 8 fields of 16-bit hexadecimal values separated by colons (:). :




Figure 3: IPv6 Address Format

HHHH is a 16-bit hexadecimal value (0000 to FFFF), while H is a 4-bit hexadecimal value. The following is an example of an IPv6 address:

```
2001:0000:0000:0200:002D:D0FF:FE48:4672
```

Note that the IPv6 address includes hexadecimal fields of zeros. To make the address less cumbersome, you can do the following:

- Omit the leading zeros; for example, 2001:0:0:200:2D:D0FF:FE48:4672.

- Compress the successive groups of zeros at the beginning, middle, or end of an IPv6 address to two colons (::) once per address; for example, 2001::200:2D:D0FF:FE48:4672.
- When specifying an IPv6 address in a command syntax, keep the following in mind: You can use the two colons (::) only once in the address to represent the longest successive hexadecimal fields of zeros.
- The hexadecimal letters in IPv6 addresses are not case-sensitive.

## *Types of IPv6 Addresses*

- **Unicast:** An address for a single interface. A packet sent to a unicast address is delivered to the interface identified by the address. There are several types of unicast addresses: Aggregatable global address (prefix 2000::/3), Site-local address (prefix FEC0::/10), Link-local address (prefix FE80::/10), IPv4-compatible address (0:0:0:0:0:0:A.B.C.D), Loopback address (0:0:0:0:0:0:1 or ::1), and Unspecified address (0:0:0:0:0:0:0 or ::).
- **Multicast:** An address for a set of interfaces belonging to different nodes. Sending a packet to a multicast address results in the delivery of the packet to all interfaces in the set. A multicast address has a fixed prefix of FF00::/8 (1111 1111). The next 4 bits define the address as a permanent or temporary address. The next 4 bits define the scope of the address (node, link, site, organization, global).
- **Anycast:** An address for a set of interfaces belonging to different nodes. Sending a packet to an anycast address results in the delivery of the packet to the closest interface identified by the address.

## PIM Sparse Mode

Brocade devices support Protocol Independent Multicast (PIM) Sparse version 2. PIM Sparse provides multicasting that is especially suitable for widely distributed multicast environments. In a PIM Sparse network, a PIM Sparse router that is connected to a host that wants to receive information for a multicast group must explicitly send a join request on behalf of the receiving host.

PIM Sparse routers are organized into domains. A PIM Sparse domain is a contiguous set of routers that all implement PIM and are configured to operate within a common boundary.

## IGMP Snooping Overview

When a device processes a multicast packet, by default, the device broadcasts the packets to all ports except the incoming port of a VLAN. Packets are flooded by hardware without going to the CPU. This behavior causes some clients to receive unwanted traffic. IGMP snooping provides multicast containment by forwarding traffic to only the ports that have IGMP receivers for a specific multicast group (destination address). A device maintains the IGMP group membership information by processing the IGMP reports and leave messages, so traffic can be forwarded to ports receiving IGMP reports.

An IGMP device is responsible for broadcasting general queries periodically, and sending group queries when it receives a leave message, to confirm that none of the clients on the port still want specific traffic before removing the traffic from the port.

IGMP snooping is a layer 2 mechanism which prevents multicast flows from flooding to all switch ports on a VLAN. The switch examines the Layer 3 IGMP packets within a VLAN by listening to the conversation between the router and hosts. The switch learns at Layer 3 which port is signaling for or leaving a multicast group. The switch discovers which interfaces are connected to hosts interested in receiving this traffic. The switch adds

or removes the ports from the Layer 2 multicast forwarding group based on the IGMP message type. Multicast streams are sent to ports that explicitly request the flow. IGMP snooping reduces bandwidth consumption by avoiding flooding the entire VLAN. The IGMP snooping feature tracks which ports are attached to multicast-capable routers to help it manage the forwarding of IGMP membership reports.

## VRRP vs. VRRP-E

Most of the parameters and default values are the same for both VRRP and VRRP-E. Some of the differences are listed as follows:

- Protocol: The Virtual Router Redundancy Protocol (VRRP) based on RFC 2338 or VRRP-Extended, the Brocade enhanced implementation of VRRP.
- Virtual Router ID (VRID): The ID of the virtual router you are creating by configuring multiple routers to back up an IP interface. You must configure the same VRID on each router that you want to use to back up the address.
- Virtual Router IP address: This is the address you are backing up.
  - VRRP: The virtual router IP address must be a real IP address configured on the VRID interface on one of the VRRP routers. This router is the IP address Owner and is the default Master. This VIP address is reachable from hosts. If the owner fails and the backup router becomes the new Master, the VIP will no longer be pinged from hosts.
  - VRRP-E: The virtual router IP address must be in the same subnet as a real IP address configured on the VRRPE interface, but cannot be the same as a real IP address configured on the interface.
- VRID MAC address: The source MAC address in VRRP or VRRP-E packets sent from the VRID interface, and the destination for packets sent to the VRID:
  - VRRP: A virtual MAC address defined as 00-00-5e-00-01-<vrid>. The Master owns the Virtual MAC address.
  - VRRP-E: A virtual MAC address defined as 02-E0-52-<hash-value>-<vrid>, where <hash-value> is a two-octet hashed value for the IP address and <vrid> is the VRID.
- Authentication type: The type of authentication the VRRP or VRRP-E routers use to validate VRRP or VRRP-E packets. The authentication type must match the authentication type the VRID's port uses with other routing protocols such as OSPF:
  - No authentication: The interfaces do not use authentication. This is the VRRP default.
  - Simple: The interface uses a simple text-string as a password in packets sent on the interface. If the interface uses simple password authentication, the VRID configured on the interface must use the same authentication type and the same password.  
Note that MD5 is not supported by VRRP or VRRP-E.
- Router type: Whether the router is an Owner or a Backup.
  - Owner (VRRP only): The router on which the real IP address used by the VRID is configured. The Owner is always the router that has the real IP address used by the VRID. All other routers for the VRID are Backups.
  - Backup: Routers that can provide routing services for the VRID but do not have a real IP address matching the VRID. With VRRPE, all routers for the VRID are Backups.
- Backup priority: A numeric value that determines a Backup's preferability for becoming the Master for the VRID. During negotiation, the router with the highest priority becomes the Master.
  - VRRP: The Owner has the highest priority (255); other routers can have a priority from 3 – 254. The default value is 255 for the Owner; 100 for each Backup.

- VRRP-E: All routers are Backups and have the same priority by default. The default value is 100 for all Backups.

If two or more Backups are tied with the highest priority, the Backup interface with the highest IP address becomes the Master for the VRID.

- Track port: Another Layer 3 Switch port or virtual interface whose link status is tracked by the VRID's interface. If the link for a tracked interface goes down, the VRRP or VRRPE priority of the VRID interface is changed, causing the devices to renegotiate for Master.
- Track priority: A VRRP or VRRPE priority value assigned to the tracked ports. If a tracked port's link goes down, the VRID port's VRRP or VRRPE priority changes.
  - VRRP: The priority changes to the value of the tracked port's priority. The value is 2 by default.
  - VRRP-E: The VRID port's priority is reduced by the amount of the tracked port's priority.
- The value is 5 by default.
- Backup preempt mode: Prevents a Backup with a higher VRRP priority from taking control of the VRID from another Backup that has a lower priority but has already assumed control of the VRID. This mode is enabled by default.
- VRRP-E slow start timer: This feature causes a specified amount of time to elapse between the time the Master is restored and when it takes over from the Backup. This interval allows time for OSPF convergence when the Master is restored. The default is disabled.

## Layer 2 and Layer 3 Multicast Address Mapping

The multicast address range of 01-00-5E-00-00-00 to 01-00-5E-7F-FF-FF has been reserved for IP multicasting. As indicated in the figure below, the high order 25 bits of the 48-bit MAC address are fixed and the low order 23 bits are variable.

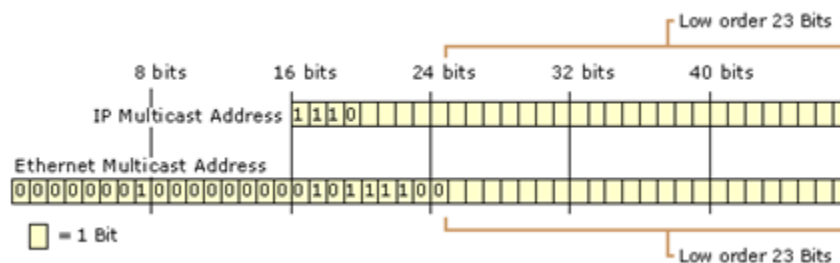


Figure 4: Multicast Address Layout

To map a Layer 3 IP multicast address to a Layer 2 MAC multicast address, the low order 23 bits of the IP multicast address are mapped directly to the low order 23 bits in the MAC multicast address. According to the class D convention, the first 4 bits of an IP multicast address are fixed. There are 5 bits in the IP multicast address that do not map to the MAC multicast address. Therefore, it is possible for a host to receive MAC multicast datagrams for groups to which it does not belong. These packets are dropped however by Layer 3 once the destination IP address has been determined.

For example, the multicast address 239.192.16.1 becomes 01-00-5E-40-10-01. To use the 23 low order bits, the first octet is not used, and only the last 7 bits of the second octet is used. The third and fourth octets are converted directly to hexadecimal numbers. The second octet, 192 in binary is 11000000. If you drop the high order bit, it becomes 1000000 or 64 (in decimal), or 0x40 (in hexadecimal). For the next octet, 16 in hexadecimal is 0x10. For the last octet, 1 in hexadecimal is 0x01. Therefore, the MAC address corresponding to 239.192.16.1 becomes 01-00-5E-40-10-01.

## PBR

Policy-Based Routing (PBR) allows you to use ACLs and route maps to selectively modify and route IP packets in hardware. The ACLs classify the traffic. Route maps that match on the ACLs set routing attributes for the traffic. A PBR policy specifies the next hop for traffic that matches the policy. Using standard ACLs with PBR, you can route IP packets based on their source IP address. With extended ACLs, you can route IP packets based on all of the clauses in the extended ACL.

You can configure the Brocade device to perform the following types of PBR based on a packet's Layer 3 and Layer 4 information:

- Select the next-hop gateway
- Send the packet to the null interface (null0)

When a PBR policy has multiple next hops to a destination, PBR selects the first live next hop specified in the policy that is up. If none of the policy's direct routes or next hops are available, the packet is routed in the normal way.

## 6 - Layer 2 Protocols

### MRP

MRP (Metro Ring Protocol) is a Brocade proprietary protocol that prevents Layer 2 loops and provides fast reconvergence in Layer 2 ring topologies. It is an alternative to STP and is especially useful in Metropolitan Area Networks (MANs). Here below shows an example of an MRP metro ring: (F: Forwarding, B: Blocking)

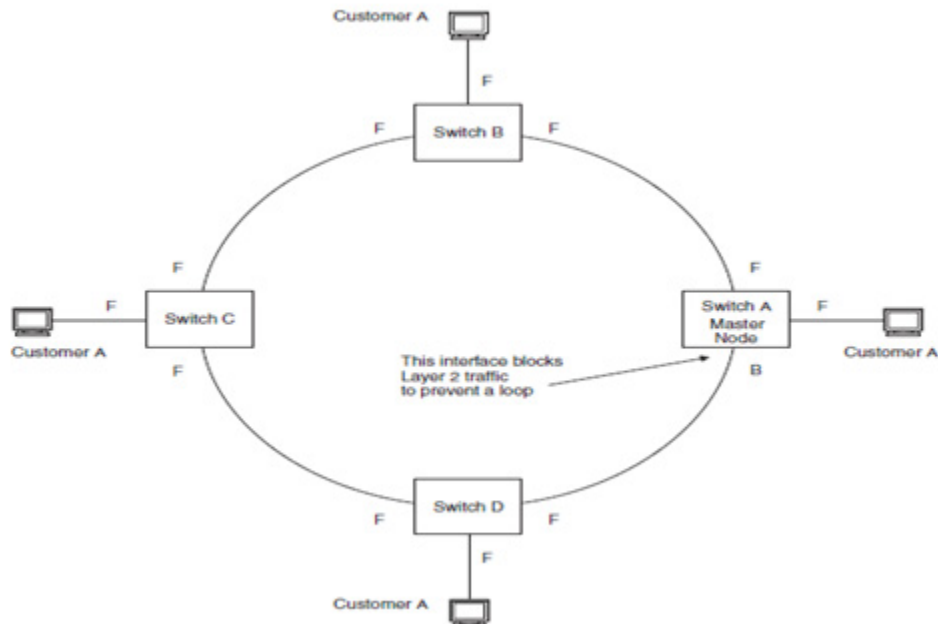



Figure 5: Metro Ring Example

The ring in this example consists of four MRP nodes (Brocade switches). Each node has two interfaces with the ring. Each node also is connected to a separate customer network. The nodes forward Layer 2 traffic to and from the customer networks through the ring. The ring interfaces are all in one port-based VLAN. Each customer interface can be in the same VLAN as the ring or in a separate VLAN.

One node is configured as the master node of the MRP ring. One of the two interfaces on the master node is configured as the primary interface; the other is the secondary interface. The primary interface originates Ring Health Packets (RHPs), which are used to monitor the health of the ring. An RHP is forwarded on the ring to the next interface until it reaches the secondary interface of the master node. The secondary interface blocks the packet to prevent Layer 2 loops. Metro Ring Protocol (MRP) was introduced in two phases: MRP Phase 1 and Phase 2.

MRP rings without shared interfaces (MRP Phase 1):

MRP Phase 1 allows you to configure multiple MRP rings, as shown below, but the rings cannot share the same link. For example, you cannot configure ring 1 and ring 2 to each have interfaces 1/1 and 1/2. Also, when you configure an MRP ring, any node on the ring can be designated as the master node for the ring. A master node can be the master node of more than one ring. Each ring is an independent ring and RHP packets are processed within each ring.



**Figure 6: Multiple Metro Rings**

In the above diagram, two nodes are each configured with two MRP rings. Any node in a ring can be the master for its ring. A node also can be the master for more than one ring.

In MRP Phase 1, a ring interface can have one of the following MRP states:

- Preforwarding (PF) – The interface can forward RHPS but cannot forward data. All ring ports begin in this state when you enable MRP.
- Forwarding (F) – The interface can forward data as well as RHPs. An interface changes from Preforwarding to Forwarding when the port's preforwarding time expires. This occurs if the port does not receive an RHP from the Master, or if the forwarding bit in the RHPs received by the port is off. This indicates a break in the ring. The port heals the ring by changing its state to Forwarding. The preforwarding time is the number of milliseconds the port will remain in the Preforwarding state before changing to the Forwarding state, even without receiving an RHP.
- Blocking (B) – The interface cannot forward data. Only the secondary interface on the Master node can be Blocking.

When MRP is enabled, all ports begin in the Preforwarding state. The primary interface on the Master node, although it is in the Preforwarding state like the other ports, immediately sends an RHP onto the ring. The secondary port on the Master node listens for the RHP.

- If the secondary port receives the RHP, all links in the ring are up and the port changes its state to Blocking. The primary port then sends another MRP with its forwarding bit set on. As each of the member ports receives the RHP, the ports changes their state to Forwarding. Typically, this occurs in sub-second time. The ring very quickly enters the fully initialized state.

- If the secondary port does not receive the RHP by the time the preforwarding time expires, a break has occurred in the ring. The port changes its state to Forwarding. The member ports also change their states from Preforwarding to Forwarding as their preforwarding timers expire. The ring is not intact, but data can still travel among the nodes using the links that are up.

MRP rings with shared interfaces (MRP Phase 2):

MRP rings can be configured to share the same interfaces as long as the interfaces belong to the same VLAN. Figure 7 displays two example diagrams of multiple MRP rings that share the same interface (Phase 2):




Figure 7: Multiple MRP Rings Sharing an Interface

To display ring information, enter the following command:

```
FastIron#show metro
```

```
Metro Ring 1
```

```
=====
```

Ring id	State	Ring role	Master vlan	Topo group	Hello time (ms)	Prefwring time (ms)
2	enabled	member	2	not conf	100	300

Ring interfaces Type	Interface	role	Forwarding state	Active interface	Interface
ethernet 1/1	primary	disabled	none		Regular
ethernet 1/2	secondary	forwarding	ethernet 2		Tunnel

RHPs sent	RHPs rcvd	TC RHPs rcvd	State changes
3	0	0	4

## RSTP Bridges and Bridge Port Roles

The 802.1W feature provides rapid traffic reconvergence for point-to-point links within a few milliseconds (0 – 500 milliseconds), following the failure of a bridge or bridge port. This reconvergence occurs more rapidly than the reconvergence provided by the 802.1D Spanning Tree Protocol or by RSTP Draft 3.

A bridge in an 802.1W rapid spanning tree topology is assigned as the root bridge if it has the highest priority (lowest bridge identifier) in the topology. Other bridges are referred to as non-root bridges. Unique roles are assigned to ports on the root and non-root bridges. Role assignments are based on the following information contained in the Rapid Spanning Tree Bridge Packet Data Unit (RST BPDU):

- Root bridge ID
- Path cost value
- Transmitting bridge ID
- Designated port ID

The 802.1W algorithm uses this information to determine if the RST BPDU received by a port is superior to the RST BPDU that the port transmits. The two values are compared in the order as given above, starting with the Root bridge ID. The RST BPDU with a lower value is considered superior. The superiority and inferiority of the RST BPDU is used to assign a role to a port.

If the value of the received RST BPDU is the same as that of the transmitted RST BPDU, then the port ID in the RST BPDUs are compared. The RST BPDU with the lower port ID is superior. Port roles are then calculated appropriately. The port's role is included in the BPDU that it transmits. The BPDU transmitted by an 802.1W port is referred to as an RST BPDU, while it is operating in 802.1W mode.

Ports can have one of the following roles:

- Root – Provides the lowest cost path to the root bridge from a specific bridge
- Designated – Provides the lowest cost path to the root bridge from a LAN to which it is connected
- Alternate – Provides an alternate path to the root bridge when the root port goes down
- Backup – Provides a backup to the LAN when the Designated port goes down
- Disabled – Has no role in the topology

Assignment of port roles: at system start-up, all 802.1W-enabled bridge ports assume a Designated role. Once start-up is complete, the 802.1W algorithm calculates the superiority or inferiority of the RST BPDU that is received and transmitted on a port.

On a root bridge, each port is assigned a *Designated port* role, except for ports on the same bridge that are physically connected together. In these types of ports, the port that receives the superior RST BPDU becomes the *Backup port*, while the other port becomes the *Designated Port*.

On non-root bridges, ports are assigned as follows:

- The port that receives the RST BPDU with the lowest path cost (based on link bandwidth) from the root bridge becomes the *Root port*.
- If two ports on the same bridge are physically connected, the port that receives the superior RST BPDU becomes the *Backup port*, while the other port becomes the *Designated port*.
- If a non-root bridge already has a Root port, then the port that receives an RST BPDU that is superior to those it can transmit becomes the *Alternate port*.
- If the RST BPDU that a port receives is inferior to the RST BPDUs it transmits, then the port becomes a *Designated port*.
- If the port is down or if 802.1W is disabled on the port, that port is given the role of *Disabled port*. Disabled ports have no role in the topology. However, if 802.1W is enabled on a port with a link down and the link of that port comes up, then that port assumes one of the following port roles: Root, Designated, Alternate, or Backup.

The switch with the lowest bridge ID (Bridge Priority + MAC address) becomes the Root Bridge. All ports on the Root Bridge should be forwarding. Each Non-Root bridge uses the following criteria to determine which ports should be forwarding or blocking:

1. Total Path Cost to the Root bridge
2. Lowest Sender's Bridge ID
3. Lower Sender's Port ID (Port Priority + Port Number)

To display a summary of 802.1W, use the `show 802-1w [vlan <vlan-id>]` command:

```
FastIron#show 802-1w
--- VLAN 1 [ STP Instance owned by VLAN 1 ] -----
VLAN 1 BPDU cam_index is 2 and the IGC and DMA master Are (HEX) 0 1 2 3

Bridge IEEE 802.1W Parameters:
Bridge Identifier hex 800000e080541700
Bridge MaxAge sec 20
Bridge Hello sec 2
Bridge Bridge FwdDly sec 15
Bridge Force Version Default
Bridge tx Hold cnt 3

RootBridgeID hex 800000e0804c9c00
RootPathCost 200000
Designated_BridgeID hex 800000e0804c9c00
Root_Port 1
MaxAge sec 20
FwdDly sec 15
FwdDly sec 2

Port IEEE 802.1W Parameters:
<--- Config Params ----->|<----- Current state ----->|
Port Pri PortPathCost P2P MAC Edge Port Role State Desig. Cost
Desig. bridge
1 128 200000 F F ROOT FORWARDING 0
800000e0804c9c00
2 128 200000 F F ALTERNATE DISCARDING 200000
800000e080548400
3 128 200000 F F DESIGNATED FORWARDING 200000
800000e080541700
4 128 200000 F F BACKUP DISCARDING 200000
800000e080541700
```

## DHCP Snooping

Dynamic Host Configuration Protocol (DHCP) snooping enables the Brocade device to filter untrusted DHCP packets in a subnet. DHCP snooping can ward off MiM attacks, such as a malicious user posing as a DHCP server sending false DHCP server reply packets with the intention of misdirecting other users. DHCP snooping can also stop unauthorized DHCP servers and prevent errors due to user mis-configuration of DHCP servers. Often DHCP snooping is used together with Dynamic ARP Inspection and IP Source Guard.

### How does DHCP Snooping Work?

When enabled on a VLAN, DHCP snooping stands between untrusted ports (those connected to host ports) and trusted ports (those connected to DHCP servers). A VLAN with DHCP snooping enabled forwards DHCP request packets from clients and discards DHCP server reply packets on untrusted ports, and it forwards DHCP server reply packets on trusted ports to DHCP clients.

## Spanning Tree Protocol (STP)

STP eliminates Layer 2 loops in networks, by selectively blocking some ports and allowing other ports to forward traffic, based on bridge and port parameters you can configure. Brocade Layer 2 and Layer 3 switches support standard STP as described in the IEEE 802.1D specification. STP is enabled by default on Layer 2 Switches but disabled by default on Layer 3 Switches. By default, each port-based VLAN on a Brocade device runs a separate spanning tree (a separate instance of STP). A Brocade device has one port-based VLAN (VLAN 1) by default that contains all the device's ports. Thus, by default each Brocade device has one spanning tree. However, if you configure additional port-based VLANs on a Brocade device, then each of those VLANs on which STP is enabled and VLAN 1 all run separate spanning trees.

With PVST, each VLAN has its own Spanning Tree instance. Each VLAN has its own root bridge. Ports blocked by one STP instance can be used by another STP instance to forward user traffic, hence achieving load balancing.

By default, each port-based VLAN on a Brocade device runs a separate spanning tree, which you can enable or disable on an individual VLAN basis. Alternatively, you can configure a Brocade device to run a single spanning tree across all ports and VLANs on the device. The Single STP feature (SSTP) is especially useful for connecting a Brocade device to third-party devices that run a single spanning tree in accordance with the 802.1Q specification.

The 802.1W RSTP provides rapid traffic reconvergence for point-to-point links within a few milliseconds (0 – 500 milliseconds), following the failure of a bridge or bridge port. This reconvergence occurs more rapidly than the reconvergence provided by the 802.1D Spanning Tree Protocol (STP).

Multiple Spanning Tree Protocol (MSTP), as defined in IEEE 802.1s, allows multiple VLANs to be managed by a single STP instance and supports per-VLAN STP. As a result, several VLANs can be mapped to a reduced number of spanning-tree instances. This ensures loop-free topology for one or more VLANs that have the similar layer-2 topology. The Brocade implementation supports up to 16 spanning tree instances in an MSTP enabled bridge which means that it can support up to 16 different Layer 2 topologies. The spanning tree algorithm used by MSTP is RSTP which provides quick convergence.

## BPDU Guard

In an STP environment, switches, end stations, and other Layer 2 devices use Bridge Protocol Data Units (BPDUs) to exchange information that STP will use to determine the best path for data flow. The BPDU guard, an enhancement to STP, removes a node that reflects BPDUs back in the network. It enforces the STP domain borders and keeps the active topology predictable by not allowing any network devices behind a BPDU guard-enabled port to participate in STP. In some instances, it is unnecessary for a connected device, such as an end station, to initiate or participate in an STP topology change. In this case, you can enable the STP BPDU guard feature on the Brocade port to which the end station is connected. STP BPDU guard shuts down the port and puts it into an errdisable state. This disables the connected device's ability to initiate or participate in an STP topology. A log message is then generated for a BPDU guard violation, and a CLI message is displayed to warn the network administrator of a severe invalid configuration. The BPDU guard feature provides a secure response to invalid configurations because the administrator must manually put the interface back in service if errdisable recovery is not enabled.

You enable STP BPDU guard on individual interfaces.(This feature is disabled by default.) Here is an example:

```
FastIron(config) interface ethe 2/1
FastIron(config-if-e1000-2/1) #stp-bpdu-guard
```

You can also use the multiple interface command to enable this feature on multiple ports at once. For example:

```
FastIron(config) #interface ethernet 1/1 to 1/9
FastIron(config-mif-1/1-1/9) #stp-bpdu-guard
```

## LLDP

Link layer discovery protocol (LLDP) is the Layer 2 network discovery protocol described in the IEEE 802.1AB standard, *Station and Media Access Control Connectivity Discovery*. This protocol enables a station to advertise its capabilities to, and to discover, other LLDP-enabled stations in the same 802 LAN segments.

LLDP enables a station attached to an IEEE 802 LAN/MAN to advertise its capabilities to, and to discover, other stations in the same 802 LAN segments. The information distributed by LLDP (the advertisement) is stored by the receiving device in a standard Management Information Base (MIB), accessible by a Network Management System (NMS) using a management protocol such as the Simple Network Management Protocol (SNMP). The information also can be viewed from the CLI, using `show lldp` commands.

## Dual-mode VLAN Ports

Configuring a tagged port as a dual-mode port allows it to accept and transmit both tagged traffic and untagged traffic at the same time. A dual-mode port accepts and transmits frames belonging to VLANs configured for the port, as well as frames belonging to the default VLAN (that is, untagged traffic).

For example, as indicated in the diagram below, port 2/11 is a dual-mode port belonging to VLAN 20. Traffic for VLAN 20, as well as traffic for the default VLAN, flows from a hub to this port. The dual-mode feature allows traffic for VLAN 20 and untagged traffic to go through the port at the same time.

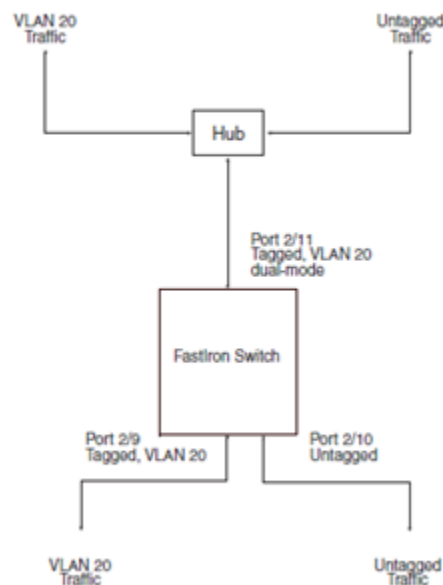


Figure 8: Dual-mode VLAN Ports

To enable the dual-mode feature on port 2/11 in the above example, enter the following commands:

```
FastIron(config)#vlan 20
FastIron(config-vlan-20)#tagged e 2/11
FastIron(config-vlan-20)#tagged e 2/9
FastIron(config-vlan-20)#int e 2/11
FastIron(config-if-e1000-2/11)#dual-mode
```

### *Specifying a default VLAN ID for a Dual-mode Port*

You can configure a dual-mode port to transmit traffic for a specified VLAN (other than the DEFAULT-VLAN) as untagged, while transmitting traffic for other VLANs as tagged, as indicated in the diagram:

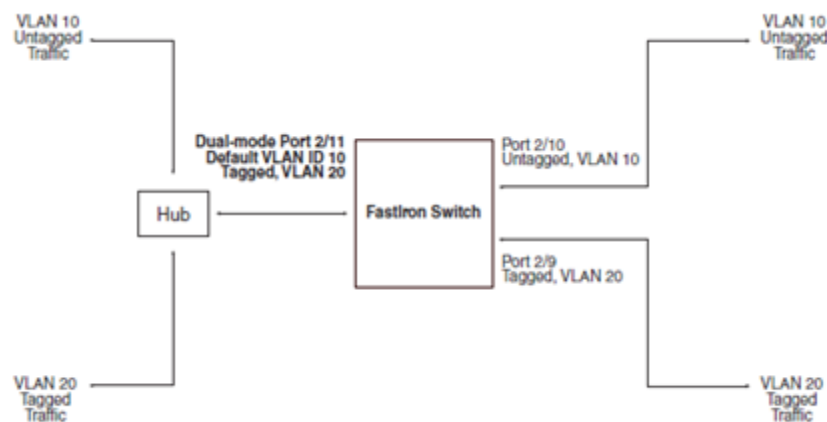


Figure 9: Dual-mode Port

Per the diagram, tagged port 2/11 is a dual-mode port belonging to VLANs 10 and 20. The default VLAN assigned to this dual-mode port is 10. This means that the port transmits tagged traffic on VLAN 20 (and all other VLANs to which the port belongs) and transmits untagged traffic on VLAN 10.

The dual-mode feature allows tagged traffic for VLAN 20 and untagged traffic for VLAN 10 to go through port 2/11 at the same time. A dual-mode port transmits only untagged traffic on its default VLAN (that is, either VLAN 1, or a user-specified VLAN ID), and only tagged traffic on all other VLANs.

The following commands configure VLANs 10 and 20 based on the diagram above. Tagged port 2/11 is added to VLANs 10 and 20, then designated a dual-mode port whose specified default VLAN is 10. In this configuration, port 2/11 transmits only untagged traffic on VLAN 10 and only tagged traffic on VLAN 20.

```
FastIron(config)#vlan 10 by port
FastIron(config-vlan-10)#untagged e 2/10
FastIron(config-vlan-10)#tagged e 2/11
FastIron(config-vlan-10)#exit
FastIron(config)#vlan 20 by port
FastIron(config-vlan-20)#tagged e 2/9
FastIron(config-vlan-20)#tagged e 2/11
FastIron(config-vlan-20)#exit
FastIron(config)#int e 2/11
FastIron(config-if-e1000-2/11)#dual-mode 10
```

Notes:

- If you do not specify a <vlan-id> in the dual mode command, the port's default VLAN is set to 1. The port transmits untagged traffic on the DEFAULT-VLAN.
- The dual-mode feature is disabled by default. Only tagged ports can be configured as dual-mode ports.
- In trunk group, either all of the ports must be dual-mode, or none of them can be. The `show vlan` command displays a separate row for dual-mode ports on each VLAN.

## Q-in-Q

Q-in-Q, defined in the IEEE 802.1Q in-Q, is a provider bridge extension in 802.1Q VLAN tag, also known as stackable VLANs. It enables service providers to use a single VLAN to support customers who have multiple VLANs. Q-in-Q allows service providers to offer IP-based services, including Metro-Ethernet in scalable implementations. Q-in-Q VLANs can also be used to provide multiple virtual connections and access to multiple services available over the Metro ISP.

802.1Q in-Q tagging provides finer granularity for configuring 802.1Q tagging, enabling you to configure 802.1Q tag-types on a group of ports. This feature allows you to create two identical 802.1Q tags (802.1Q in-Q tagging) on a single device. This is an example application with 802.1Q in-Q tagging:

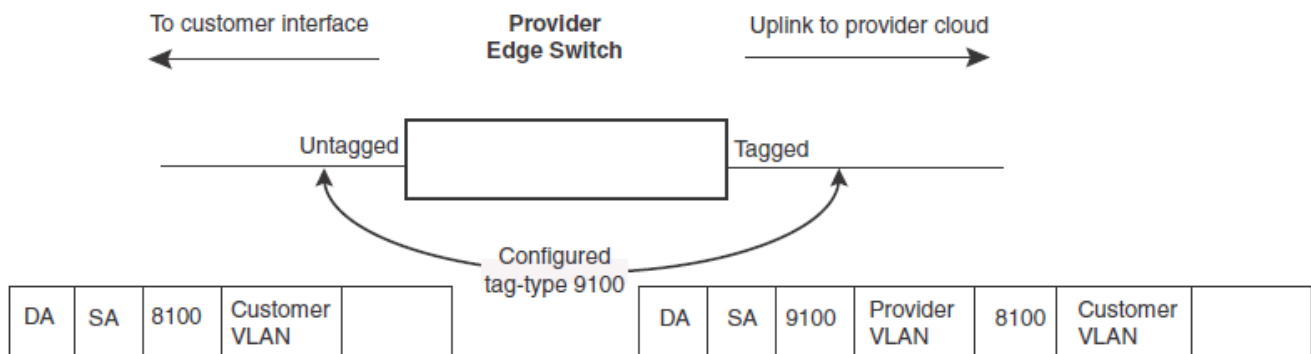


Figure 10: Q-in-Q Tagging

As shown in the diagram above, the untagged ports (to customer interfaces) accept frames that have any 802.1Q tag other than the configured tag-type 9100. These packets are considered untagged on this incoming port and are re-tagged when they are sent out of the uplink towards the provider. The 802.1Q tag-type on the uplink port is 8100, so the Brocade device will switch the frames to the uplink device with an additional 8100 tag, thereby supporting devices that only support this method of VLAN tagging.

For FGS, FLS, and FWS devices, the frame size is limited to 1522 (not 1530) bytes. To allow frames larger than 1522, you must enable jumbo frames. To globally enable jumbo support, enter commands such as the following:

```
FGS Switch(config)#jumbo
FGS Switch(config)#write memory
FGS Switch(config)#end
FGS Switch#reload
```

## VLAN

You can configure the following types of VLANs on FastIron devices:

- Layer 2 port-based VLAN - a set of physical ports that share a common, exclusive Layer 2 broadcast domain.
- Each port-based VLAN can contain either tagged or untagged ports. A port cannot be a member of more than one port-based VLAN unless the port is tagged. *802.1Q tagging* allows the port to add a four-byte tag field, which contains the VLAN ID, to each packet sent on the port.
- You also can configure port-based VLANs that span multiple devices by tagging the ports within the VLAN. The tag enables each device that receives the packet to determine the VLAN the packet belongs to. *802.1Q tagging* applies only to Layer 2 VLANs, not to Layer 3 VLANs. Here is a configuration example:

```
FastIron(config) #vlan 4
FastIron(config-vlan-4) #untag e 3 to 4
FastIron(config-vlan-4) #tagged e 5
FastIron(config-vlan-4) #exit
FastIron(config) #vlan 10
FastIron(config-vlan-4) #untag e 8 to 9
FastIron(config-vlan-4) #tagged e 5
```

- Layer 3 protocol VLANs - a subset of ports within a port-based VLAN that share a common, exclusive broadcast domain for Layer 3 broadcasts of the specified protocol type. If you want some or all of the ports within a port-based VLAN to be organized according to Layer 3 protocol, you must configure a Layer 3 protocol-based VLAN within the port-based VLAN. You can configure each of the following types of protocol-based VLAN within a port-based VLAN. All the ports in the Layer 3 VLAN must be in the same Layer 2 VLAN.
- IP subnet VLANs - a subset of ports in a port-based VLAN that share a common, exclusive subnet broadcast domain for a specified IP subnet
- IPv6 VLANs - a subset of ports in a port-based VLAN that share a common, exclusive network broadcast domain for IPv6 packets
- IPX network VLANs - a subset of ports in a port-based VLAN that share a common, exclusive network broadcast domain for a specified IPX network
- AppleTalk cable VLANs - a subset of ports in a port-based-based VLAN that share a common, exclusive network broadcast domain for a specified AppleTalk cable range

## ***Private VLAN***

A private VLAN is a VLAN that has the properties of standard Layer 2 port-based VLANs but also provides additional control over flooding packets on a VLAN. The figure below shows an example of an application using a private VLAN.

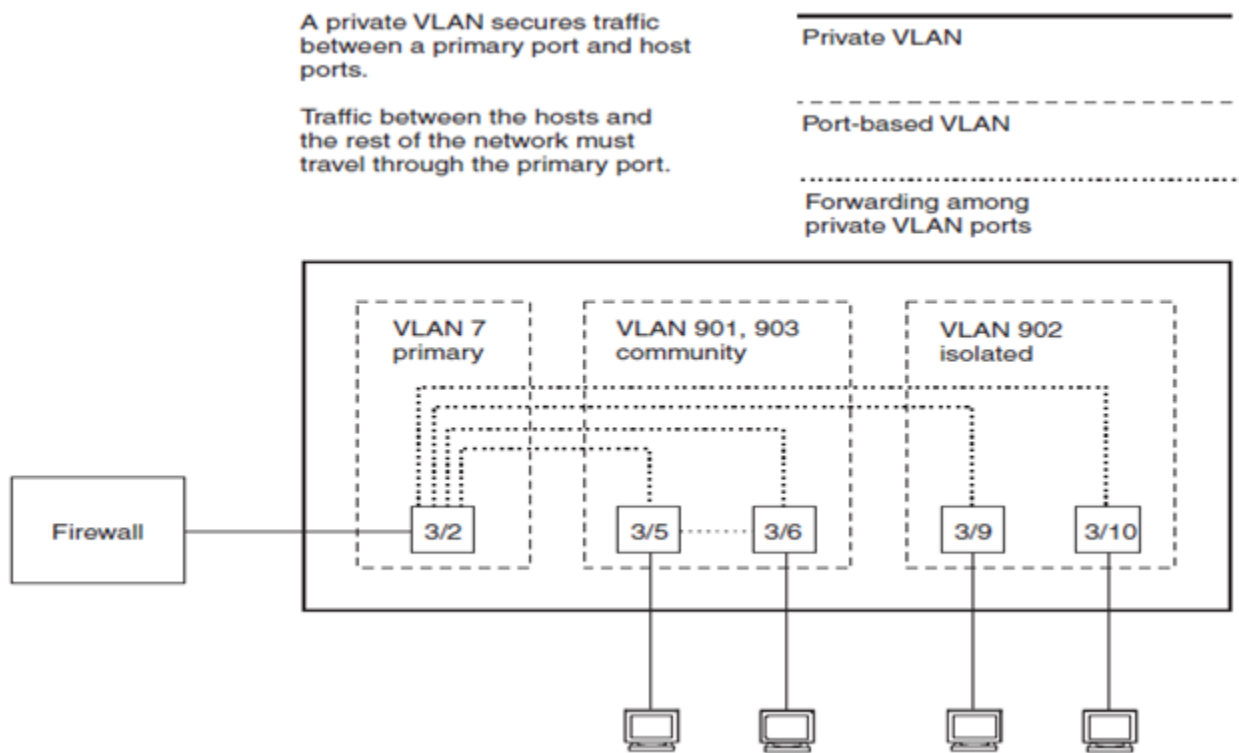


Figure 11: Private VLAN

This example uses a private VLAN to secure traffic between hosts and the rest of the network through a firewall. Five ports in this example are members of a private VLAN. The first port (port 3/2) is attached to a firewall. The next four ports (ports 3/5, 3/6, 3/9, and 3/10) are attached to hosts that rely on the firewall to secure traffic between the hosts and the rest of the network. In this example, two of the hosts (on ports 3/5 and 3/6) are in a community private VLAN, and thus can communicate with one another as well as through the firewall. The other two hosts (on ports 3/9 and 3/10), are in an isolated VLAN and thus can communicate only through the firewall. The two hosts are secured from communicating with one another even though they are in the same VLAN.

By default, the private VLAN does not forward broadcast or unknown-unicast packets from outside sources into the private VLAN. If needed, you can override this behavior for broadcast packets, unknown-unicast packets, or both.

You can configure a combination of the following types of private VLANs:

- Primary – The primary private VLAN ports are “promiscuous”. They can communicate with all the isolated private VLAN ports and community private VLAN ports in the isolated and community VLANs that are mapped to the promiscuous port.
- Isolated – Broadcasts and unknown unicasts received on isolated ports are sent only to the primary port. They are not flooded to other ports in the isolated VLAN.
- Community – Broadcasts and unknown unicasts received on community ports are sent to the primary port and also are flooded to the other ports in the community VLAN.

Each private VLAN must have a primary VLAN. The primary VLAN is the interface between the secured ports and the rest of the network. The private VLAN can have any combination of community and isolated VLANs.

Figure 4 compares private VLANs and standard port-based VLANs:

**TABLE 4 Private and Standard Port-based VLANs Comparison**

Forwarding Behavior	Private VLANs	Standard VLANs
All ports within a VLAN constitute a common Layer broadcast domain	No	Yes
Broadcasts and unknown unicasts are forwarded to all the VLAN's ports by default	No (isolated VLAN) Yes (community VLAN)	Yes
Known unicasts	Yes	Yes

## 7 - Monitoring, Maintenance, and Troubleshooting

### OSPF External Route Summarization

When the Layer 3 Switch is an OSPF Autonomous System Boundary Router (ASBR), you can configure it to advertise one external route as an aggregate for all redistributed routes that are covered by a specified address range. When you configure an address range, the range takes effect immediately. All the imported routes are summarized according to the configured address range. Imported routes that have already been advertised and that fall within the range are flushed out of the AS and a single route corresponding to the range is advertised.

If a route that falls within a configured address range is imported by the Layer 3 Switch, no action is taken if the Layer 3 Switch has already advertised the aggregate route; otherwise the Layer 3 Switch advertises the aggregate route. If an imported route that falls within a configured address range is removed by the Layer 3 Switch, no action is taken if there are other imported routes that fall within the same address range; otherwise the aggregate route is flushed. You can configure up to 32 address ranges. The Layer 3 Switch sets the forwarding address of the aggregate route to zero and sets the tag to zero.

If you delete an address range, the advertised aggregate route is flushed and all imported routes that fall within the range are advertised individually. If an external LSDB overflow condition occurs, all aggregate routes are flushed out of the AS, along with other external routes. When the Layer 3 Switch exits the external LSDB overflow condition, all the imported routes are summarized according to the configured address ranges.

Note that if you use redistribution filters in addition to address ranges, the Layer 3 Switch applies the redistribution filters to routes first, then applies them to the address ranges. Also note that if you disable redistribution, all the aggregate routes are flushed, along with other imported routes.

To configure a summary address for OSPF routes, enter commands such as the following:

```
FastIron(config-ospf-router) #summary-address 10.1.0.0 255.255.0.0
```

The command in this example configures summary address 10.1.0.0, which includes addresses 10.1.1.0, 10.1.2.0, 10.1.3.0, and so on. For all of these networks, only the address 10.1.0.0 is advertised in external LSAs.

### OSPF Internal Route Summarization

You may optionally assign a range for an area on the ABR. Ranges allow a specific IP address and mask to represent a range of IP addresses within an area, so that only that reference range address is advertised to the network, instead of all the addresses within that range. Each area can have up to 32 range addresses.

For example, to define an area range for subnets on 193.45.5.1 and 193.45.6.2, enter the following command on the ABR:

```
FastIron(config) #router ospf
FastIron(config-ospf-router) #area 1 range 193.45.0.0 255.255.0.0
```

Syntax: area <num> | <ip-addr> range <ip-addr> <ip-mask>

The `area <num> | <ip-addr>` parameter specifies the area number, whose internal specific networks will be summarized. The range `<ip-addr>` parameter specifies the IP address portion of the range. The software compares the address with the significant bits in the mask. All network addresses that match this comparison are summarized in a single route advertised by the router. The `<ip-mask>` parameter specifies the portions of the IP address that a route must contain to be summarized in the summary route. In the example above, all networks that begin with 193.45 are summarized into a single route.

## BGP Route Flap Dampening

A “route flap” is the change in a route’s state, from up to down or down to up. When a route’s state changes, the state change causes changes in the route tables of the routers that support the route. Frequent changes in a route’s state can cause Internet instability and add processing overhead to the routers that support the route.

Route flap dampening is a mechanism that reduces the impact of route flap by changing a BGP4 router’s response to route state changes. When route flap dampening is configured, the Layer 3 Switch suppresses unstable routes until the route’s state changes reduce enough to meet an acceptable degree of stability.

Route flap dampening is disabled by default. You can enable the feature globally or on an individual route basis using route maps. The Layer 3 Switch applies route flap dampening only to routes learned from EBGP neighbors.

The route flap dampening mechanism is based on penalties. When a route exceeds a configured penalty value, the Layer 3 Switch stops using that route and also stops advertising it to other routers. The mechanism also allows a route’s penalties to reduce over time if the route’s stability improves. The route flap dampening mechanism uses the following parameters:

- Suppression threshold
- Half-Time
- Reuse Threshold
- Maximum suppression time

To display route dampening statistics or all the dampened routes, enter the following command:

```
FastIron#show ip bgp flap-statistics
```

```
Total number of flapping routes: 414
```

```
Status Code >:best d:damped h:history *:valid
```

Network	From	Flaps	Since	Reuse	Path
h> 192.50.206.0/23	166.90.213.77	1	0 :0 :13	0 :0 :0	65001 4355 1 701
h> 203.255.192.0/20	166.90.213.77	1	0 :0 :13	0 :0 :0	65001 4355 1 7018
h> 203.252.165.0/24	166.90.213.77	1	0 :0 :13	0 :0 :0	65001 4355 1 7018
h> 192.50.208.0/23	166.90.213.77	1	0 :0 :13	0 :0 :0	65001 4355 1 701
h> 133.33.0.0/16	166.90.213.77	1	0 :0 :13	0 :0 :0	65001 4355 1 701
*> 204.17.220.0/24	166.90.213.77	1	0 :1 :4	0 :0 :0	65001 4355 701 62

The fields in the above output are explained as follows:

- `Total number of flapping routes`: Total number of routes in the Layer 3 Switch’s BGP4 route table that have changed state and thus have been marked as flapping routes.

- `Status code`: Indicates the dampening status of the route, which can be one of the following:
  - `>` - This is the best route among those in the BGP4 route table to the route's destination.
  - `d` - This route is currently dampened, and thus unusable.
  - `h` - The route has a history of flapping and is unreachable now.
  - `*` - The route has a history of flapping but is currently usable.
- `Network`: The destination network of the route.
- `From`: The neighbor that sent the route to the Layer 3 Switch.
- `Flaps`: The number of flaps (state changes) the route has experienced.
- `Since`: The amount of time since the first flap of this route.
- `Reuse`: The amount of time remaining until this route will be un-suppressed and thus be usable again.
- `Path`: Shows the AS-path information for the route.

## OSPF Network Types and DR and BDR Election

In a network that has multiple routers attached, OSPF elects one router to serve as the designated router (DR) and another router on the segment to act as the backup designated router (BDR). This arrangement minimizes the amount of repetitive information that is forwarded on the network by forwarding all messages to the designated router and backup designated routers responsible for forwarding the updates throughout the network.

In a network with no designated router and no backup designated router, the neighboring router with the highest priority is elected as the DR, and the router with the next largest priority is elected as the BDR. If the DR goes off-line, the BDR automatically becomes the DR. The router with the next highest priority becomes the new BDR. If two neighbors share the same priority, the router with the highest router ID is designated as the DR. The router with the next highest router ID is designated as the BDR. When multiple routers on the same network are declaring themselves as DRs, then both priority and router ID are used to select the designated router and backup designated routers. When only one router on the network claims the DR role despite neighboring routers with higher priorities or router IDs, this router remains the DR. This is also true for BDRs.

One important OSPF process is *Adjacency*. Adjacency occurs when a relationship is formed between neighboring routers for the purpose of exchanging routing information. Adjacent OSPF neighbor routers go beyond the simple Hello packet exchange; they exchange database information. In order to minimize the amount of information exchanged on a particular segment, one of the first steps in creating adjacency is to assign a Designated Router (DR) and a Backup Designated Router (BDR). The Designated Router ensures that there is a central point of contact, thereby improving convergence time within a multi-access segment.

In an OSPF point-to-point network, where a direct Layer 3 connection exists between a single pair of OSPF routers, there is no need for DR or BDR, as is the case in OSPF multi-access networks. Without the need for DR or BDR, a point-to-point network establishes adjacency and converges faster. The neighboring routers become adjacent whenever they can communicate directly. In contrast, in broadcast and non-broadcast multi-access (NBMA) networks, the DR and BDR become adjacent to all other routers attached to the network.

OSPF supports the following network types: Broadcast, Point-to-Point, Point-to-Multipoint, and NBMA.

## OSPF Interface: Passive and Ignore

To set an interface as OSPF passive or ignore, use the following command at the interface configuration context:

```
[no] ip address <ip-addr>/<mask-bits> [ospf-ignore | ospf-passive | secondary]
```

The `ospf-ignore` | `ospf-passive` parameters modify the Layer 3 Switch defaults for adjacency formation and interface advertisement. Use one of these parameters if you are configuring multiple IP subnet addresses on the interface but you want to prevent OSPF from running on some of the subnets:

- `ospf-passive` - This option disables adjacency formation with OSPF neighbors. By default, when OSPF is enabled on an interface, the software forms OSPF router adjacencies between each primary IP address on the interface and the OSPF neighbor attached to the interface.
- `ospf-ignore` - This option disables OSPF adjacency formation and also disables advertisement of the interface into OSPF. The subnet is completely ignored by OSPF.

When you configure an OSPF interface to be passive, that interface does not send or receive OSPF route updates. By default, all OSPF interfaces are active and thus can send and receive OSPF route information. Since a passive interface does not send or receive route information, the interface is in effect a stub network. Note that this option affects all IP subnets configured on the interface. The `ospf-passive` option disables adjacency formation but does not disable advertisement of the interface into OSPF. To disable advertisement in addition to disabling adjacency formation, you must use the `ospf-ignore` option.

## Dynamically Refreshing BGP Routes and Placing BGP Policy Changes Into Effect

To request a dynamic refresh of all routes from a neighbor, enter a command such as the following:

```
FastIron(config-bgp-router)#clear ip bgp neighbor 192.168.1.170 soft in
```

This command asks the neighbor to send its BGP4 table (Adj-RIB-Out) again. The Layer 3 Switch applies its filters to the incoming routes and adds, modifies, or removes BGP4 routes as necessary.

Syntax: `clear ip bgp neighbor all | <ip-addr> | <peer-group-name> | <as-num> [soft-outbound | soft[in|out]]`

The `all` | `<ip-addr>` | `<peer-group-name>` | `<as-num>` specifies the neighbor.

The `<ip-addr>` parameter specifies a neighbor by its IP interface with the Layer 3 Switch.

The `<peer-group-name>` specifies all neighbors in a specific peer group.

The `<as-num>` parameter specifies all neighbors within the specified AS. The `all` parameter specifies all neighbors.

The `soft-outbound` parameter updates all outbound routes by applying the new or changed filters, but sends only the existing routes affected by the new or changed filters to the neighbor.

The `soft [in | out]` parameter specifies whether you want to refresh the routes received from the neighbor or sent to the neighbor:

- `soft in` does one of the following:
  - If you enabled soft reconfiguration for the neighbor or peer group, `soft in` updates the routes by comparing the route policies against the route updates that the Layer 3 Switch has stored. Soft reconfiguration does not request additional updates from the neighbor or otherwise affect the session with the neighbor.

- If you did not enable soft reconfiguration, soft in requests the neighbor’s entire BGP4 route table (Adj-RIB-Out), then applies the filters to add, change, or exclude routes.
- If a neighbor does not support dynamic refresh, soft in resets the neighbor session.
- `soft out` updates all outbound routes, then sends the Layer 3 Switch’s entire BGP4 route table (Adj-RIB-Out) to the neighbor, after changing or excluding the routes affected by the filters. If you do not specify in or out, the Layer 3 Switch performs both options.

The `soft-outbound` parameter updates all outbound routes by applying the new or changed filters, but sends only the existing routes affected by the new or changed filters to the neighbor. The `soft out` parameter updates all outbound routes, then sends the Layer 3 Switch’s entire BGP4 route table (Adj-RIB-Out) to the neighbor, after changing or excluding the routes affected by the filters.

To dynamically resend all the Layer 3 Switch’s BGP4 routes to a neighbor, enter a command such as the following.

```
FastIron(config-bgp-router) #clear ip bgp neighbor 192.168.1.170 soft out
```

This command applies its filters for outgoing routes to the Layer 3 Switch’s BGP4 route table (Adj-RIB-Out), changes or excludes routes accordingly, then sends the resulting Adj-RIB-Out to the neighbor.

The Brocade Layer 3 Switch does not automatically update outbound routes using a new or changed outbound policy or filter when a session with the neighbor goes up or down. Instead, the Layer 3 Switch applies a new or changed policy or filter when a route is placed in the outbound queue (Adj-RIB-Out). To place a new or changed outbound policy or filter into effect, you must enter a `clear ip bgp neighbor` command regardless of whether the neighbor session is up or down. You can enter the command without optional parameters or with the `soft out` or `soft-outbound` option. Either way, you must specify a parameter for the neighbor (`<ip-addr>`, `<as-num>`, `<peer-group-name>`, or `all`).

To place policy changes into effect, enter a command such as the following:

```
FastIron(config-bgp-router) #clear ip bgp neighbor 10.10.200.102 soft in
```

This command does not tear down the BGP peer session. It updates the routes by comparing the route policies against the route updates that the Layer 3 Switch has stored. The command does not request additional updates from the neighbor or otherwise affect the session with the neighbor. If you do not specify “in”, the command applies to both inbound and outbound updates. The following sections describe how to dynamically refresh BGP4 routes to place new or changed filters into effect. If you need to tear down and re-establish the BGP session and resend all BGP routes between neighbors, use the hard clear `clear ip bgp neighbor x.x.x.x` or `clear ip bgp neighbor all` command.

## Allocating Memory for More VLANs or Virtual Routing Interfaces

Brocade Layer 2 and Layer 3 Switches support up to 4095 VLANs. In addition, Layer 3 switches support up to 512 virtual routing interfaces. The number of VLANs and virtual routing interfaces supported on your product depends on the device and, for Chassis devices, the amount of DRAM on the management module. [Figure 5](#) displays the default and configurable maximum numbers of VLANs and virtual routing interfaces for Layer 2 and Layer 3 switches:

**TABLE 5 VLAN and Virtual Routing Interface Maximums**

VLANs		Virtual Routing Interfaces	
Default Maximum	Configurable Maximum	Default Maximum	Configurable Maximum
64	4094	255	512

You may increase the number of VLANs you can configure. Although you can specify up to 4095 VLANs, you can configure only 4094 VLANs. VLAN ID 4094 is reserved for use by the Single Spanning Tree feature. To increase the maximum number of VLANs you can configure, enter the following commands:

```
FastIron(config) #system-max vlan 2048
FastIron#reload
```

To increase the maximum number of virtual routing interfaces you can configure, enter the following commands:

```
FastIron(config) #system-max virtual-interface 512
FastIron#reload
```

## Specify Types of OSPF Syslog Messages to Log

You can specify which kinds of OSPF-related Syslog messages are logged. By default, the only OSPF messages that are logged are those indicating possible system errors. If you want other kinds of OSPF messages to be logged, you can configure the Brocade device to log them. For example, to specify that all OSPF-related Syslog messages be logged, enter the following commands:

```
FastIron(config) #router ospf
FastIron(config-ospf-router) #log all
```

Syntax: [no] log all | adjacency | bad\_packet [checksum] | database | memory | retransmit

The log command has the following options:

- The `all` option causes all OSPF-related Syslog messages to be logged. If you later disable this option with the `no log all` command, the OSPF logging options return to their default settings.
- The `adjacency` option logs essential OSPF neighbor state changes, especially on error cases. This option is disabled by default.
- The `bad_packet checksum` option logs all OSPF packets that have checksum errors. This option is enabled by default.
- The `bad_packet` option logs all other bad OSPF packets. This option is disabled by default.
- The `database` option logs OSPF LSA-related information. This option is disabled by default.
- The `memory` option logs abnormal OSPF memory usage. This option is enabled by default.
- The `retransmit` option logs OSPF retransmission activities. This option is disabled by default.

## Port Mirroring and Monitoring

Port mirroring is a method of monitoring network traffic that forwards a copy of each incoming or outgoing packet from one port on a network switch to another port where the packet can be analyzed. Port mirroring may be used as a diagnostic tool or debugging feature, especially for preventing attacks. Port mirroring can be managed locally or remotely.

Configure port mirroring by assigning a port from which to copy all packets, and a “mirror” port where the copies of the packets are sent (also known as the monitor port). A packet received on, or issued from, the first port is forwarded to the second port as well. Attach a protocol analyzer on the mirror port to monitor each segment separately. The analyzer captures and evaluates the data without affecting the client on the original port. The mirror port may be a port on the same switch with an attached RMON probe, a port on a different switch in the same hub, or the switch processor.

To configure port monitoring on an individual port on a Brocade device, enter commands similar to the following.

```
FastIron(config) #mirror-port ethernet 1/2/4
FastIron(config) #interface ethernet 1/2/11
FastIron(config-if-e1000-11) #monitor ethernet 1/2/4 both
```

Traffic on port e 1/2/11 will be monitored, and the monitored traffic will be copied to port e 1/2/4, the mirror port.

Syntax: [no] mirror-port ethernet [<stack-unit>/<slotnum>/]<portnum> [input | output]

Syntax: [no] monitor ethernet [<stack-unit>/<slotnum>/]<portnum> both | in | out

- The <portnum> parameter for mirror-port ethernet specifies the port to which the monitored traffic will be copied.
- The <portnum> parameter for monitor ethernet specifies the port on which traffic will be monitored.
- The input and output parameters configure the mirror port exclusively for ingress or egress traffic. If you do not specify one, both types of traffic apply.
- The both, in, and out parameters specify the traffic direction you want to monitor on the mirror port. There is no default.

To display the port monitoring configuration, enter the show monitor and show mirror commands. You may also configure ACL-based inbound mirroring, MAC filter-based mirroring, and VLAN-based mirroring.

## SNMP

SNMP is the protocol developed to manage nodes (servers, workstations, routers, switches, etc.) on an IP network. Currently, there are three versions of SNMP defined: SNMP v1, v2, and v3. SNMP is a set of protocols for managing complex networks. SNMP sends messages, called protocol data units (PDUs), to different parts of a network. SNMP-compliant devices, called agents, store data about themselves in Management Information Bases (MIBs) and return this data to the SNMP requesters. In order to secure access to management functions, you may use ACLs, or restrict SNMP access to a specific IP address or VLAN. You may also disable SNMP access on switches as well.

Here is an example of using ACL to control SNMP access:

```
FastIron(config) #access-list 25 deny host 209.157.22.98 log
FastIron(config) #access-list 25 deny 209.157.23.0 0.0.0.255 log
FastIron(config) #access-list 25 deny 209.157.24.0 0.0.0.255 log
FastIron(config) #access-list 25 permit any
FastIron(config) #access-list 30 deny 209.157.25.0 0.0.0.255 log
FastIron(config) #access-list 30 deny 209.157.26.0/24 log
FastIron(config) #access-list 30 permit any
FastIron(config) #snmp-server community public ro 25
FastIron(config) #snmp-server community private rw 30
```

Syntax: snmp-server community <string> ro | rw <num>

- The <string> parameter specifies the SNMP community string the user must enter to gain SNMP access.
- The ro parameter indicates that the community string is for read-only (“get”) access.

- The `rw` parameter indicates the community string is for read-write (“set”) access.
- The `<num>` parameter specifies the number of a standard ACL and must be from 1 - 99. These commands configure ACLs 25 and 30, then apply the ACLs to community strings. ACL 25 is used to control read-only access using the `public` community string. ACL 30 is used to control read-write access using the `private` community string.
- When `snmp-server community` is configured, all incoming SNMP packets are validated first by their community strings and then by their bound ACLs.

SNMP Version 3 (SNMPv3) adds security and remote configuration capabilities to the previous versions. Using SNMPv3, users can securely collect management information from their SNMP agents without fear that the data has been tampered with. Also, confidential information, such as SNMP set packets that change a device's configuration, can be encrypted to prevent their contents from being exposed on the wire. Also, the group-based administrative model allows different users to access the same SNMP agent with varying access privileges. The SNMPv3 architecture introduces the User-based Security Model (USM) for message security and the View-based Access Control Model (VACM) for access control. The architecture supports the concurrent use of different security, access control, and message processing models such as: Security, authentication and privacy, authorization and access control, Administrative Framework, naming of entities, people and policies, usernames and key management, notification destinations, proxy relationships, and remotely configurable via SNMP operations.

SNMPv3 also introduces the ability to dynamically configure the SNMP agent using SNMP SET commands against the MIB objects that represent the agent's configuration. This dynamic configuration support enables addition, deletion, and modification of configuration entries either locally or remotely.

## Recovering from a Lost Password

Recovery from a lost password requires direct access to the serial port and a system reset. Follow the steps given below to recover from a lost password:

1. Start a CLI session over the serial interface to the device.
2. Reboot the device.
3. At the initial boot prompt at system startup, enter `b` to enter the boot monitor mode.
4. Enter `no password` at the prompt. (You cannot abbreviate this command.) This command will cause the device to bypass the system password check.
5. Enter `boot system flash primary` at the prompt.
6. After the console prompt reappears, assign a new password.

