IBM iFlow Director® Version 2.1

User’s Guide

IBM iFlow Director® Version 2.1

User’s Guide

Note: Before using this information and the product it supports, read the general information in the *Safety information and Environmental Notices and User Guide* documents on the IBM *Documentation* CD and the *Warranty Information* document that comes with the product.

First Edition (July 2012)

Contents

Preface

The *IBM iFlow Director 2.1 User's Guide* describes how to configure and use IBM iFlow Director firmware.

Who Should Use This Book

This guide is intended for network installers and system administrators engaged in configuring and maintaining a network. The administrator should be familiar with Ethernet concepts, IP addressing, load balancing, and Access Control List (ACL) filtering.

How This Book Is Organized

Chapter 1, "The IBM iFlow Director," provides an overview of iFlow Director features.

Chapter 2, "Layer 2 Mode Hashing," describes how iFlow Director provides transparent, "bump in the wire," traffic direction and load balancing.

Chapter 3, "Layer 3 Mode Hashing," describes how iFlow Director provides a Layer 3 hashing for server load balancing with a variety of routing options and high-availability.

Chapter 4, "Command Reference," describes the CLI menus and commands that are specific to iFlow Director.

Chapter 5, "ISCLI Reference," describes the ISCLI commands that are specific to iFlow Director.

Index," includes pointers to the description of the key words used throughout the book.

Typographic Conventions

The following table describes the typographic styles used in this book.

Typeface or Symbol	Meaning	Example
`AaBbCc123`	This type is used for names of commands, files, and directories used within the text. It also depicts on-screen computer output and prompts.	View the `readme.txt` file. `Main#`
`AaBbCc123`	This bold type appears in command examples. It shows text that must be typed in exactly as shown.	`Main#` **`sys`**
<AaBbCc123>	This italicized type appears in command examples as a parameter placeholder. Replace the indicated text with the appropriate real name or value when using the command. Do not type the brackets. This also shows book titles, special terms, or words to be emphasized.	To establish a Telnet session, enter: `host#` **`telnet`** *<IP address>* Read your *User's Guide* thoroughly.

Typeface or Symbol	**Meaning**	**Example**
`[ ]`	Command items shown inside brackets are optional and can be used or excluded as the situation demands. Do not type the brackets.	`host# ls [-a]`
AaBbCc123	This block type depicts menus, buttons, and other controls that appear in Web browsers and other graphical interfaces.	Click the Save button.

Chapter 1. The IBM iFlow Director

The following topics are discussed in this chapter:

Introducing the IBM iFlow Director

IBM iFlow Director is a high performance, low-latency, integrated 10Gb Ethernet flow-balancing switch with policy-based traffic steering capability. It delivers high-availability, scalability, and lower cost of ownership for appliance vendors to offer BladeCenter-based solutions for applications such as Security Gateways, Wireless Gateways, Lawful Interception & Network Surveillance, and Traffic Management & Service Differentiation.

With ten 10Gb uplink ports and fourteen 10Gb internal blade-server ports, each iFlow Director provides 100Gbps of non-blocking throughput. The system throughput can be scaled by adding blade servers to the IBM BladeCenter H or HT chassis.

iFlow Director high-availability features such as server health checks, failover, and more, coupled with industry-leading BladeCenter hardware innovations (redundant mid-plane, redundant power domains, redundant management modules for platform monitoring) improve the overall system reliability for Enterprise, Data Center, and Carrier Grade deployments.

iFlow Director Applications

iFlow Director can be configured to provide the following solutions:

- Layer 2 “bump-in-the-wire” hashing mode
- Layer 3 hashing mode

Layer 2 Mode

The IBM BladeCenter H or HT chassis with IBM iFlow Director can be deployed in a transparent in-line Layer 2 mode (bump in the wire) or in a Layer 3 mode.

Typically, a transparent in-line mode deployment uses a dual iFlow Director configuration, with blade servers sandwiched by the two switches. One switch is the front-end that connects to the Internet, and one switch is the back-end that connects to internal resources.

Servers between the switches are placed into service application groups, binding together servers that perform the same function. ACL filters on the switches steer traffic to each application sequentially.

For details on this solution, see “Layer 2 Mode Hashing” on page 5.

Layer 3 Mode

Network applications such as caches, firewalls, Intrusion Prevention Systems, Content Gateways, and other Deep Packet Inspection (DPI) devices are usually hosted on application-specific devices. These single-purpose appliances process traffic in-line at the network perimeter and within the core networks. To meet capacity requirements, an iFlow Director helps consolidate the function of stand-alone appliances, load-balancers, and Ethernet switches into an IBM BladeCenter chassis.

The iFlow Director steers traffic from the client network and the Internet to the application blades. When corresponding traffic from the applications blades comes back to the switch, the iFlow Director routes the traffic to the appropriate client destination or Internet destination using either standard routing (via static routes or OSPF) or policy-based routing.

For redundancy, dual iFlow Directors may be deployed using VRRP where one switch is active and the other is a backup, taking over in case of a component or link failure.

For details on this solution, see “Layer 3 Mode Hashing” on page 35.

Unlocking iFlow Director

You need a software license key to unlock the iFlow Director feature set. The iFlow license key is uniquely tied to the serial number of the IBM iFlow Director and is not transferable to another IBM iFlow Director. The license key grants you permission to access all the iFlow Director functions and features on ports purchased for your IBM iFlow Director.

Two demo license keys are available for the purpose of evaluating iFlow Director. The demo license grants you permission to access functions and features of iFlow Director for a predetermined period of time. After the demo license expires, your iFlow Director configurations are erased and the iFlow functions and features are disabled until a license key is entered.

To unlock iFlow Director, use the following command:

```
>> # /oper/swkey/key/enakey
```

When prompted, enter the software feature and the license key:

```
Enter The Software Feature: ibmiflow
Enter License Key: <key code>
```

Chapter 2. Layer 2 Mode Hashing

The following topics are discussed in this chapter:

Layer 2 Hashing Features

iFlow Director uses hash-based load balancing to distribute network traffic across 1 to 14 blade servers at wire-speed. Hundreds of policies consisting of ingress port, Layer 2, Layer 3-4 IPv4 header fields can be configured to steer matching traffic toward a port, trunk, or server load-balancing group.

The system throughput can be scaled by adding more blade servers and up to four switches running iFlow Director.

iFlow Director includes high-availability features, such as:

- Server health-checks (link, ping, ARP, HTTP, IPS, TCP, TCP-script, etc.) for rapid failure detection.
- Failover options to allow distributing flows from a failed blade to other active or backup blades.
- Traffic redirection through alternate ports (for example, fail-open bypass) if an application group fails.
- Uplink Failure Detection and controlled failover with NIC teaming.

These features, coupled with industry-leading BladeCenter hardware innovations (redundant mid-plane, redundant power domains, redundant management modules for platform monitoring) improve the overall system reliability for Enterprise, Data Center and Carrier Grade deployments.

How Layer 2 Mode Hashing Works

Network applications such as firewalls, Intrusion Prevention Systems, Content Gateways, and other Deep Packet Inspection (DPI) devices are usually hosted on application-specific boxes. These single-purpose appliances process traffic in-line at the network perimeter and within the core networks. To meet capacity requirements, stand-alone load-balancer boxes are typically deployed along with these appliances and Layer 2 switches, creating an "appliance sprawl." iFlow Director allows you to consolidate the stand-alone appliances, load-balancers and Ethernet switching into an IBM BladeCenter chassis, as shown in Figure 1.

Figure 1. Enhanced Application Network Integration with iFlow Director

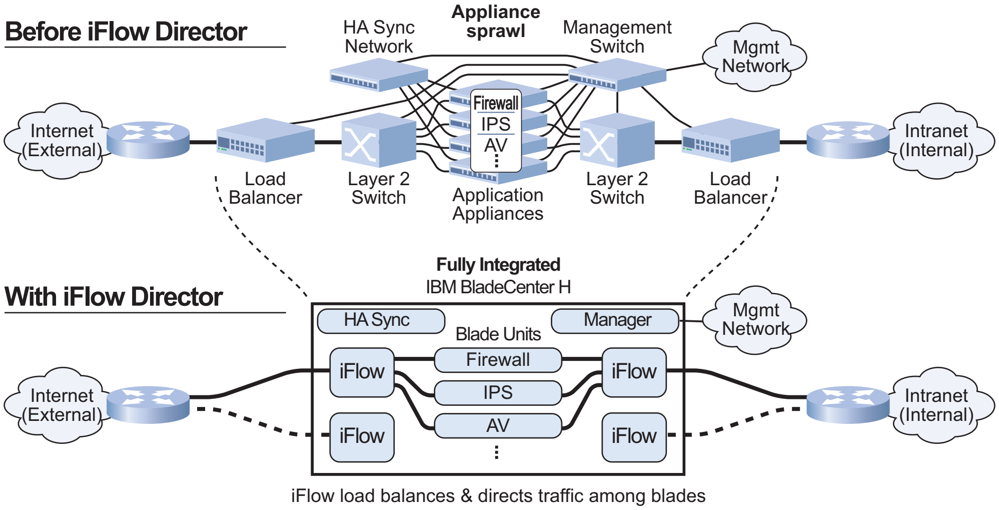

Depending on the application requirements, the IBM BladeCenter H or HT chassis with iFlow Director can be deployed in a transparent in-line mode (bump in the wire), or in a Layer 2 mode. Typically, a transparent in-line mode deployment uses a dual iFlow Director configuration, with blade servers sandwiched by the two switches. One switch is the front-end that connects to the Internet, and one switch is the back-end that connects to internal resources.

Service application groups bind together servers that perform the same function. ACL filters steer traffic to each application sequentially.

An Example of the Layer 2 Hashing Process

A typical iFlow Director scenario works as follows:

1. Traffic from the Internet enters the switch through an external port.
2. An ACL filter redirects the traffic to a pre-defined application group of blade servers.
3. Traffic is load-balanced among the servers in the application group using a hash algorithm. Depending on the traffic type (Layer 2 or Layer 3), the user-configurable hash algorithm may be based on one of the following:
   - Source IP (SIP)
   - Destination IP (DIP)
   - Source MAC (SMAC)
   - Destination MAC (DMAC)
   - SIP + DIP
   - SMAC + DMAC

   See "Selecting the Load-Balancing Metric" on page 22 for details.
4. When the traffic leaves the application group, the switch performs Layer 2 switching or ACL redirection, and forwards the traffic to the Intranet.

   The switch acts as a bump in the wire. The switch maintains flow persistency by correctly sending the traffic between the Internet and the Intranet. Flows established among the same source and destination addresses (IP or MAC) are sent to the same server in the application group.

   If a server fails, existing flows from the failed server are processed by other servers in the application group.

Creating SLB Application Groups

An *application group* is an aggregation of individual blade servers running a common application, such as a firewall or Intrusion Detection System. Incoming traffic is directed to the application group, and load-balanced among the servers in the group. Server Load Balancing (SLB) reduces the processing load on each individual server and provides fault tolerance.

An application group that contains eight servers or fewer is called a normal application group, and an application group that contains more than eight servers is called a *jumbo application group*.

Two modes of load-balancing can be configured in iFlow Director:

- **Basic** SLB mode supports one or more normal application groups in a switch with basic load-balancing that scales optimally for 1, 2, 4, or 8 servers in a group.
- **Expanded** SLB mode supports a single application group (either normal or jumbo) in a switch with enhanced load-balancing that scales optimally from 1 to 14 servers.

The configuration in Figure 2 consists of a single application group that contains eight servers. This network design can apply to both basic application groups (see page 11) and expanded application groups (see page 13).

Traffic from PC 11 ingresses Switch 1 port EXT8 from the external network, and its destination is PC 17 on Switch 2 port EXT8 on the internal network.

An ACL redirects the incoming traffic to application group 1. One server in the application group is selected to handle the traffic, and propagates it to Switch 2. When Switch 2 receives the traffic it will be forwarded via Layer 2 to its destination (PC 17).

Persistency can be maintained by setting the trunk hash to use both SIP and DIP. Both request and response traffic is handled by the same server.

Figure 2. Single Application Group with Server Load Balancing

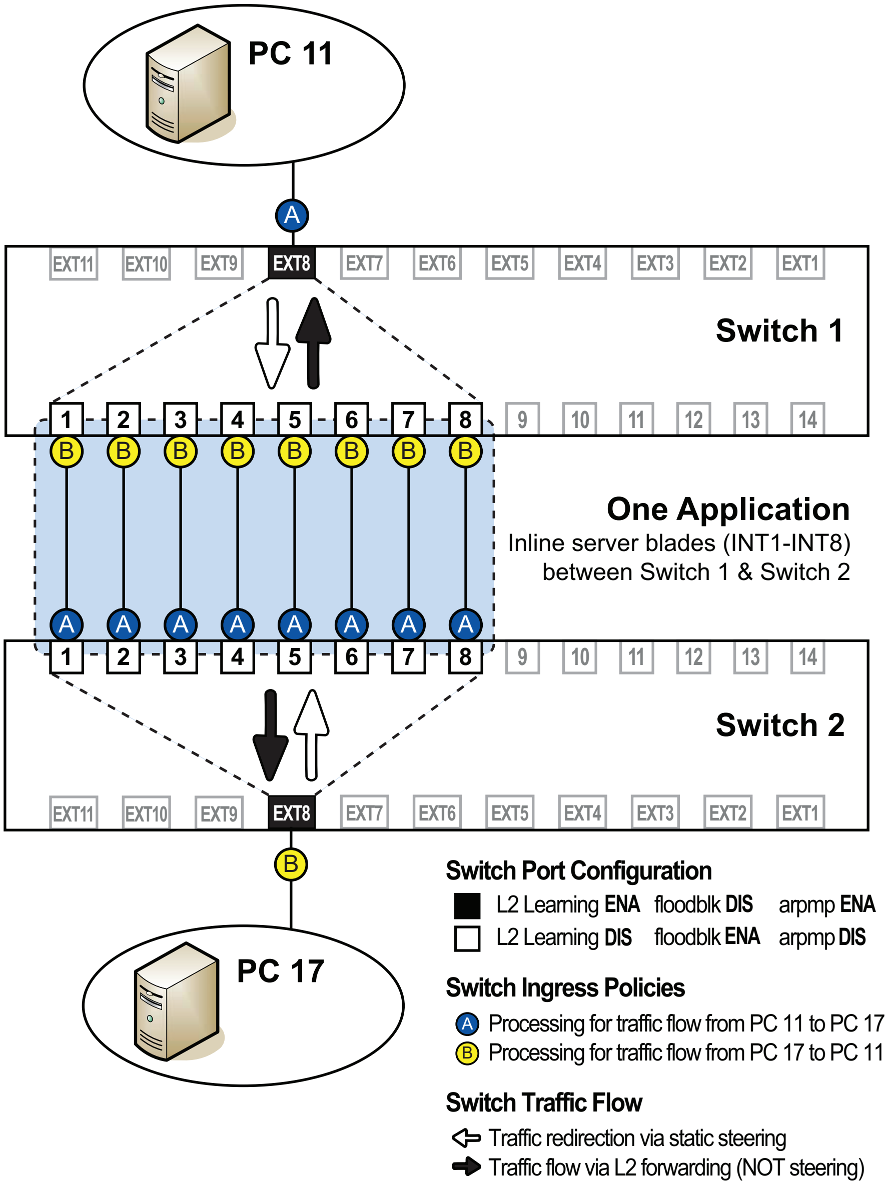
PC 11
A
EXT11
EXT10
EXT9
EXT8
EXT7
EXT6
EXT5
EXT4
EXT3
EXT2
EXT1
Switch 1
1
2
3
4
5
6
7
8
9
10
11
12
13
14
B
One Application
Inline server blades (INT1-INT8)
between Switch 1 & Switch 2
A
Switch 2
B
PC 17
Switch Port Configuration
■ L2 Learning ENA floodblk DIS arpmp ENA
□ L2 Learning DIS floodblk ENA arpmp DIS
Switch Ingress Policies
A Processing for traffic flow from PC 11 to PC 17
B Processing for traffic flow from PC 17 to PC 11
Switch Traffic Flow
⇦ Traffic redirection via static steering
➡ Traffic flow via L2 forwarding (NOT steering)

Basic Application Groups

The following procedure demonstrates how to configure a basic SLB application group, as shown in Figure 2.

1. Define port processing behavior, as shown in the following example.

```
/cfg/port INT1
     learn dis
     arpmp dis
     floodblk ena
/cfg/port INT2
     learn dis
     arpmp dis
     floodblk ena
/cfg/port INT3
     learn dis
     arpmp dis
     floodblk ena
/cfg/port INT4
     learn dis
     arpmp dis
     floodblk ena
/cfg/port INT5
     learn dis
     arpmp dis
     floodblk ena
/cfg/port INT6
     learn dis
     arpmp dis
     floodblk ena
/cfg/port INT7
     learn dis
     arpmp dis
     floodblk ena
/cfg/port INT8
     learn dis
     arpmp dis
     floodblk ena
```

2. Enable SLB

```
/cfg/slb/on
```

3. Select the trunk hashing method.

```
/cfg/slb/method trunk
```

Note: The ECMP method is the default setting and need not be explicitly configured unless switch settings have been previously altered. The trunk method (Layer 2 hashing) and ECMP method (Layer 3 hashing) are mutually exclusive. When the `trunk` method is specified, Layer 3 hashing features (page 35) are not supported.

4. Create the primary servers.

```
/cfg/slb/real 1
     ena
     port INT1
/cfg/slb/real 2
     ena
     port INT2
/cfg/slb/real 3
     ena
     port INT3
/cfg/slb/real 4
     ena
     port INT4
/cfg/slb/real 5
     ena
     port INT5
/cfg/slb/real 6
     ena
     port INT6
/cfg/slb/real 7
     ena
     port INT7
/cfg/slb/real 8
     ena
     port INT8
```

5. Define an SLB group.

```
/cfg/slb/group 1
     add 1
     add 2
     add 3
     add 4
     add 5
     add 6
     add 7
     add 8
```

6. Define an SLB application group.

   Create an application manager for each SLB group and enable the error handling remap option. Define the mode as `inline` for blades that pass traffic across dual NIC interfaces.

```
/cfg/slb/app 1
     group 1
     mode inline
     errhand
     remap ena
```

7. Configure a trunk group to service the SLB application.

```
/cfg/l2/trunk 1
     ena
     slbapp 1
```

8. Define an ACL redirection filter for the application group.

```
/cfg/acl/acl 1/target
     dest slbapp
     slbapp 1
/cfg/acl/acl 1
     action redirect
     stats ena
```

9. Define an ACL group to service the redirection filter.

```
/cfg/acl/group 1
     add 1
```

10. Apply the ACL group to the uplink port to steer traffic to the application group.

    Note: The switch prevents you from directly adding an ACL filter for SLB directly to any switch port. This operation must always be done indirectly via an ACL group.

```
/cfg/port EXT8/acl/add grp 1
```

11. Set the hash metric to use SIP and DIP for Layer 3 traffic.

```
/cfg/l2/thash/l3thash
     sip enable
     dip enable
```

Expanded Application Groups

An expanded application group provides better load distribution. Eight trunks are aggregated to expand the number of hash buckets to 64. Use the expand command (`/c/slb/app` *x*`/expand e`) to provide expanded trunk hashing to an application that contains eight or fewer real servers.

An expanded SLB application group requires eight trunk groups. If there are fewer than eight trunk groups to support the application, the switch displays the following message when you apply the configuration:

```
Warning: Expanded SLB application 1 must use 8 trunks.
```

However, if you use SNMP or the browser-based interface (BBI) to configure the application, no warning message is displayed on the switch console but a syslog message is generated. When a switch configuration with fewer than 8 trunk groups is imported via TFTP or FTP in ISCLI, the warning message is logged in the syslog.

The following steps demonstrate how to configure a single expanded application group, using eight trunk groups.

Perform the following steps to configure the example shown in Figure 2.

1. Define port processing behavior, as shown in the following example.

```
/cfg/port INT1
     learn dis
     arpmp dis
     floodblk ena
/cfg/port INT2
     learn dis
     arpmp dis
     floodblk ena
/cfg/port INT3
     learn dis
     arpmp dis
     floodblk ena
/cfg/port INT4
     learn dis
     arpmp dis
     floodblk ena
/cfg/port INT5
     learn dis
     arpmp dis
     floodblk ena
/cfg/port INT6
     learn dis
     arpmp dis
     floodblk ena
/cfg/port INT7
     learn dis
     arpmp dis
     floodblk ena
/cfg/port INT8
     learn dis
     arpmp dis
     floodblk ena
```

2. Enable SLB.

```
/cfg/slb/on
```

3. Select the trunk hashing method.

```
/cfg/slb/method trunk
```

Note: The `trunk` setting (Layer 2 hashing) and `ecmp` setting (Layer 3 hashing) are mutually exclusive. When `trunk` method is specified, Layer 3 hashing features (page 35) are not supported.

4. Define SLB real servers.

```
/cfg/slb/real 1/ena
     port INT1
/cfg/slb/real 2/ena
     port INT2
/cfg/slb/real 3/ena
     port INT3
/cfg/slb/real 4/ena
     port INT4
/cfg/slb/real 5/ena
     port INT5
/cfg/slb/real 6/ena
     port INT6
/cfg/slb/real 7/ena
     port INT7
/cfg/slb/real 8/ena
     port INT8
```

5. Define an SLB group.

```
/cfg/slb/group 1
     add 1
     add 2
     add 3
     add 4
     add 5
     add 6
     add 7
     add 8
```

6. Define an SLB application group.

   Create an application manager for each SLB group and enable the error handling remap option. Define the mode as `inline` for blades that pass traffic across dual NIC interfaces. Enable expanded mode.

```
/cfg/slb/app 1
     group 1
     mode inline
/cfg/slb/app 1/errhand
     remap ena
     expand ena
```

7. Configure 8 trunk groups to service the expanded SLB application.

```
/cfg/l2/trunk 1
     ena
     slbapp 1
/cfg/l2/trunk 2
     ena
     slbapp 1
/cfg/l2/trunk 3
     ena
     slbapp 1
/cfg/l2/trunk 4
     ena
     slbapp 1
/cfg/l2/trunk 5
     ena
     slbapp 1
/cfg/l2/trunk 6
     ena
     slbapp 1
/cfg/l2/trunk 7
     ena
     slbapp 1
/cfg/l2/trunk 8
     ena
     slbapp 1
```

8. Define an ACL redirection filter to perform server load balancing for the application group.

```
/cfg/acl/acl 1/target
     dest slbapp
     slbapp 1
/cfg/acl/acl 1
     action redirect
     stats ena
```

9. Define an ACL group to service the redirection filter.

```
/cfg/acl/group 1
     add 1
```

10. Apply the ACL group to the uplink port to steer traffic to the application group.

    Note: The switch prevents you from directly adding an ACL filter for SLB directly to any switch port. This operation must always be done indirectly via an ACL group.

```
/cfg/port EXT8/acl/add grp 1
```

11. Set the hash metric to use SIP and DIP for Layer 3 traffic.

```
/cfg/l2/thash/l3thash
     sip enable
     dip enable
```

Using Manual OSP Distribution

Manual OSP Distribution (MOD) allows an external application to control the trunk method hash distribution and synchronization of application ports by manually assigning switch ports to the application hash buckets.

Note: The ECMP hash method is not supported with MOD.

SLB application hash buckets can be controlled via Chassis Internal Network (CIN) by using a management application running SNMP.

Use the following command to enable Manual OSP Distribution:

```
/cfg/slb/mod ena
```

When MOD is enabled:

- All SLB health checks are disabled.
- All application hash buckets are initialized to an empty state.

The switch performs additional checks and operations as follows:

- Prevents SLB real server operations commands to be executed.
- SLB health check and Application Management configuration options are ignored (no operational effect).
- Any SLB-related configuration changes that occur during runtime force all application hash buckets to be re-initialized to the default empty state.

The following sections demonstrate how to operationally control the application hash buckets manually:

- “Using the CLI to Control the Hash Buckets” on page 18
- “Using SNMP to Control the Hash Buckets” on page 20

Using the CLI to Control the Hash Buckets

The following example shows the general steps used to operationally control the MOD hash buckets.

1. When the `mod` command is enabled, application hash buckets are initialized to an empty state:

```
/oper/slb/app 1/cur
 Current Application state:

SLB APP 1: Cur Application Distribution Matrix table
           <empty>
```

2. Set one bucket to a real sever port:

```
/oper/slb/app 1/bucket 1/port INT1
```

   Or set all 64 buckets to a real server port:

```
/oper/slb/app 1/bucket 1-64/port INT1
```

   No actual runtime operation occurs until you issue the `update` command.

```
/oper/slb/app 1/cur
Current Application state:

SLB APP 1: Cur Application Distribution Matrix table
           <empty>

SLB APP 1: New Application Distribution Matrix table

1) bucket: 1      2      3      4      5      6      7      8
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
2) bucket: 9      10     11     12     13     14     15     16
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
3) bucket: 17     18     19     20     21     22     23     24
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
4) bucket: 25     26     27     28     29     30     31     32
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
5) bucket: 33     34     35     36     37     38     39     40
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
6) bucket: 41     42     43     44     45     46     47     48
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
7) bucket: 49     50     51     52     53     54     55     56
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
8) bucket: 57     58     59     60     61     62     63     64
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
```

3. When you are satisfied with the new bucket changes, use the `update` command to operationally commit the changes to the switch.

```
/oper/slb/app 1/update
```

4. You can edit changes to specific buckets (optional).

```
/oper/slb/app 1/bucket 5/port INT5
/oper/slb/app 1/bucket 27/port INT7
/oper/slb/app 1/cur
Current Application state:
SLB APP 1: Cur Application Distribution Matrix table

1) bucket: 1      2      3      4      5      6      7      8
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
2) bucket: 9      10     11     12     13     14     15     16
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
3) bucket: 17     18     19     20     21     22     23     24
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
4) bucket: 25     26     27     28     29     30     31     32
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
5) bucket: 33     34     35     36     37     38     39     40
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
6) bucket: 41     42     43     44     45     46     47     48
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
7) bucket: 49     50     51     52     53     54     55     56
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
8) bucket: 57     58     59     60     61     62     63     64
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1

SLB APP 1: New Application Distribution Matrix table

1) bucket: 1      2      3      4      5      6      7      8
           INT1   INT1   INT1   INT1   INT5   INT1   INT1   INT1
2) bucket: 9      10     11     12     13     14     15     16
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
3) bucket: 17     18     19     20     21     22     23     24
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
4) bucket: 25     26     27     28     29     30     31     32
           INT1   INT1   INT7   INT1   INT1   INT1   INT1   INT1
5) bucket: 33     34     35     36     37     38     39     40
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
6) bucket: 41     42     43     44     45     46     47     48
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
7) bucket: 49     50     51     52     53     54     55     56
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
8) bucket: 57     58     59     60     61     62     63     64
           INT1   INT1   INT1   INT1   INT1   INT1   INT1   INT1
```

The changes do not take affect until you use the `update` command to operationally commit the changes to the switch.

5. Use the following command to reset all outstanding changes that have not yet been updated to the switch (optional).

```
/oper/slb/app 1/reset
```

6. Use the following command to clear all buckets (optional).

```
/oper/slb/app 1/clear
```

No actual runtime operation occurs until you issue the update command to operationally commit the changes to the switch.

Note: If the application distribution matrix table is empty, all incoming packets will be discarded at the ingress port which has an ACL for redirection to an empty application distribution matrix. However, ACL statistics still show all the packet counts for all 8 trunk groups, because these packets match the ACLs at the ingress port.

Using SNMP to Control the Hash Buckets

The following example shows the general steps used to operationally control the MOD hash buckets.

1. When the `mod` command is enabled, application hash buckets are initialized to an empty state:

```
  GET: slbCurOperApplicationBucketString.1
  OID: .1.3.6.1.4.1.26543.2.5.10.4.2.1.5.1
Value: 0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,
       0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0
```

2. Set one bucket to a real sever port:

```
  SET: slbNewOperApplicationBucketPort.1.1
  OID: .1.3.6.1.4.1.26543.2.5.10.4.4.1.3.1.1
Value: 1
```

Or set all 64 buckets to a real server port:

```
  SET: slbNewOperApplicationBucketString.1
  OID: .1.3.6.1.4.1.26543.2.5.10.4.2.1.6.1
Value: 1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,
       1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1
```

No actual runtime operation occurs until you issue the `update` command.

```
  GET: slbCurOperApplicationBucketString.1
  OID: .1.3.6.1.4.1.26543.2.5.10.4.2.1.5.1
Value: 0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,
       0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0
```

3. When you are satisfied with the new bucket changes, use the `update` command to operationally commit the changes to the switch.

```
  SET: slbOperApplicationUpdate.1
  OID: .1.3.6.1.4.1.26543.2.5.10.4.2.1.2.1
Value: exec(2)

  GET: slbCurOperApplicationBucketString.1
  OID: .1.3.6.1.4.1.26543.2.5.10.4.2.1.5.1
 Value: INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,
        INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,
        INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,
        INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,
        INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,
        INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,
        INT1,INT1,INT1,INT1
```

4. You can edit changes to specific buckets (optional).

```
  SET: slbNewOperApplicationBucketPort.1.5
  OID: .1.3.6.1.4.1.26543.2.5.10.4.4.1.3.1.5
Value: 5

  SET: slbNewOperApplicationBucketPort.1.27
  OID: .1.3.6.1.4.1.26543.2.5.10.4.4.1.3.1.27
Value: 7
```

The following display shows the edited matrix table.

```
  GET: slbNewOperApplicationBucketString.1
  OID: .1.3.6.1.4.1.26543.2.5.10.4.2.1.6.1
Value: INT1,INT1,INT1,INT1,INT5,INT1,INT1,INT1,INT1,INT1,
       INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,
       INT1,INT1,INT1,INT1,INT1,INT1,INT7,INT1,INT1,INT1,
       INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,
       INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,
       INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,INT1,
       INT1,INT1,INT1,INT1
```

The changes do not take affect until you use the `update` command to operationally commit the changes to the switch.

5. Use the following command to reset all outstanding changes that have not yet been updated to the switch (optional).

```
  SET: slbOperApplicationReset.1
  OID: .1.3.6.1.4.1.26543.2.5.10.4.2.1.3.1
Value: exec(2)
```

6. Use the following command to clear all the buckets (optional).

```
  SET: slbOperApplicationClear.1
  OID: .1.3.6.1.4.1.26543.2.5.10.4.2.1.4.1
Value: exec(2)
```

No actual runtime operation occurs until you issue the update command to operationally commit the changes to the switch.

Note: If the application distribution matrix table is empty, all incoming packets will be discarded at the ingress port which has an ACL for redirection to an empty application distribution matrix. However, ACL statistics still show all the packet counts for all 8 trunk groups, because these packets match the ACLs at the ingress port.

Server Load Balancing within an Application Group

Server Load Balancing (SLB) allows you to configure the switch to balance network traffic among a pool of available servers in an application group. iFlow Director uses hash-based load balancing.

SLB provides the following benefits:

- Increased efficiency:

  With SLB, your switch is aware of the shared services provided by your application group and can then balance user traffic among the available servers.

- Increased reliability:

  If any server in an application group fails, the remaining servers continue to provide access to vital applications and data.
- Increased scalability:

  As users are added and the server pool's capabilities are saturated, new servers can be added to the application group transparently.

Identical content must be available to each server in the application group. For example:

- Static applications and data can be duplicated on each real server in the application group.
- Each real server has access to the same data through use of a shared file system or back-end database server.

Selecting the Load-Balancing Metric

Traffic in a trunk group is statistically distributed among member ports using a *hash* process where various address and attribute bits from each transmitted frame are recombined to specify the particular trunk port the frame will use. iFlow Director uses trunk hashing to provide a variety of hashing options.

To achieve the most even traffic distribution, select options that exhibit a wide range of values for your particular network. Avoid hashing on information that is not usually present in the expected traffic or which does not vary.

Use the trunk hash commands (`/cfg/l2/thash`) to configure the hash-based load balancing behavior. You can select a minimum of one or a maximum of two parameters to create one of the following configurations:

- Source IP (SIP)
- Destination IP (DIP)
- Source MAC (SMAC)
- Destination MAC (DMAC)
- SIP + DIP
- SMAC + DMAC

Configuring VLAN-Based Forwarding

The switch can insert a pseudo VLAN tag onto packets entering the switch from an uplink port. The pseudo VLAN tag includes a VLAN ID that matches the port VLAN ID (PVID). This additional VLAN tag is inserted even if the original packet already has a VLAN tag. The pseudo VLAN tag identifies the receiving port.

Software running in an SLB application group can use the PVID to apply its security policies.

When traffic is sent out from an SLB application group, the switch can use an ACL that checks the pseudo VLAN tag and directs traffic to the appropriate egress port. As the packet leaves the switch, the pseudo VLAN tag with a VLAN ID that matches the egress port PVID is removed.

Use the following command to add the pseudo VLAN tag (VLAN ID = 101) on ingress packets.

```
/cfg/port x/tagpvid e
```

Use the following command to remove the pseudo VLAN tag (VLAN ID = 101) on egress packets.

```
/cfg/port x/tagpvid d
```

Figure 3 illustrates the concept of VLAN-based forwarding.

Figure 3. VLAN-based Forwarding

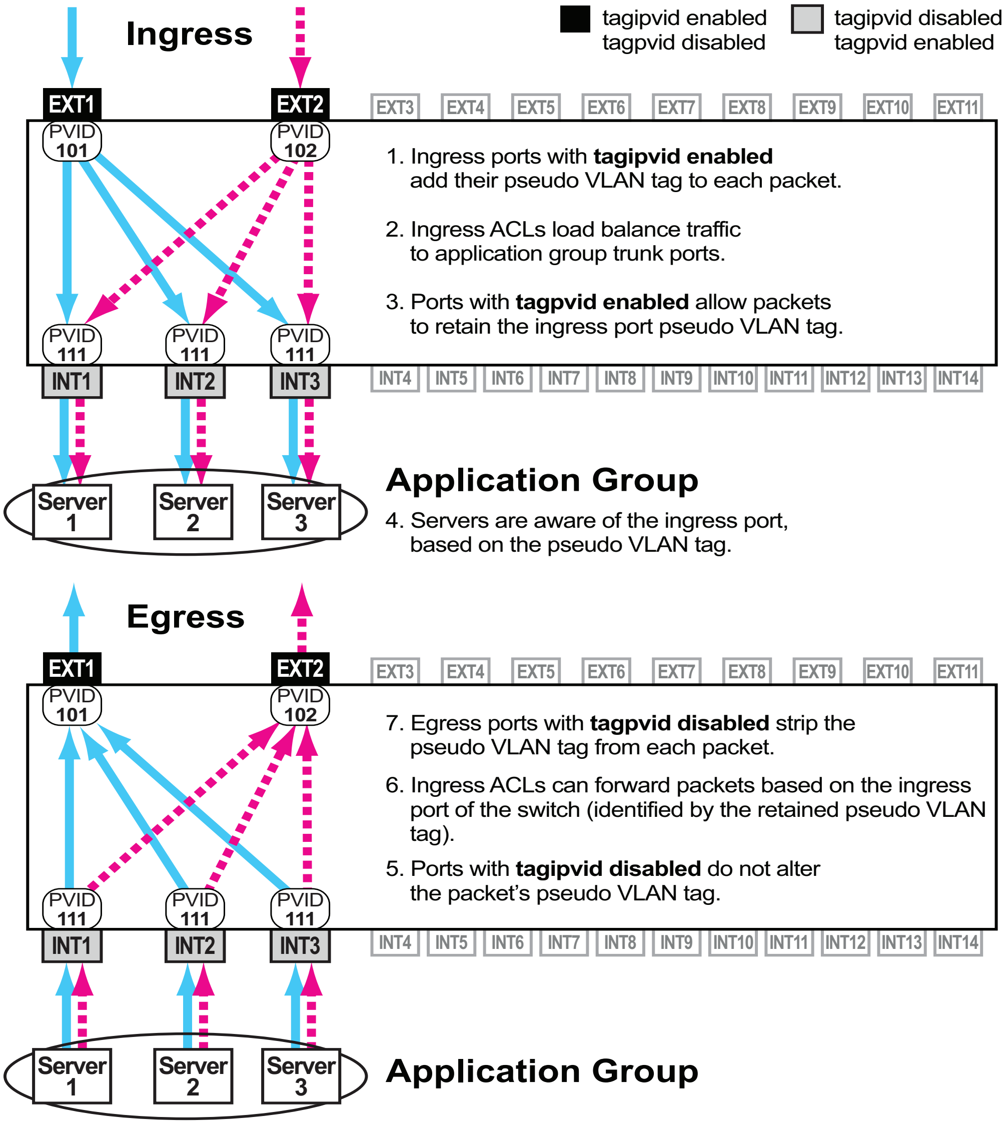

1. Ingress ports with **tagipvid enabled** add their pseudo VLAN tag to each packet.

2. Ingress ACLs load balance traffic to application group trunk ports.

3. Ports with **tagpvid enabled** allow packets to retain the ingress port pseudo VLAN tag.

4. Servers are aware of the ingress port, based on the pseudo VLAN tag.

7. Egress ports with **tagpvid disabled** strip the pseudo VLAN tag from each packet.

6. Ingress ACLs can forward packets based on the ingress port of the switch (identified by the retained pseudo VLAN tag).

5. Ports with **tagipvid disabled** do not alter the packet’s pseudo VLAN tag.

Note: Figure 3 depicts bidirectional processing of packets that ingress and egress the switch.

The following example shows the general steps used to configure VLAN-based forwarding as shown in Figure 3.

1. Configure switch uplink ports.

   a. General parameters.

```
/cfg/port EXT1
     pvid 101
     tag ena
     learn dis
     tagiskip ena
     tageskip ena
     arpmp dis
     floodblk ena
/cfg/port EXT2
     pvid 102
     tag ena
     learn dis
     tagiskip ena
     tageskip ena
     arpmp dis
     floodblk ena
```

   b. Specify that the port should add its PVID as a pseudo VLAN tag to packets at ingress.

```
/cfg/port EXT1/tagipvid ena
/cfg/port EXT2/tagipvid ena
```

   c. Specify that the port should strip the pseudo VLAN tag from packets at egress.

```
/cfg/port EXT1/tagpvid dis
/cfg/port EXT2/tagpvid dis
```

Note: The `tagpvid` option is disabled by default and need not be explicitly configured unless the settings have been previously enabled.

2. Configure SLB application ports.

   a. General parameters

```
/cfg/port INT1
     pvid 111
     tag ena
     learn dis
     tagiskip ena
     tageskip ena
     arpmp dis
     floodblk ena
/cfg/port INT2
     pvid 111
     tag ena
     learn dis
     tagiskip ena
     tageskip ena
     arpmp dis
     floodblk ena
/cfg/port INT3
     pvid 111
     tag ena
     learn dis
     tagiskip ena
     tageskip ena
     arpmp dis
     floodblk ena
```

   b. Specify that the port will not change or add a pseudo VLAN tag at ingress.

```
/cfg/port INT1/tagipvid dis
/cfg/port INT2/tagipvid dis
/cfg/port INT3/tagipvid dis
```

Note: The `tagipvid` option is disabled by default and need not be explicitly configured unless the settings have been previously enabled.

   c. Specify that the port will retain the pseudo VLAN tag at egress.

```
/cfg/port INT1/tagpvid ena
/cfg/port INT2/tagpvid ena
/cfg/port INT3/tagpvid ena
```

3. Define an SLB application group.

   a. Enable SLB and define application group servers.

```
/cfg/slb/on
     method trunk
/cfg/slb/real 1
     ena
     port INT1
/cfg/slb/real 2
     ena
     port INT2
/cfg/slb/real 3
     ena
     port INT3
```

b. Add the servers to an SLB group.

```
/cfg/slb/group 1
     add 1
     add 2
     add 3
```

c. Create an application group for the SLB group and enable the error handling remap option.

```
/c/slb/app 1
     group 1
     errhand/remap ena
```

d. Configure a trunk group to service the SLB application.

```
/cfg/l2/trunk 1
     ena
     slbapp 1
```

e. Configure an ACL redirection filter for the SLB application.

```
/cfg/acl/acl 1/target
     dest slbapp
     slbapp 1
/cfg/acl/acl 1
     action redirect
     stats ena
```

f. Create an ACL group for the ACL redirection filter.

```
/cfg/acl/group 1/add 1
```

g. Apply the ACL group to the uplink ports, to steer traffic to the application group.

```
/cfg/port EXT1/acl/add grp 1
/cfg/port EXT2/acl/add grp 1
```

4. Create port mapping ACLs for the SLB application ports. Use PVID as a port identifier to match against the packet VLAN ID, and redirect traffic to the appropriate egress port.

a. Configure ACLs to match against PVID 101, and redirect to port EXT1.

```
/cfg/acl/acl 101/target
     dest port
     port EXT1
/cfg/acl/acl 101/ethernet/vlan 101 0xfff
/cfg/acl/acl 101/action redirect
```

b. Configure ACLs to match on PVID 102, to redirect traffic to port EXT2.

```
/cfg/acl/acl 102/target
     dest port
     port EXT2
/cfg/acl/acl 102/ethernet/vlan 102 0xfff
/cfg/acl/acl 102/action redirect
```

c. Configure an ACL to drop unexpected traffic.

```
/cfg/acl/acl 200/ethernet/etype any
/cfg/acl/acl 200
     action deny
     stats ena
```

d. Add port mapping ACLs to a group to be applied to all SLB application ports.

```
/cfg/acl/group 100
     add 101
     add 102
     add 200
```

e. Apply the port mapping ACL group to the SLB application ports.

```
/cfg/port INT1/acl/add grp 100
/cfg/port INT2/acl/add grp 100
/cfg/port INT3/acl/add grp 100
```

5. Apply and save the configuration.

```
apply
save
```

Server Health-Checks Options

Health checking allows you to verify server accessibility in an application group. As content grows and information is distributed across servers in an application group, content health checks ensure end-to-end availability.

Use the following command to configure the health check option for an application group:

```
>> # /cfg/slb/group <SLB group number>/health
```

The following server health checks are available in iFlow Director:

- Link State (default)
- IPS (Layer 2 application-specific health check)

Note: Link and IPS health check methods are available when the trunk hashing method is selected (`/cfg/slb/method` is set to `trunk`). These options are not available when the ECMP hashing method is selected (`/cfg/slb/method` is set to `ecmp`).

- ICMP ping
- TCP (Transmission Control Protocol)
- TCP script-based health checks
- Application-specific health-checks with configurable TCP Port Numbers (for example, HTTP could use port 8080 instead of 80)
  - HTTP (HyperText Transfer Protocol)
  - HTTP Head message
  - SSL (Secure Socket Layer)
  - SMTP (Simple Mail Transfer Protocol)
  - SIP-register (Session Initiated Protocol)
  - DNS (Domain Name Server)
  - ARP (Address Resolution Protocol)

Note: If the MOD option is enabled (`/cfg/slb/mod enable`), all health checking is disabled.

Each real server's health can be checked by type, service port, and content configured for the application group to which it belongs. The health-check setting applies to every real server in the group, though health check packets are only sent out through the switch ports configured for the real servers (`/cfg/slb/real` *x*`/port`), not all ports.

For all Layer 4 health-checks, and for ping and ARP, the service port also must be configured:

```
>> # /cfg/slb/group <SLB group number>/svport <service port number>
```

Health checks are performed on intervals of 1-60 seconds, specified under the real server configuration, as follows:

```
>> # /cfg/slb/real <server number>/inter <1-60, or 400 ms>
```

The 400 millisecond interval is available only for link, IPS, or ping health checks.

Health Check Configuration Example

The following example shows the general steps used to configure health checks for an existing SLB group (configured as shown in prior examples in this chapter).

1. Configure the health check type for the SLB group.

```
/cfg/slb/group <1-14>/health [link|ping|ips|tcp|http|httphead|smtp|ssl|
udpdns|arp|sip-register|script <script number>]
```

2. Configure the service port used to send data to the server. The `svport` option is ignored for health checks that do not require it.

```
/cfg/slb/group <1-14>/svport <0-65534>
```

3. Define the content string (optional, for types HTTP, HTTPhead, SMTP, UDPDNS, SIP-register).

```
/cfg/slb/group <1-14>/content <text string>
```

The `content` option is used to configure type-specific data, such as a URL (`www.ibm.com`) for DNS health checks, a filename and file path (`index.html`) for HTTP health checks, a username for SIP register health checks, an email address (`joe@ibm.com`) for SMTP health checks, and so on. The maximum string length for the content field is:

– SMTP = 1-64 characters

– SIP =1-32 characters

– HTTP/HTTP Head/UDPDNS = 1-127 characters

4. If using a Layer 3 or Layer 4 health-checking type (ping, TCP, HTTP, etc.), define the server's IPv4 address.

```
/cfg/slb/real <1-84>/rip <IP address>
```

Note: An IP address is not required for Link health checking (the default), or IPS health checking.

Scriptable Health Checks

Scriptable health checks allow you to create a customized TCP health check for service types not specifically implemented. The script can mimic a standard service request and response sequence and check if the behaviors match what is expected.

Use the following command to configure a health check script:

```
/cfg/slb/script <1-16>
```

Consider the following guidelines when you configure a health-check script:

- The script must begin with the `open` command to a specific service port.
- The `send` command is used to define a text string to send to the specified service port.

- The expect command is used to define a text string to verify that the first part of the text string received matches the string configured in the script. The script fails if the exact string is not included in the server response.
- It is recommended that each script contain only one open command.
- The switch considers the health check to be successful only if the entire script has been executed and properly matched.
- The following commands are ignored when the switch performs a scriptable health check:

  /cfg/slb/group <*1-14*>/svport

  /cfg/slb/group <*1-14*>/content

The following example script checks the HTTP service:

```
/cfg/slb/script 1
open "80"
send "GET / HTTP/1.0\\r\\n\\r\\n"
expect "HTTP/1.1 200"
close
```

Note: The backslash (\ ) is used as an escape character in the script. If you want to include a backslash in the text string, you must preface it with a second backslash.

Layer 2 Mode Guidelines

When configuring iFlow Director, consider the following guidelines:

- During a configuration apply to SLB or ACL-dependent parameters, you might experience temporary traffic disruptions, because the apply operation causes the switch to re-initialize and reinstall these modules. You might experience the following types of disruptions:
  - Temporary packet discards
  - Temporary traffic flooding via Layer 2
  - Temporary hash disruptions
  - Inaccurate statistical accounting
- If you use SNMP to set the next boot configuration file to factory default and perform a save operation, the switch automatically changes the next-boot setting to the active configuration block instead of the factory configuration block. It is recommended that you immediately reboot the switch after setting the configuration block to factory default.
- To avoid errors when using SNMPwalk, use the `-Cc` option with the SNMPwalk command. For example:

  ```
  snmpwalk -Cc -v 1 -t 60 -c public -m ALL -M $miblocation
  10.13.5.103 $1
  ```

Statistics

The following guidelines apply to statistics with iFlow Director:

- Statistics for ECMP routes over the management interface are not supported.
- The following command displays statistics for each real server port in each SLB application when the SLB applications are configured for aggregation:

  ```
    /stats/slb/app all
  ```

  However, if a backup real server shared among the SLB applications becomes active, the statistics on this real server port are counted from all the SLB applications that it is serving. For example, if the shared backup real server port receives 100 packets from SLB application 1 and 200 packets from SLB application 2, the command above displays a packet count of 300 for the real server in each SLB application.

Access Control Lists (ACLs)

The following guidelines apply to ACLs with iFlow Director:

- ACL redirection action takes precedence over port mirroring actions. If you configure port mirroring, packets are not mirrored if a redirection ACL is configured on the port.
- Statistics for ACL redirection (`/stats/acl/dump`) show ACL matches, but not necessarily actions. In some cases, the statistics increment because an ACL matched the packet header, but no redirection took place. For example, a packet might be dropped on ingress, so no redirection occurred.
- When a 1Gb port is configured with metering but does not have flow control enabled, the switch can drop Out-of-Profile packets normally. However, when flow control is enabled on the port, it will take precedence on the port to limit the rate of traffic at 1Gbps. In this case, metering does not take effect if the committed rate is greater or equal to 1,000,000 kbps (1Gbps).

- When a redirect ACL is configured to run on cluster mode, traffic intended for Layer 3 routing (traffic with a DMAC that matches the switch MAC) will be redirected to an SLB application group.

  In the following example ECMP configuration, there must be a forward traffic ACL and a reverse traffic ACL:

```
/cfg/acl/acl 1/target/dest slbapp
/cfg/acl/acl 1/target/slbapp 1
/cfg/acl/acl 1/ethernet/vlan 100
/cfg/acl/acl 1/ipv4/sip 100.100.100.0 255.255.255.0
/cfg/acl/acl 1/tcpudp/dport 80
/cfg/acl/acl 1/action redirect
/cfg/acl/acl 1/stats ena
```

```
/cfg/acl/acl 10/target/dest slbapp
/cfg/acl/acl 10/target/slbapp 1
/cfg/acl/acl 10/ethernet/vlan 101
/cfg/acl/acl 10/ipv4/sip/101.101.101.0 255.255.255.0
/cfg/acl/acl 10/tcpudp/dport 80
/cfg/acl/acl 10/action redirect
/cfg/acl/acl 10/stats ena
```

Health Checks

The following guidelines apply to health checks with iFlow Director:

- When writing a health-check script, you must use numeric values for service ports, because aliases are not accepted (for example, use `80`, not `http`).
- When configuring health check content within an SLB group, the double quotation mark (") is used as a special escape character. Therefore, it cannot be included in the body of the content.
- Health checks are performed asynchronously on peer switches when using the trunk hashing method. This can cause hash buckets to become out-of-sync between the switches. To improve hash synchronization, do one of the following:
  - Use the IPS health check option and enable the SLB remap option (`/c/slb/app` *x*`/errhand/remap`).
  - Enable the Manual OSP Distribution (`/c/slb/mod ena`) and manually perform the hash matrix updates across the peer switches with an external application.
- Health checks can cause an application to go down and come up repeatedly (flap) if the health check interval or retry values are too small. To avoid this issue, increase the value of the real server health check interval (`/cfg/slb/real` *x*`/inter`) or the retry option (`/cfg/slb/real` *x*`/retry`).
- ISCLI: If the switch is configured with health checks that depend on a specific service-port value (`smtp`, `sip-register`, or `udpdns`), you must first change the health-check type to `link`. Then you can set the service port to the appropriate value and configure the appropriate health-check type.

Chapter 3. Layer 3 Mode Hashing

The following IBM iFlow Director Layer 3 Mode Hashing topics are discussed in this chapter:

Layer 3 Mode Features

Server Load Balancing

With Server Load Balancing (SLB), the switch can balance network traffic among a pool of available application blades. iFlow Director uses a configurable hash algorithm to perform wire-speed SLB for an application group with up to 84 servers.

SLB provides the following benefits:

- **Increased efficiency:** The switch is aware of the shared services provided by your application group and can then balance user traffic among the available servers.
- **Increased reliability:** If any server in an application group fails, the remaining servers continue to provide access to vital applications and data.
- **Increased scalability:** As users are added and the server pool's capabilities are saturated, new servers can be added to the application group transparently.

Identical content must be available to each server in the application group. For example:

- Static applications and data can be duplicated on each real server in the application group.
- Each real server has access to the same data through use of a shared file system or back-end database server.

Routing

iFlow Director includes versatile routing options:

- Standard routing using IPv4 or IPv6 static routes.
- Standard routing using IPv4 OSPF or IPv6 OSPFv3.
- Policy-based routing for IPv4 that can be tailored to any qualifiers configurable on a normal ACL.

High-Availability

iFlow Director includes high availability features, such as:

- Virtual Router Redundancy Protocol (VRRP) for placing dual iFlow Director switches in an active-standby configuration.
- Server health-checks (ping, ARP, HTTP, TCP, TCP-script) for rapid failure detection in IPv4 or IPv6 networks.
- Failover options allow distributing flows from a failed blade to other active or backup blades.

These features, coupled with industry-leading BladeCenter hardware innovations (redundant mid-plane, redundant power domains, redundant management modules for platform monitoring) improve the overall system reliability for Enterprise, Data Center and Carrier Grade deployments.

How Layer 3 Mode Works

Network applications such as caches, firewalls, Intrusion Prevention Systems, Content Gateways, and other Deep Packet Inspection (DPI) devices are usually hosted on application-specific devices. These single-purpose appliances process traffic in-line at the network perimeter and within the core networks. To meet capacity requirements, an iFlow Director helps consolidate the function of stand-alone appliances, load-balancers and Ethernet switches into an IBM BladeCenter chassis, as shown in Figure 4. .

Figure 4. Application Network with iFlow Director

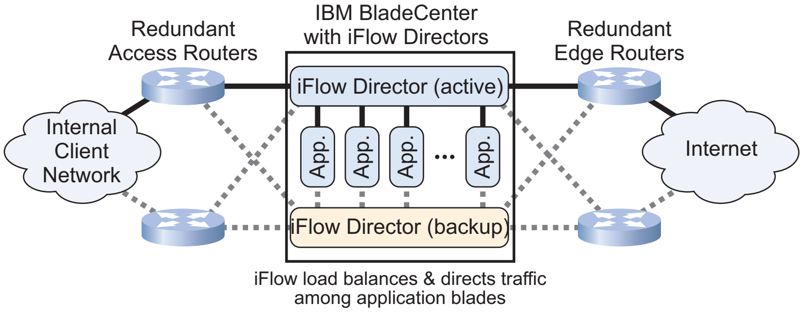

The IBM BladeCenter H or HT chassis with iFlow Director is deployed between the internal client network and the external network.

The iFlow Director steers traffic from the client network and the Internet to the application blades. When corresponding traffic from the applications blades comes back to the switch, the iFlow Director routes the traffic to the appropriate client destination or Internet destination using either standard routing (via static routes or OSPF) or policy-based routing.

For redundancy, dual iFlow Directors may be deployed (as shown) using VRRP. One switch is active and the other is a backup, taking over in case of a component or link failure.

An Example of the Layer 3 Hashing Process

A typical iFlow Director scenario works as follows:

1. Both the access router and edge router are configured with the iFlow Director as their next hop gateway.
2. Client traffic to the Internet enters from the internal network. For example, a data request from a user mobile device is collected by a backhaul network and forwarded to the internal access router. The access router then forwards the traffic to the iFlow Director as the next hop toward the Internet.
3. A Server Load Balancing (SLB) ACL filter on the iFlow Director distributes client traffic to a pre-configured SLB *application group*, an aggregation of individual blade servers running a common application, such as a web cache, firewall, or Intrusion Detection System. The traffic is load-balanced among the servers in the application group using a hash algorithm based on the client Source IP (SIP). Or, depending on the application, the Destination IP (DIP) may be used instead.
4. The selected application server processes the intercepted client request and responds accordingly, either with a response to the original client or a request to the Internet or another network destination.

5. When the traffic leaves the application group, it arrives at the internal port of the iFlow Director, which forwards the traffic to the Internet using either standard routing (static routes or OSPF) or using policy-based routing.
6. Internet traffic, such as responses to client requests, enters from the external network. The edge router forwards this traffic to the iFlow Director as the next hop to the client network.
7. Another SLB ACL filter on the iFlow Director redirects the Internet traffic to the application group. To maintain flow persistency, the hash algorithm uses the Destination IP (DIP) for return traffic, ensuring that traffic from or to the same client IP address is send to the same server in the application group. Or depending on the application, the Source IP (SIP) may be used instead.
8. The selected application server processes the intercepted Internet response and responds to the original client.
9. When the response traffic leaves the application group, it arrives at the internal port of the iFlow Director, which forwards the traffic to the client using either standard routing (static routes or OSPF), or using policy-based routing.
10. If a server fails, existing flows from the failed server are processed by other servers in the application group.
11. For additional redundancy, a second iFlow Director can be used in the IBM BladeCenter chassis and synchronized using VRRP. One switch is active and the other is a backup, taking over in case of a component or link failure.

Basic Layer 3 Mode Configuration

Consider the following network topology:

Figure 5. Simplified Server Load Balancing Example

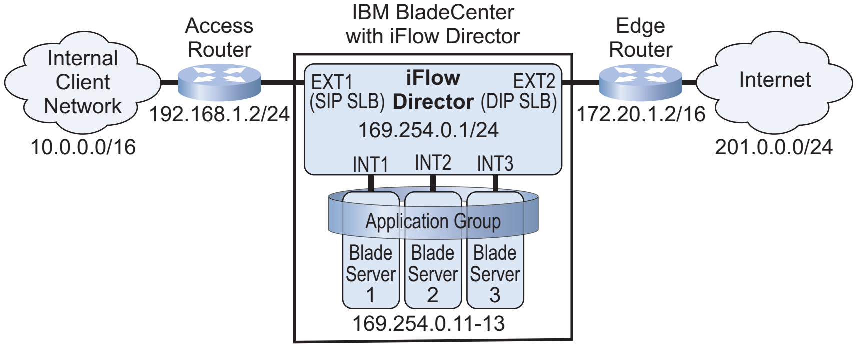

In this simplified example, a single iFlow Director is placed between the client network and the Internet. The iFlow Director is configured with an application group comprised of three blade servers running one specialized application (such as a web-cache).

The access router is connected to the iFlow Director on port EXT1. An SLB ACL on this port performs load balancing based on a hash of the client's source IP (SIP) address. For traffic coming in the other direction, as from the edge router connected on port EXT2, load balancing is based on a hash of the destination IP (DIP). Because the SIP of a client request is the same as the DIP of the corresponding Internet response, the same blade server is selected to process both directions of traffic for any particular client, maintaining session persistence.

In this scenario, the iFlow Director must be configured with the following:

- Network interfaces and port link behavior.
- SLB settings for the blade servers participating in the application group.
- SLB ACLs for directing client and Internet traffic to the application group.
- Routing, so that resulting traffic from the application group can be forwarded to the appropriate destination.

Note: The access router and the edge router must be configured with the iFlow Director as their next hop gateway.

Configure Network Interfaces and Port Link Behavior

1. Define the switch IP interface:

```
/cfg/l3/if 1
     addr 169.254.0.1
     maskplen 255.255.255.0
     ena
```

Note: In this example, IPv4 interfaces are used. However, IPv6 interfaces are also supported.

2. Define an IP interface for the client network and Internet:

```
/cfg/l3/if 2
     addr 192.168.1.1
     maskplen 255.255.255.0
     ena
/cfg/l3/if 3
     addr 172.20.1.1
     maskplen 255.255.0.0
     ena
```

3. Define the port link behavior for the external network ports:

```
/cfg/port EXT1/floodblk ena
/cfg/port EXT2/floodblk ena
```

4. Define the port link behavior for the application server blades:

```
/cfg/port INT1
     gig
     fctl none
     speed 10000
     auto off
/cfg/port INT2
     gig
     fctl none
     speed 10000
     auto off
/cfg/port INT3
     gig
     fctl none
     speed 10000
     auto off
```

Configure Server Load Balancing

The following procedure demonstrates how to configure an SLB application group comprising server blades 1 through 3.

1. Turn on the SLB feature and select the ECMP hashing method.

```
/cfg/slb/on
     method ecmp
```

Note: The `ecmp` setting is the default and need not be explicitly configured unless switch settings have been previously altered. Also, the `ecmp` setting (Layer 3 hashing) and the `trunk` setting (Layer 2 hashing) are mutually exclusive. When `ecmp` is selected, Layer 2 hashing (page 5) is not supported.

2. Define one real server for each internal server port which participates in the application group.

```
/cfg/slb/real 1
     ena
     port INT1
     rip 169.254.0.11
/cfg/slb/real 2
     ena
     port INT2
     rip 169.254.0.12
/cfg/slb/real 3
     ena
     port INT3
     rip 169.254.0.13
```

Note: The server ports (port INT1, INT2, and INT3) are optional; only the RIP is required for each real server.

3. Define an SLB group comprised of the configured real server IDs.

```
/cfg/slb/group 1
     add 1
     add 2
     add 3
```

4. Define health checking for the SLB group.

```
/cfg/slb/group 1
     health http
```

Health checks can be performed using a variety of mechanisms such as ping, ARP, and TCP characteristics and scripts. For expanded options, see “Server Health-Check Options” on page 60.

5. Add the SLB group to the application group.

```
/cfg/slb/app 1
     group 1
```

Configure SLB ACLs

1. Define an SLB ACL filter for the client network.

```
/cfg/acl/acl 10/target
     dest slbapp
     slbapp 1
/cfg/acl/acl 10/tcpudp/dport 100 0xffff
/cfg/acl/acl 10
     action redirect
     stats ena
```

Note: In this example, IPv4 ACLs are configured. IPv6 ACLs are also supported, but only when the `action` is `permit`, `deny`, or `setprio`.

2. Set the ECMP hash metric to use the SIP for the EXT1 port receiving the client network traffic.

```
/cfg/port EXT1/ecmphash sip
```

3. Define an ACL group to service the SLB filter.

```
/cfg/acl/group 1
     add 10
```

4. Apply the ACL group to the client port to steer traffic to the application group.

```
/cfg/port EXT1/acl
     add grp 1
```

Note: The switch prevents you from directly adding an SLB ACL filter directly to any switch port. This ACL must be added indirectly via the ACL group.

5. Define an SLB ACL filter for the Internet side of the network.

```
/cfg/acl/acl 20/target
     dest slbapp
     slbapp 1
/cfg/acl/acl 20/tcpudp/dport 100 0xffff
/cfg/acl/acl 20
     action redirect
     stats ena
```

6. Set the ECMP hash metric to use the DIP for the EXT2 port receiving the Internet traffic.

```
/cfg/port EXT2/ecmphash dip
```

Note: This ACL maintains session persistence by ensuring that return traffic (from the Internet to the client) is directed to the same application server that originally handled the client's request.

7. Define an ACL group to service the SLB filter.

```
/cfg/acl/group 2
     add 20
```

8. Apply the ACL group to the Internet port to steer traffic to the application group.

```
/cfg/port EXT2/acl
     add grp 2
```

Advanced Layer 3 Mode Configuration

The IBM iFlow Director can selectively load-balance traffic on Layers 2, 3, and 4 using port-based ECMP hashing, ingress VLAN retention, and ACL qualifiers to steer packets to the blade servers. Using PBR within the ACL, all ACL qualifiers such as protocol or SIP and DIP address range can be used as criteria. If the specified qualifiers are not met, traffic is either dropped or handled as specified.

Consider the following network topology:

Figure 6. Multi-Chassis Server Load Balancing Example

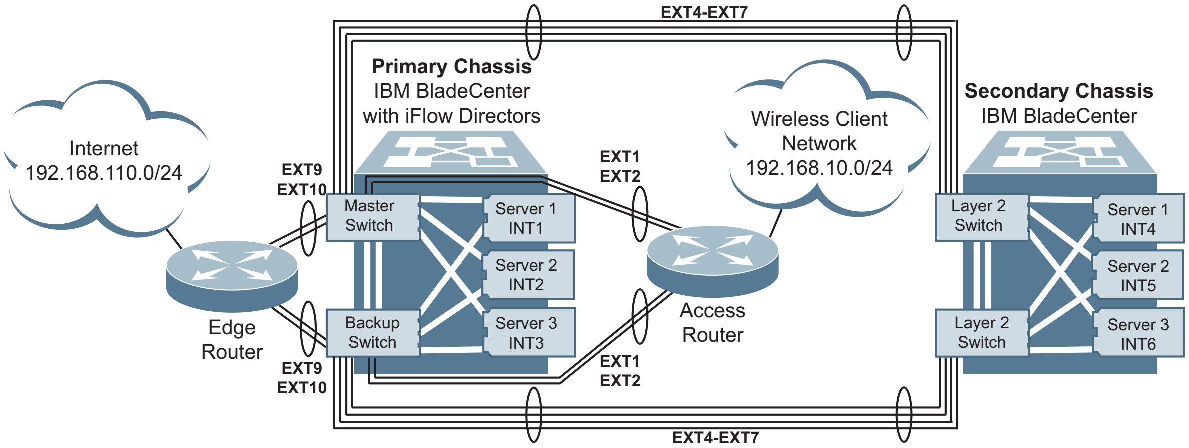

In this example, two iFlow Directors are used for load balancing. A primary chassis with an iFlow Director running as a VRRP master switch and a backup switch is placed between the client network and the Internet. A secondary chassis with two Layer 2 switches is connected to it through a series of trunks. The iFlow Directors in the primary chassis are each configured with an application group comprised of six blade servers (three blade servers in the primary chassis, and three blade servers in the secondary chassis) running one specialized application (such as a web cache).

The access router is connected to the iFlow Director VRRP master and backup switch in the primary chassis, separately, via a trunk consisting of ports EXT1 and EXT2, and a secondary chassis is connected to the primary chassis through a trunk consisting of ports EXT4-EXT7. EXT1 and EXT2 ports are using the default hash setting based on the client's source IP (SIP) address, with ACLs to redirect the ingress traffic to the primary and secondary chassis blade servers group. For traffic coming in the other direction, as from the edge router connected on ports EXT9 and EXT10, load balancing is based on a hash of the destination IP (DIP). Because the SIP of a client request is the same as the DIP of the corresponding Internet response, the same blade server is selected to process both directions of traffic for any particular client, maintaining session persistence.

There are four switches involved in this configuration example:

- two iFlow Director switches in the primary chassis:
  - the VRRP Master
  - the VRRP Backup
- two ordinary Layer 2 switches in the secondary chassis

The configurations of the two iFlow Director switches are basically the same. Whenever there is a difference, the respective commands are marked with Switch 1 (for VRRP Master switch) or Switch 2 (for VRRP Backup switch). In the secondary chassis, the first Layer 2 switch connecting to the Master is Switch 3, and the secondary Layer 2 switch connecting to the Backup is Switch 4.

The following are key features for configuration of these four switches:

- The requirement for this configuration is that the blade servers are accepting and returning untagged traffic with various subnets by setting `retivlan` to `disabled` and using `nhvlan`.
- Traffic coming from the Wireless Client and Internet external ports will not retain the ingress VLAN tag when forwarded to the internal ports. When the traffic send to the blade servers, they will be untagged on VLAN 3000, which is the health check VLAN. After the traffic were processed by the blade servers, they are returned back to the iFlow Director switch, which will be sent out tagged with egress interface VLAN according to the nexthop ACLs or dynamic/static route specified.
- If the blade servers returned traffic match the selective traffic ACL defined subnet 30.2.0.0, then send the traffic to the next hop router with IP address 169.244.144.100.
- If the blade servers returned traffic match the selective traffic ACL defined subnet 130.2.0.0, then send the traffic to the next hop router with IP address 169.234.144.100.
- Trunks are used on the Wireless Client and Internet side, with ACL group filters applied to those ports.
- This configuration example is not for VLAN PBR deployment.

iFlow Director VRRP Switch Configuration (Switch 1 and Switch 2)

The following steps show the iFlow Director VRRP switch configuration for Switch 1 and Switch 2.

Note: The following example configuration is one of the more commonly used, and may differ from yours, depending on your specific needs.

Configure Port Link Behavior

1. Define the internal port link behavior for the application server blades in the primary chassis to accept VLAN tagged packets:.

```
/cfg/port INT1/gig
        speed 10000
        auto off
        fctl none
        tag enable
/cfg/port INT2/gig
        speed 10000
        auto off
        fctl none
        tag enable
/cfg/port INT3/gig
        speed 10000
        auto off
        fctl none
        tag enable
```

2. Define the Wireless Client external ports link to accept VLAN tagged packets:

```
/cfg/port EXT1       /* Default is retivlan=disabled, that is to not retain
                     /* the ingress VLAN tag
        tag enable   /* Default is ecmphash=sip, hence no need to configure
        floodblk enable
/cfg/fgport EXT2
        tag ena
        floodblk enable
```

3. Define the Internet external ports link to accept VLAN tagged packets:

```
/cfg/port EXT9
        tag enable
        ecmphash dip /* Ingress traffic from Internet to ecmphash using dip
        floodblk ena
/cfg/port EXT10
        tag enable
        ecmphash dip /* Ingress traffic from Internet to ecmphash using dip
        floodblk enable
```

4. Define the inter-chassis external ports link behavior to accept VLAN tagged packets.

```
/cfg/port EXT4
        tag enable
/cfg/port EXT5
        tag enable
/cfg/port EXT6
        tag enable
/cfg/port EXT7
        tag enable
```

Configure VLANs onto the Internal and External Ports

1. Create VLAN 10 for the wireless client traffic coming into the wireless side external trunk ports (EXT1 and EXT2), the inter-chassis trunk ports connecting to the secondary chassis (EXT4 to EXT7), and the blade server ports (INT1 to INT3):

```
/cfg/l2/vlan 10
        enable
        add INT1-INT3
        add EXT1-EXT2
        add EXT4-EXT7
```

2. Create an additional VLAN 1000 for EXT1 and EXT2 port only:

```
/cfg/l2/vlan 1000
        enable
        add EXT1-EXT2
```

3. Create VLAN 110 for the Internet traffic coming into the Internet side external trunk ports (EXT9 and EXT10), the inter-chassis trunk ports connecting to the secondary chassis (EXT4 to EXT7), and the blade server ports (INT1 to INT3):

```
/cfg/l2/vlan 110
        enable
        add INT1-INT3
        add EXT9-EXT10
        add EXT4-EXT7
```

4. Create an additional VLAN 2000 for EXT9 and EXT10 port only:

```
/cfg/l2/vlan 2000
        enable
        add EXT9-EXT10
```

5. Assign VLAN 3000 for the internal blade servers health check and communication. Set this VLAN as the untagged VLAN for those related ports:

```
/cfg/l2/vlan 3000
        enable
        add INT1-INT3
        add EXT4-EXT7
/cfg/port INT1
        pvid 3000
/cfg/port INT2
        pvid 3000
/cfg/port INT3
        pvid 3000
/cfg/port EXT4
        pvid 3000
/cfg/port EXT5
        pvid 3000
/cfg/port EXT6
        pvid 3000
/cfg/port EXT7
        pvid 3000
```

6. Assign VLAN 4000 to EXT11 port for VRRP advertisement communication between the Master switch and the Backup switch. Set this VLAN as the untagged VLAN for the port:

```
/cfg/port EXT11
        pvid 4000
/cfg/l2/vlan 4000
        enable
        add EXT11
/cfg/l2/vlan 1
        rem EXT11
```

Turn Off Spanning Trees

Since this pair of IBM iFlow Director switches are within a controlled environment, turn off spanning trees to avoid unnecessary BPDUs:

```
/cfg/l2/nostp enable
/cfg/l2/stg 1/off
```

Create Trunks

Create Wireless Client, Internet, and inter-chassis trunks:

```
/cfg/l2/trunk 1  /* Wireless Client trunk
        enable
        add EXT1
        add EXT2
/cfg/l2/trunk 2  /* Internet trunk
        enable
        add EXT9
        add EXT10
/cfg/l2/trunk 3  /* Inter-chassis trunk
        enable   /* Although default l3thash is SIP+DIP, it is implemented
                 /* as SIP only
        add EXT4
        add EXT5
        add EXT6
        add EXT7
```

Turn on SLB and Create Real Servers for the Blade Servers

1. In a multi-chassis environment, it is not necessary to configure a port for each real server.

```
/cfg/slb
        On
/cfg/slb/real 1
        enable
        rip 169.254.144.1
/cfg/slb/real 2
        enable
        rip 169.254.144.2
/cfg/slb/real 3
        enable
        rip 169.254.144.3
/cfg/slb/real 4    /* From here on are the servers in the secondary chassis
        enable
        rip 169.254.144.4
/cfg/slb/real 5
        enable
        rip 169.254.144.5
/cfg/slb/real 6
        enable
        rip 169.254.144.6
```

2. Add all real servers to a group and define the type of health check used:

```
/cfg/slb/group 1
        health http
        content "index.html"
        svport 80
        add 1
        add 2
        add 3
        add 4
        add 5
        add 6
```

3. Assign the real server group into the SLB application group that belongs to next hop VLAN 3000:

```
/cfg/slb/app 1
        group 1
        nhvlan 3000
```

Note: You must assign a value to `nhvlan` when `retivlan` (the default setting is to not retain the ingress VLAN tag) is set to `disabled` on the external wireless client and Internet external ports.

Configure Selective Traffic SLB and PBR ACLs

1. Define selective traffic to be load-balanced to SLB application group 1. Also define the selective traffic next hop routing ACLs with the respective next hop router. If there is no match on the selective traffic routing ACLs, use the dynamic or static route (as shown in "Configure Static Routes" on page 51).

```
/cfg/acl/acl 39/target/dest slbapp
/cfg/acl/acl 39/target/slbapp 1
/*** Redirect wireless packets with this destination MAC to App Grp 1
/cfg/acl/acl 39/ethernet/dmac 00:00:5e:00:01:02 ff:ff:ff:ff:ff:ff
/cfg/acl/acl 39/action redirect
/cfg/acl/acl 39/stats enable
/*** Selective traffic ACL, redirect specific destination subnet 30.2.0.0
/*** to next hop
/cfg/acl/acl 40/target/dest nexthop
/cfg/acl/acl 40/target/nexthop 169.244.144.100
/cfg/acl/acl 40/ipv4/dip 30.2.0.0 255.255.0.0
/cfg/acl/acl 40/action redirect
/cfg/acl/acl 40/stats enable
/cfg/acl/acl 109/target/dest slbapp
/cfg/acl/acl 109/target/slbapp 1
/*** Redirect Internet packets with this destination MAC to App Grp 1
/cfg/acl/acl 109/ethernet/dmac 00:00:5e:00:01:03 ff:ff:ff:ff:ff:ff
/cfg/acl/acl 109/action redirect
/cfg/acl/acl 109/stats enable
/*** Selective traffic ACL. redirect specific destination subnet 130.2.0.0
/*** to next hop
/cfg/acl/acl 110/target/dest nexthop
/cfg/acl/acl 110/target/nexthop 169.234.144.100
/cfg/acl/acl 110/ipv4/dip 130.2.0.0 255.255.0.0
/cfg/acl/acl 110/action redirect
/cfg/acl/acl 110/stats enable
```

2. Put the ACLs into ACL groups:

```
/cfg/acl/group 39
     add 39
/cfg/acl/group 40 /* PBR ACL group for Wireless
     add 40
/cfg/acl/group 109
     add 109
/cfg/acl/group 110 /* PBR ACP group for Internet
     add 110
```

3. Apply the ACL group to respective internal and external ports for selective traffic and PBR:

```
/cfg/port INT1/acl/add grp 40    /* Apply PBR ACL for wireless 30.2.0.0 traffic
/cfg/port INT1/acl/add grp 110   /* Apply PBR ACL for Internet 130.2.0.0 traffic
/cfg/port INT2/acl/add grp 40    /* Apply PBR ACL for wireless 30.2.0.0 traffic
/cfg/port INT2/acl/add grp 110   /* Apply PBR ACL for Internet 130.2.0.0 traffic
/cfg/port INT3/acl/add grp 40    /* Apply PBR ACL for wireless 30.2.0.0 traffic
/cfg/port INT3/acl/add grp 110   /* Apply PBR ACL for Internet 130.2.0.0 traffic
/cfg/port EXT1/acl/add grp 39
/cfg/port EXT2/acl/add grp 39
/cfg/port EXT4/acl/add grp 40    /* Apply PBR ACL for wireless 30.2.0.0 traffic
                                 /* for second chassis
/cfg/port EXT4/acl/add grp 110   /* Apply PBR ACL for Internet 130.2.0.0 traffic
                                 /* for second chassis
/cfg/port EXT5/acl/add grp 40    /* Apply PBR ACL for wireless 30.2.0.0 traffic
                                 /* for second chassis
/cfg/port EXT5/acl/add grp 110   /* Apply PBR ACL for Internet 130.2.0.0 traffic
                                 /* for second chassis
/cfg/port EXT6/acl/add grp 40    /* Apply PBR ACL for wireless 30.2.0.0 traffic
                                 /* for second chassis
/cfg/port EXT6/acl/add grp 110   /* Apply PBR ACL for Internet 130.2.0.0 traffic
                                 /* for second chassis
/cfg/port EXT7/acl/add grp 40    /* Apply PBR ACL for wireless 30.2.0.0 traffic
                                 /* for second chassis
/cfg/port EXT7/acl/add grp 110   /* Apply PBR ACL for Internet 130.2.0.0 traffic
                                 /* for second chassis
/cfg/port EXT9/acl/add grp 109
/cfg/port EXT10/acl/add grp 109
```

Create an IP Interface for Each VLAN

On each VLAN for internal blade server health check, wireless client and Internet, create an IP interface:

```
/cfg/l3/if 1                                /* Blade server health check interface
        enable
        addr 169.254.144.253                /* Switch 2 uses: 169.254.144.252
        mask 255.255.255.0
        broad 169.254.144.255
        vlan 3000
/cfg/l3/if 2                                /* Wireless Client second interface
        enable
        addr 169.244.144.253                /* Switch 2 uses: 169.244.144.252
        vlan 1000
/cfg/l3/if 3                                /* Internet second interface
        enable
        addr 169.234.144.253                /* Switch 2 uses: 169.234.144.252
        vlan 2000
/cfg/l3/if 10                               /* Wireless Client first interface
        enable
        addr 192.168.10.253                 /* Switch 2 uses: 192.168.10.252
        vlan 10
/cfg/l3/if 110                              /* Internet first interface
        enable
        addr 192.168.110.253                /* Switch 2 uses: 192.168.110.252
        vlan 110
/cfg/l3/if 126                              /* VRRP advertisement interface
        enable
        addr 1.1.1.253                      /* Switch 2 uses: 1.1.1.252
        vlan 4000
```

Configure Static Routes

Define the static routes to be used for the Wireless Client and Internet side:

```
/cfg/l3/route
      add 30.1.0.0 255.255.0.0 169.244.144.100    /* Wireless Client static routes
      add 130.1.0.0 255.255.0.0 169.234.144.100   /* Internet static routes
```

VRRP Router Configuration

1. Globally turn on VRRP and configure each VRRP router for the internal server group, the wireless client, and the Internet:

```
/cfg/l3/vrrp/on
/cfg/l3/vrrp/vr 1    /* VRRP router for internal blade servers application group
        enable
        vrid 1
        if 1
        addr 169.254.144.254
/cfg/l3/vrrp/vr 2    /* VRRP router for Wireless Client second interface
        enable
        vrid 2
        if 2
        addr 169.244.144.254
/cfg/l3/vrrp/vr 3    /* VRRP router for Internet second interface
        enable
        vrid 3
        if 3
        addr 169.234.144.254
/cfg/l3/vrrp/vr 10   /* VRRP router for Wireless Client first interface
        enable
        vrid 10
        if 10
        addr 192.168.10.254
/cfg/l3/vrrp/vr 110  /* VRRP router for Internet first interface
        enable
        vrid 110
        if 110
        addr 192.168.110.254
```

2. Put all the VRRP routers into a VRRP group so all VRRP advertisements can be confined to one Interface:

```
/cfg/l3/vrrp/group
        enable
        vrid 1
        if 126 restricted   /* All VRRP advertisements will use Interface 126
        garp 250
```

3. Configure the VRRP group to use ports for calculating the change of VRRP Master role:

```
/cfg/l3/vrrp/group/track
         ports enable
/cfg/l3/vrrp/group/track/portlist 1/addport EXT1-EXT2
/cfg/l3/vrrp/group/track/portlist 2/addport EXT9-EXT10
```

Layer 2 Switch Configuration (Switch 3 and Switch 4)

The following steps show the Layer 2 switch configuration for Switch 3. Whenever there is a difference, the respective commands are marked with Switch 4 for the Layer 2 switch connecting to the VRRP Backup.

Note: The following example configuration is one of the more commonly used and may differ from yours, depending on your specific needs.

Configure Port Link Behavior

1. Define the internal port link behavior for the application server blades following the same parameters as in the primary chassis to accept VLAN tagged packets.

```
/cfg/port INT1/gig
        speed 10000
        auto off
        fctl none
        tag enable
/cfg/port INT2/gig
        speed 10000
        auto off
        fctl none
        tag enable
/cfg/port INT3/gig
        speed 10000
        auto off
        fctl none
        tag enable
```

2. Define the inter-chassis external ports link behavior to accept VLAN tagged packets:

```
/cfg/port EXT4
        tag enable
/cfg/port EXT5
        tag enable
/cfg/port EXT6
        tag enable
/cfg/port EXT7
        tag enable
```

Configure VLANs onto the Internal and External Ports

1. Create VLAN 10 to handle the wireless client traffic coming into the inter-chassis trunk ports connecting to the primary chassis (EXT4 to EXT7) and the blade server ports (INT1 to INT3):

```
/cfg/l2/vlan 10
        enable
        add INT1-INT3
        add EXT4-EXT7
```

2. Create VLAN 110 to handle the Internet traffic coming into the inter-chassis trunk ports connecting to the primary chassis (EXT4 to EXT7) and the blade server ports (INT1 to INT3):

```
/cfg/l2/vlan 110
        ena
        add INT1-INT3
        add EXT4-EXT7
```

3. Assign VLAN 3000 for the internal blade servers health check and communication. Set this VLAN as the untagged VLAN for those related ports:

```
/cfg/l2/vlan 3000
        ena
        add INT1-INT3
        add EXT4-EXT7
/cfg/port INT1
        pvid 3000
/cfg/port INT2
        pvid 3000
/cfg/port INT3
        pvid 3000
/cfg/port EXT4
        pvid 3000
/cfg/port EXT5
        pvid 3000
/cfg/port EXT6
        pvid 3000
/cfg/port EXT7
        pvid 3000
```

Turn off Spanning Trees

Since this pair of iFlow switches are within a controlled environment, turn off spanning trees to avoid unnecessary BPDUs:

```
/cfg/l2/nostp enable
/cfg/l2/stg 1/off
```

Create Trunks

Create the inter-chassis trunk:

```
/cfg/l2/trunk 1     /* Inter-chassis trunk
        enable      /* Default is l3thash is SIP+DIP, and it will hash in SIP+DIP.
        add EXT4
        add EXT5
        add EXT6
        add EXT7
```

Routing

The iFlow Director supports the following options for routing traffic from the application servers to the appropriate client or Internet destination:

- Standard routing using static routes
- Standard routing using OSPF.
- Policy-based routing, using multiple VLANs.

Note: These routing features may be used in combination.

Static Routes

Static routes are a standard Layer 3 routing feature. Using this feature, you can define forwarding routes as follows:

- For IPv4 networks:

```
>> # /cfg/l3/route/add <destination> <mask> <gateway> [<interface>]
```

- For IPv6 networks:

```
>> # /cfg/l3/route6/add <destination> <prefix length> <gateway> [<interface>]
```

These commands specify that all traffic to the destination network (the range formed by the destination network IP address and mask or prefix) is forwarded through the designated route address.

Using the example from Figure 5 on page 39, IPv4 static routes would be configured as follows:

```
/cfg/l3/route
     add 10.0.0.0 255.255.0.0 192.168.1.2
     add 201.0.0.0 255.255.255.0 172.20.1.2
apply
save
```

With this configuration, traffic to the client network (10.0.0.0/16) is forwarded through the access router (192.168.1.2), and traffic to the Internet (201.0.0.0/24) is forwarded through the edge router (172.20.1.2).

Note: When a PBR ACL is in the Drop state, traffic is dropped regardless of the static routes or OSPF configuration.

OSPF Routing

OSPF is another standard Layer 3 routing feature. On the iFlow Director, the VRRP standby switch has OSPF interfaces with cost incremented to 65535. This advertises the master switch as the best router to use for OSPF routes towards access and edge routers.

For detailed information on OSPF, refer to your full product documentation.

Policy-Based Routing

The iFlow Director can perform Layer 3 policy-based routing for traffic from the application servers. When you specify a next hop IP address, the iFlow Director can control the destination of outbound traffic with settings provided using the policy-based routing commands. Policy-based routing takes precedence over other configured static routes or OSPF routing.

To use policy-based routing for the example shown in Figure 5 on page 39 (instead of static routes as shown in "Static Routes" on page 55), we will assume that the client network is assigned to VLAN 10 and the Internet is assigned to VLAN 20.

In this scenario, the following configuration is required in addition to the basic configuration.

1. Create a PBR ACL and set the next hop IP address to the IP address of an adjacent router.

   IBM N/OS CLI:

   ```
   /cfg/acl/acl 110/target/dest nexthop
   /cfg/acl/acl 110/target/nexthop 192.168.110.2
   ```

   ISCLI:

   ```
   access-control list 100 target destination nexthop 192.168.1.2
   ```

2. Define a qualifier for the PBR ACL. In this example, the qualifier is defined as VLAN 10 with a VLAN mask of 0xfff.

   IBM N/OS CLI:

   ```
   /cfg/acl/acl 110/ethernet/vlan 10 0xfff
   ```

   ISCLI:

   ```
   access-control list 100 ethernet vlan 10 oxfff
   ```

3. Set the filter action to "redirect".

   IBM N/OS CLI:

   ```
   /cfg/acl/acl 110/action redirect
   ```

   ISCLI:

   ```
   access-control list 100 action redirect
   ```

4. Enable statistics.

   IBM N/OS CLI:

   ```
   /cfg/acl/acl 110/stats enable
   ```

   ISCLI:

   ```
   access-control list 100 statistics
   ```

5. Apply the configured PBR ACL to applicable server-connected ports.

   IBM N/OS CLI:

   ```
   /cfg/acl/group 5/add 110
   /cfg/port INT1/acl/add grp 5
   /cfg/port INT2/acl/add grp 5
   /cfg/port INT3/acl/add grp 5
   /cfg/port INT4/acl/add grp 5
   ```

   ISCLI:

   ```
   interface port INT1-INT4
   access-control group 5
   exit
   ```

6. Configure PBR health checking.

   The first command sets the number of seconds between health checks, the second sets the number of consecutive failed health checks required to determine that the link is down, and the third sets the number of consecutive positive health checks required to determine that the link is up.

   IBM N/OS CLI:

   ```
   /cfg/slb/pbr/interval 5
   /cfg/slb/pbr/retry 2
   /cfg/slb/pbr/restr 2
   ```

   ISCLI:

   ```
   slb pbr interval 5
   slb pbr retry 2
   slb pbr restore 2
   ```

Note: The current release of IBM iFlow Director does not support PBR on IPv6; the next release will have this feature.

7. (IBM N/OS CLI only) Apply and save the configuration

   ```
   apply
   save
   ```

Next Hop VLAN Forwarding

The iFlow Director can be configured to forward traffic entering the switch to a specific next hop VLAN. This involves enabling the option and programming the the ECMP table entries hosting the SLB application with the next hop VLAN interface index. Traffic heading into the iFlow Director ports either exits untagged or with the PVID tag, depending on the port PVID configuration.

When this option is enabled, the port-based operation that prevents the VLAN translation on the iFlow ports is disabled. By default, this option is set to 1. This preserves the ingress VLAN ID.

In the following example, the wireless VLAN is 1000 and the Internet VLAN is 2000. To set up next hop VLAN forwarding with these VLANs:

1. Enable next hop VLAN forwarding.

   ```
   /cfg/slb/app 1/nhvlan 3000
   ```

   In this example, 3000 is the VLAN number assigned as the next-hop VLAN for the SLB application group of the local or remote internal server blades.

2. Create and enable the VLAN being used for next hop VLAN forwarding.

   ```
   /cfg/l2/vlan 3000
        ena
   ```

3. Add the bladeserver ports to this VLAN.

   ```
   /cfg/l2/vlan 3000
        add INT1-INT3
   ```

4. Set up local blade server ports with the next hop VLAN PVID so they can receive VLAN 3000 untagged packets.

   ```
   /cfg/port INT1
        pvid 3000
   /cfg/port INT2
        pvid 3000
   /cfg/port INT3
        pvid 3000
   ```

5. Set up the next hop VLAN interface.

   Note: An interface is usually configured for the blade server health check and communication purposes and is on the same next hop VLAN.

   ```
   /cfg/l3/if 1
        ena
        addr 169.254.144.253
        mask 255.255.255.0
        broad 169.254.144.255
        vlan 3000
   ```

6. Create interfaces and static routes based on the VLAN routing environment.

```
/cfg/l3/route
     add 30.1.0.0 255.255.0.0 169.244.144.100        /* For Wireless-side
     add 130.1.0.0 255.255.0.0 169.234.144.100       /* For Internet-side
/cfg/l3/if 2                                         /* For Wireless-side
     ena
     addr 169.244.144.253
     vlan 1000
/cfg/l3/if 3                                         /* For Internet-side
     ena
     addr 169.234.144.253
     vlan 2000
```

7. Since the real servers are not handling VLAN tagged packets, set the retain ingress VLAN flag to `disable` (the default value) on the external ports on which the incoming traffic is arriving.

```
/cfg/port EXT1-EXT2/retivlan disable        /* For wireless side
/cfg/port EXT9-EXT10/retivlan disable       /* For Internet side
```

Notes:

- “Retain ingress VLAN flag” is a specific port-based setting. If it is enabled, the redirect ACL applied on those external ports from the wireless or Internet client will forward traffic with VLAN tags onto the real server application group. There is no need for the real server ports to be configured with the same VLAN membership as the Wireless or Internet sides.
- If retain ingress VLAN flag is disabled (the default value), the real server ports must be assigned to the next hop VLAN (VLAN 3000 in this exampe), which must not be used by the wireless or Internet external ports. The real server ports can be configured as untagged, tagged, or tagpvid enabled. Traffic from wireless or Internet clients with redirect ACL will be forwarded with all VLANs untagged onto the real server application group. When the real servers send back the packets to the IBM iFlow switch, they will be VLAN tagged respectively with the correct VLAN based on the IP subnet.
- For this example, the real servers are not handling VLAN tagged packets, hence the retain ingress VLAN flag is set to `disable`.

Server Health-Check Options

Health checking allows you to verify server accessibility in an application group. As content grows and information is distributed across servers in an application group, content health-checks ensure end-to-end availability.

Use the following command to configure the health-check option for an SLB group:

```
>> # /cfg/slb/group <SLB group number>/health
```

The following server health-checks are available in iFlow Director:

- ICMP ping
- TCP (Transmission Control Protocol)
- TCP script-based health checks
- Application-specific health-checks with configurable TCP Port Numbers (for example, HTTP could use port 8080 instead of 80)
  - HTTP (HyperText Transfer Protocol)
  - HTTP Head message
  - SSL (Secure Socket Layer)
  - SMTP (Simple Mail Transfer Protocol)
  - SIP-register (Session Initiated Protocol)
  - DNS (Domain Name Server)
  - ARP (Address Resolution Protocol)

Note: IPS and link health checking do not apply for Layer 3 hash applications.

Each real server's health can be checked by type, service port, and content configured for the application group to which it belongs. Note that the health check packets are only sent out through the switch port configured for the real server, not all ports.

The health-check setting applies to every real server in the group.

Health checks are performed on intervals of 1-60 seconds (the default is 2 seconds), specified under the real server configuration, as follows:

```
>> # /cfg/slb/real <server number>/inter <1-60, or 400 ms>
```

The 400 millisecond interval is available only for ping health checks.

Health Check Configuration Example

The following example shows the general steps used to configure health checks for an existing SLB group (configured as shown in "Configure Server Load Balancing" on page 40).

1. Configure the health check type for the SLB group.

```
/cfg/slb/group <1-84>/health [ping|tcp|http|httphead|smtp|ssl| udpdns|arp|sip-
register|script <script number>]
```

2. Configure the service port used to send data to the server. The `svport` option is ignored for health checks that do not require it.

   ```
   /cfg/slb/group <1-14>/svport <0-65534>
   ```

3. Define the content string (optional, for types HTTP, HTTPhead, SMTP, UDPDNS, SIP-register).

   ```
   /cfg/slb/group <1-14>/content <text string>
   ```

The `content` option is used to configure type-specific data, such as a URL (`www.ibm.com`) for DNS health checks, a filename and file path (`index.html`) for HTTP health checks, a username for SIP register health checks, an email address (`joe@ibm.com`) for SMTP health checks, and so on. The maximum string length for the content field is:

- SMTP = 1-64 characters
- SIP =1-32 characters
- HTTP/HTTP Head/UDPDNS = 1-127 characters

Scriptable Health Checks

Scriptable health checks allow you to create a customized TCP health check for service types not specifically implemented. The script can mimic a standard service request and response sequence, and check if the behaviors match what is expected.

Use the following command to configure a health check script:

```
>> # /cfg/slb/script <1-16>
```

Consider the following guidelines when you configure a health-check script:

- The script must begin with the `open` command to a specific service port.
- The `send` command is used to define a text string to send to the specified service port.
- The `expect` command is used to define a text string to verify that the first part of the text string received matches the string configured in the script. The script fails if the exact string is not included in the server response.
- It is recommended that each script contain only one `open` command.
- The switch considers the health check to be successful only if the entire script has been executed and properly matched.
- The following commands are ignored when the switch performs a scriptable health check:

  ```
  /cfg/slb/group <1-14>/svport
  /cfg/slb/group <1-14>/content
  ```

The following example script checks the HTTP service:

```
/c/slb/script 1
open "80"
send "GET / HTTP/1.0\\r\\n\\r\\n"
expect "HTTP/1.1 200"
close
```

Note: The backslash (\ ) is used as an escape character in the script. If you want to include a backslash in the text string, you must preface it with a second backslash.

Using VRRP for iFlow Redundancy

Where high-availability networks are required, dual iFlow Directors can be installed in the BladeCenter chassis. Using VRRP, one switch is in active mode and the other is in standby mode, ready to take over from the active switch in the event of a component or link failure.

Consider the following example:

Figure 7. Application Network with iFlow Director

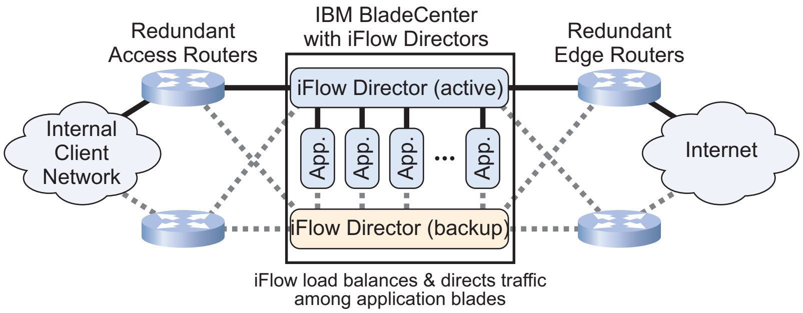

See your product documentation for details on VRRP configurations.

Note: When a VRRP configuration is enabled, the access router and edge router should use the iFlow Director's virtual router (VIR) address as the next hop, instead of the switch IP interface address.

Layer 3 Hash Table Synchronization

The iFlow Director can be configured to automatically synchronize the hash across two peer switches in a redundant, VRRP environment. Both switches must have the same port memberships, and both must have the same Layer 3 health check type.

When hash synchronization is enabled, VRRP is used to match the standby switch's hash table to that of the active switch. Synchronization is performed periodically, or whenever the active switch changes its hash table.

The following command must be executed on each switch to enable hash synchronization:

```
/cfg/slb/sync ena
```

VRRP Tracking Extensions

The iFlow Directors use the VRRP election process to determine which switch will be the master, and which the standby. This election process is based the priority value assigned to each switch. Normally, this priority value is assigned by the administrator to indicate the preference that a specific switch perform as master.

However, depending on the network, there may be situations where the master is still operational even though it has lost access to an important next-hop router or has lost connection on all of its uplink ports. In such situations, it can be useful to preempt the election and promote the standby switch to take over.

To ensure that the master switch offers optimal network performance to neighboring devices, VRRP tracking allows the pre-configured priority value to be dynamically modified by the switch according to various switch health criteria. In addition to regular VRRP tracking options such as IP interfaces and physical ports (see your complete switch documentation), iFlow Director permits tracking for multiple port groups and for next-hop routers.

VRRP priority tracking can be configured from the Virtual Router Group Priority Tracking menu (`/cfg/l3/vrrp/group/track`).

Port Group Tracking

Although standard VRRP permits tracking the physical ports within the VLAN of a particular IP interface, it may be preferable to configure the iFlow Director to track a set of ports regardless of VLAN. This is particularly useful when using the VRRP Virtual Router Group feature, since virtual router instances may belong to different VLANs. It may also be used to assist failover when all ports for a particular service or function are lost, though other ports remain active.

iFlow Director supports tracking two port groups. The switch's priority value is increased by a configurable amount for each active port in each group. However, to assist failover in the event either group is compromised, if all ports in one or both groups are down, the switch's priority value will be reduced to its original base priority value.

Router Tracking

iFlow Director supports tracking routers. The switch's priority value is increased by a configurable amount for each reachable adjacent, next-hop router. This is useful for assisting failover when the master switch is up, but has lost router access. The base priority is increased for each available router on the access or Internet side of the network. However, if all routers on one or both sides are down, the switch's priority value will be reduced to its original base priority value.

Port Group and Router Tracking Combination

Port group tracking and router tracking can be used in combination. When both features are enabled, the base priority value will be increased by the configured weights for each port and router, as long as at least one item is available in each configured port group or router zone. If any configured group has zero items, the switch's priority value will be reduced to its original base priority value.

Additional VRRP Tracking Considerations

When configuring VRRP tracking, consider the following guidelines:

- For tracking to effect virtual router selection, preemption must be enabled.
- The virtual router priority value will not exceed 254, regardless of port group or router tracking increments.
- Management ports cannot be included in port tracking groups.
- Each physical port can be assigned to no more than one port tracking group. No specific port can be included in both port tracking groups.
- If tracking ports in a trunk, all ports belonging to the same static or dynamic trunk should be placed in the same port tracking group.

- Unlike the default VLAN port tracking, port group tracking is based only upon the physical link state of the ports, not on the forwarding state. Ports that are part of a dynamic trunk are considered up when LACP is up on the port. All ports of a static or dynamic trunk are considered up if at least one of the member ports is active.

Information Commands

Two useful information commands are as follows:

- ECMP (`/info/slb/ecmp`) displays the hash table objects.
- Binding (`/info/slb/bind` *<IPv4 address>* or `/info/slb/bind6` *<IPv6 address>*) returns the internal port number that the hash would select for the specified IP address.

Layer 3 Hashing Configuration Guidelines

When configuring iFlow Director, consider the following guidelines:

- When the configuration is applied, changes to SLB or ACL-dependent parameters might result in temporary traffic disruptions because the apply operation causes the switch to re-initialize these features. You might experience the following types of disruptions:
  - Temporary packet discards
  - Temporary traffic flooding via Layer 2
  - Temporary hash disruptions
  - Inaccurate statistical accounting
- If you use SNMP to set the next boot configuration file to factory default, and perform a save operation, the switch automatically changes the next-boot setting to the active configuration block instead of the factory configuration block. It is recommended that you immediately reboot the switch after setting the configuration block to factory default.
- To avoid errors when using SNMPwalk, use the `-Cc` option with the SNMPwalk command. For example:

```
snmpwalk -Cc -v 1 -t 60 -c public -m ALL -M $miblocation
10.13.5.103 $1
```

Statistics

- Statistics for ECMP routes over the management interface are not supported.
- The following command displays statistics for each real server port in each SLB application when the SLB applications are configured for aggregation:

```
/stats/slb/app all
```

Access Control Lists (ACLs)

When configuring ACLs, consider the following guidelines:

- SLB ACL distribution takes precedence over port mirroring actions. If you configure port mirroring, packets are not mirrored if an SLB ACL is configured on the port.
- SLB ACL statistics (`/stats/acl/dump`) show ACL matches, but not necessarily actions. In some cases, statistics increment when ACLs matched the packet headers, but no redirection takes place. For example, a packet might be dropped on ingress, so no redirection occurred.
- When a 1Gb port is configured with metering but does not have flow control enabled, the switch can drop Out-of-Profile packets normally. However, when flow control is enabled on the port, it will take precedence on the port to limit the rate of traffic at 1Gbps. In this case, metering does not take effect if the committed rate is greater or equal to 1,000,000 kbps (1Gbps).

Health Checks

- When writing a health-check script, you must use numeric values for service ports, because aliases are not accepted (for example, use `80` instead of `http`).
- When configuring health check content within an SLB group, the double quotation mark (") is used as a special escape character. Therefore, it cannot be included in the body of the content.

- Health checks are performed asynchronously on peer VRRP switches. This can cause hash tables to become out-of-sync between the switches unless hash synchronization is configured (see "Layer 3 Hash Table Synchronization" on page 63).
- Health checks can cause an application to go on- or offline repeatedly (flap) if the health check interval or retry values are too small. To avoid this issue, increase the value of the real server health check interval (`/cfg/slb/real` *x*`/inter`) or the retry option (`/cfg/slb/real` *x*`/retry`).

Chapter 4. Command Reference

The following sections include detailed information about iFlow Director commands. For more detailed information about switch commands, refer to the *Command Reference* for your switch.

The following topics are discussed in this chapter:

/info

Information Menu

```
[Information Menu]
     sys      - System Information Menu
     l2       - Layer 2 Information Menu
     l3       - Layer 3 Information Menu
     slb      - Server Load Balancing Information Menu
     qos      - QoS Menu
     acl      - ACL Information Menu
     rmon     - Show RMON information
     link     - Show link status
     port     - Show port information
     transcvr - Show Port Transceiver status
     swkey    - Show enabled software features
     dump     - Dump all information
```

The following table briefly describes the information commands used with the iFlow Director.

Table 1. Information Options (/info)

Command Syntax and Usage
`slb` Displays the Server Load Balancing Information menu. To view menu options, see page 71.
`acl` Displays the ACL Information menu, which allows you to view the current configuration profile for each Access Control List (ACL) and ACL Group. For more information, see your *Command Reference*.
`swkey` Displays license information for iFlow Director.

/info/slb
SLB Information

```
[Server Load Balancing Information Menu]
     real     - Show real server information
     group    - Show real server group information
     app      - Show application management information
     bind     - Show real server selected by hash with ipv4 address
     bindv6   - Show real server selected by hash with ipv6 address
     ecmp     - Show ECMP table entries
     pbr      - Show PBR health check configuration and information
     dump     - Show all server load balancing information
```

The following table describes the Server Load Balancing (SLB) information.

Note: The `bind`, `bindv6`, and `ecmp` options are only applicable when `/cfg/slb/method` is set to `ecmp`.

Table 2. SLB Information Options (/info/slb)

Command Syntax and Usage
`real` *<real server number (1-84)>* Displays real server information.
`group` *<group number (1-14)>* Displays information for the selected real server group.
`app` *<application number (1-14)>* Displays information for the selected application.
`bind` *<IPv4 address> <mask>* Displays real server port selected by ECMP hashing on an IPv4 address.
`bindv6` *<IPv6 address>* Displays real server port selected by ECMP hashing on an IPv6 address.
`ecmp` Displays the ECMP hashing table.
`dump` Displays all SLB information for the switch.

/stats/slb
SLB Statistics

```
[Server Load Balancing Statistics Menu]
     app      - SLB Application statistics
     clear    - Clear SLB Application statistics
```

The following table describes the Server Load Balancing (SLB) statistics commands.

Table 3. SLB Statistics Options (/stats/slb)

Command Syntax and Usage
`app` *<application number (1-14)>*`|all` Displays statistics for the selected application.
`clear` *<application number (1-14)>*`|all` Clears SLB application statistics.

/cfg/slb
SLB Configuration

The following sections describe the configuration commands that are specific to iFlow Director firmware. For more information about configuration commands, see your *Command Reference*.

```
[Server Load Balancing Menu]
     app      - Application Menu
     real     - Real Server Menu
     group    - Real Server Group Menu
     script   - Scriptable Health Check Menu
     pbr      - Policy Based Routing Health Check Configuration Menu
     method   - Set SLB hashing method
     mod      - Enable/disable Manual OSP Distribution
     sync     - Enable/disable inter-switch synchronization
     on       - Globally turn Server Load Balancing ON
     off      - Globally turn Server Load Balancing OFF
     hcwtint  - Set server health check delay time on failover
     initarp  - Set server ARP delay time on failover
     cur      - Display current Server Load Balancing configuration
```

Note: The `sync` option is only applicable when `/cfg/slb/method` is set to `ecmp`. The `mod` option is only applicable when `/cfg/slb/method` is set to `trunk`.

Table 4. Server Load Balancing Configuration Options (/cfg/slb)

Command Syntax and Usage
`app` *<application number (1-14)>* Displays the Application Management menu. To view menu options, see page 74.
`real` *<real server number (1-84)>* Displays the menu for configuring real servers. To view menu options, see page 77.
`group` *<real server group number (1-14)>* Displays the menu for placing real servers into real server groups. To view menu options, see page 79.
`script` *<1-16>* Displays the Scriptable Health-Check menu. To view menu options, see page 81.
`pbr` Displays the Policy Based Routing Health Check Configuration menu. To view menu options, see page 83.
`method ecmp|trunk` Configures the SLB hashing method. The default setting is `ecmp`.
`mod e|d` Enables or disables manual hash table distribution. When enabled, you must manually set the hash port matrix. For more information, see "Operations-Level SLB Application Options" on page 115. The default setting is `disabled`. **Note**: When this option is enabled, health checks by the switch are disabled.
`sync e|d` Enables or disables hash table synchronization between two peer switches. This setting is disabled by default.
`on` Globally turns on Server Load Balancing and Application Redirection.
`off` Globally disables Server Load Balancing. All configuration information will remain in place (if `applied` or `saved`), but the processes will no longer be active in the switch.
`hcwtint` *<0-60>* Sets the health check delay time (in seconds) on failover. The default value is `4`.

Table 4. Server Load Balancing Configuration Options (/cfg/slb) (continued)

Command Syntax and Usage
`initarp` *<1-60>* Sets the ARP delay time (in seconds) on failover. The default value is `2`.
`cur` Displays the current Server Load Balancing configuration.

/cfg/slb/app *<application number>*
SLB Application Management Configuration

```
[Application Management 1 Menu]
     errhand  - Application Error Handling Menu
     name     - Set application name
     group    - Set real server group
     mode     - Set application mode
     nhvlan   - Set next hop VLAN ID
     aggr     - Set application aggregation
     expand   - Enable/disable application hash expansion
     del      - Delete application
     cur      - Display current application configuration
```

Table 5. Application Management Options (/cfg/slb/app

Command Syntax and Usage
errhand Displays the Application Error Handling menu. To view menu options, see page 76.
`name` *<string, maximum 31 characters>*|`none` Defines an alias for each application.
`group` *<group number (1-14)>*|`none` Set the SLB group ID to be managed by this application. The default value is `none`.
`mode inline`|`hairpin` Sets the application operation mode to the following: – inline (default) – application will receive traffic from a local port and forward traffic inline to a remote port. – hairpin – application will receive and forward traffic on its local port.
`nhvlan` *<VLAN number (1-4094)>* Sets the next hop VLAN ID.

Table 5. Application Management Options (/cfg/slb/app (continued)

Command Syntax and Usage
`aggr e|d` Enables or disables application aggregation. When enabled, this application is designated as a sub-application that will be aggregated with other sub-applications to form a super-application. The default setting is `disabled`.
`expand e|d` Enables or disables hash expansion for applications that have 8 or fewer real servers. When enabled, the number of trunks required to support the application is increased to 8 and the number of hash table entries is increased to 64. The default setting is `disabled`.
`del` Deletes this Application Management instance.
`cur` Displays the current configuration parameters for this Application Management instance.

Note: The `errhand`, `aggr`, and `expand` options are only applicable when `/cfg/slb/method` is set to `trunk`.

/cfg/slb/app *<application number>*/errhand

Application Error Handling Configuration

```
[Application Error Handling Menu]
  xdelay   - Set transition delay period
  preempt  - Enable/disable preemption for failback
  remap    - Enable/disable traffic redistribution
  cur      - Display current application error handling configuration
```

Note: The options in this menu are only applicable when `/cfg/slb/method` is set to `trunk`.

Table 6. Error Handling Options (/cfg/slb/app/errhand)

Command Syntax and Usage
`xdelay` *<0-60>* Sets the application transition delay (in seconds) to define the period of instability while health checking real servers in the application group during the following conditions: – When creating a new SLB application for the first time. – When booting up with an existing SLB application configuration. – When re-populating the application group with an empty port membership. – When changing the real server membership in an SLB application during runtime. – When applying SLB-related configuration changes. The default value is 20 seconds.
`preempt e|d` Enables or disables pre-emption, which provides the ability to fail back to the primary application servers in the event that service is restored while the corresponding backup server is active. The default value is `enabled`.
`remap e|d` Enables or disables traffic redistribution, or re-mapping of the server group load distribution in the event of server failures within an application group. The default value is `disabled`.
`cur` Displays the current error handling configuration for this application.

/cfg/slb/real *<server number>*
Real Server SLB Configuration

```
[Real Server 1  Menu]
    name     - Set real server name
    port     - Set the port the server connects to
    rip      - Set IP addr of real server
    inter    - Set interval between health checks
    retry    - Set number of failed attempts to declare server DOWN
    restr    - Set number of successful attempts to declare server UP
    backup   - Set backup real server
    ena      - Enable real server
    dis      - Disable real server
    del      - Delete real server
    cur      - Display current real server configuration
```

This menu is used for configuring information about real servers that participate in a server pool for Server Load Balancing or Application Redirection. The required parameters are:

- Real server port number
- Real server enabled (disabled by default)

Table 7. Real Server Configuration Options (/cfg/slb/real)

Command Syntax and Usage
`name` *<string, maximum 31 characters>*|`none` Defines an alias for each real server. This will enable the network administrator to quickly identify the server by a natural language keyword value.
`port` *<port alias or number>* Defines the port to which this real server connects. The default setting is `INT1`.
`rip` *<real server IP address>* Sets the IP address of the real server, in dotted decimal format to define the IP address for Layer 3 health checks on the specified port number to determine if the server is up. The administrator will be warned if the server does not respond. **Note**: When you define a Real Server IP address, configuring the Real Server Port is optional. When using ECMP, there is a dynamic check regarding the port where the server is connected, so there is no need to specify the connected port (real server port).
`inter` *<1-60 seconds, or 400 for milliseconds>* Sets the interval between real server health verification attempts. Determining the health of each real server is a necessary function for SLB. The `inter` option lets you choose the time between health checks. The range is from 1 to 60 seconds, or 400 milliseconds. The default interval is 2 seconds.
`retry` *<number of consecutive health checks (2-63)>* Sets the number of failed health check attempts required before declaring this real server inoperative. The default value is 4 attempts.

Table 7. Real Server Configuration Options (/cfg/slb/real)

Command Syntax and Usage
`restr` *<number of consecutive health checks (1-63)>* Sets the number of successful health check attempts required before declaring a server operational. The default value is 8 attempts.
`backup` *<real server number (1-84)>* | `none` Sets the real server used as the backup server for this real server. To prevent loss of service if a particular real server fails, use this option to assign a backup real server number. Then, if the primary real server becomes inoperative, the switch will activate the backup real server. The same backup server may be assigned to more than one real server within the same application group.
`ena` You *must* perform this command to enable this real server for SLB service.
`dis` Disables this real server from SLB service. This option *does not* perform a graceful server shutdown.
`del` Deletes this real server from the SLB configuration. This removes the real server from operation within its real server groups. Use this command with caution, as it will delete any configuration options that have been set for this real server. This option *does not* perform a graceful server shutdown.
`cur` Displays the current configuration information for this real server.

/cfg/slb/group <*real server group number*>
Real Server Group SLB Configuration

```
[Real server group 1 Menu]
      health  - Set health check type
      name    - Set real server group name
      realthr - Set real server failure threshold
      svport  - Set health check service port
      content - Set health check content
      add     - Add real server
      rem     - Remove real server
      del     - Delete real server group
      cur     - Display current group configuration
```

This menu is used for combining real servers into real server groups. Each real server group must consist of all the real servers which provide a specific service for load balancing. Each group must consist of at least one real server. Real server groups are used both for Server Load Balancing and Application Redirection.

Table 8. Real Server Group Configuration Options (/cfg/slb/group)

Command Syntax and Usage
`health link|ping|ips|tcp|http|httphead|smtp|ssl|udpdns|arp|sip-register|script`<*n*> Sets the type of health checking performed, as follows: - `link` - `ping` **Note:** The `link` and `ping` options are only applicable when `/cfg/slb/method` is set to `trunk`. – `ips` (Intrusion Protection System, a Layer 2 application-specific health check) - `tcp` - `http` - `httphead` - `smtp` - `ssl` - `udpdns` - `arp` - `sip-register` - `script` The default setting is `link`.
`name` <*1-31 characters*>`|none` Defines an alias for each Real Server Group. This will enable the network administrator to quickly identify the server group by a natural language keyword value.

Table 8. Real Server Group Configuration Options (/cfg/slb/group) (continued)

Command Syntax and Usage
`realthr` *<(1-84, 0 for disabled)>* Specifies the real server failure threshold for triggering an application bypass action, as defined by the following command: `/cfg/acl/acl #/target/apperr` For more details, see page 100. If the number of operational servers is equal to or less than the threshold, the bypass action is initiated if it is defined. When the number of operational ports rises above the threshold, the application reverts to normal behavior. A SYSLOG message is generated when the # of operational real server ports used within the SLB application group hash matrix have failed down to a non-zero threshold value. A value of 0 (zero) is used to disable the SYSLOG alert. The default threshold value is 0.
`svport` *<0-65534>* Configures the service port to send health-check data on to the server, if necessary. This option is used only for health-check types that require a service port, as follows: – TCP – HTTP – HTTP Head – SMTP – SSL – DNS – SIP-register The default setting is 0 (zero).
`content` *<filename>*|*//<host>/<filename>*|`none` Defines supporting information for the selected health-check type, as follows: – `udpdns`: Define a URL, such as `www.ibm.com` (1-127 characters) – `http`: Define a filename or path, such as `index.html` (1-127 characters) – `sip-register`: Define a username, such as `Smith` (1-32 characters) – `smtp`: Define an email address, such as `joe@ibm.com` (1-64 characters) The default value is `none`.
`add` *<real server number (1-84)>* Adds a real server to this real server group.
`rem` *<real server number (1-84)>* Remove a real server from this real server group.
`del` Deletes this real server group from the SLB configuration.
`cur` Displays the current configuration parameters for this real server group.

`/cfg/slb/script` *<1-16>*
Scriptable Health Check Configuration

```
[Health Script 1 Menu]
     open     - Add open command to end of script
     send     - Add send command to end of script
     expect   - Add expect command to end of script
     close    - Add close command to end of script (TCP only)
     rem      - Remove last command from script
     del      - Delete script
     cur      - Display current script configuration
```

The maximum size of a script is defined by the total number of characters allowed in the script (6144 characters).

Table 9. Scriptable Health-Check Options (/cfg/slb/script)

Command Syntax and Usage
`open` *<0-65534>* Initiates a TCP three-way handshake to the selected service port. Some well known port numbers/names are: **Number** **Name** 20 ftp-data 21 ftp 23 telnet 25 smtp 53 dns 69 tftp 80 http 443 https 1812 http
`send` *<1-502 characters>* Sends a text string to the server port.
`expect` *<1-502 characters>* Waits for a string from the server and checks if the argument matches the first part of the string received. This command fails if the expected string is not found in the server response. If a match is found, the next script command is executed. If a match is not found, the system waits for the next packet to arrive. If no matching string arrives by the next health check interval, the script health check fails and the current connection is terminated.
`close` Closes an open TCP connection.
`rem` Removes the last line from the script.
`del` Deletes the script.

Table 9. Scriptable Health-Check Options (/cfg/slb/script) (continued)

Command Syntax and Usage
`cur` Displays the current health-check script parameters.

/cfg/slb/pbr
Policy-Based Routing Health Check Configuration

```
[Policy Based Routing Menu]
     interval - Set interval between Router health checks
     retry    - Set number of failed attempts to declare router DOWN
     restr    - Set number of successful attempts to declare router UP
     cur      - Display current Policy Based Router configuration
```

Table 10. Policy-Based Router Mapping Options (/cfg/slb/pbr)

Command Syntax and Usage
`interval` *<1-60>* Configures the number of seconds between health checks. The default value is `5` seconds.
`retry` *<1-63>* Configures the number of consecutive failed health checks required to determine that the next hop router is down. The default value is `1`.
`restr` *<1-63>* Configures the number of consecutive positive health checks required to determine that the next hop router is up. The default value is `1`.
`cur` Displays the current policy-based routing configuration.

/cfg/port *<port alias or number>*
Port iFlow Configuration

```
[Port INT1 Menu]
     errdis   - ErrDisable Menu
     gig      - Gig Phy Menu
     udld     - UDLD Menu
     oam      - OAM Menu
     aclqos   - Acl/Qos Configuration Menu
     stp      - STP Menu
     8021ppri - Set default 802.1p priority
     pvid     - Set default port VLAN id
     name     - Set port name
     bpdugrd  - Enable/disable BPDU Guard
     dscpmrk  - Enable/disable DSCP remarking for port
     rmon     - Enable/disable RMON for port
     learn    - Enable/disable FDB Learning for port
     tag      - Enable/disable VLAN tagging for port
     tagpvid  - Enable/disable tagging on pvid
     tagiskip - Enable/disable skipping ingress VLAN tag enforcement
     tageskip - Enable/disable skipping egress VLAN tag enforcement
     arpmp    - Enable/disable ARP copy to MP
     retivlan - Enable/disable port retain ingress vlan
     ecmphash - Port ecmp hash setting
     floodblk - Enable/disable Port flood blocking
     brate    - Set BroadCast Threshold
     mrate    - Set MultiCast Threshold
     drate    - Set Dest. Lookup Fail Threshold
     trust    - Set port as DHCP Snooping trusted or untrusted port
     dhrate   - Set DHCP packets rate limit for port
     ena      - Enable port
     dis      - Disable port
     cur      - Display current port configuration
```

The following table describes the port configuration commands that are specific to iFlow Director. For more information about configuring ports, see your *Command Reference*.

Table 11. Port Options (/cfg/port/)

Command Syntax and Usage
`aclqos` Displays the ACL/QoS Configuration menu. To view menu options, see page 86.
`tagpvid disable|enable` Disables or enables VLAN tag persistence. When disabled, the VLAN tag is removed from packets whose VLAN tag matches the port PVID. The default setting is `disabled`.
`tagipvid e|d` Enables or disables the option to insert a VLAN tag to packets that ingress a port using the PVID. When enabled, packets (untagged or single tagged) that ingress a port have an additional VLAN tag inserted, with the VLAN ID set to the port PVID. The default setting is `disabled`.

Table 11. Port Options (/cfg/port/) (continued)

Command Syntax and Usage
`tagiskip e|d` Skips the enforcement of VLAN tag memberships on the ingress port. If enabled, tagged packets are allowed to ingress the port if the VLAN tag in the packet does not belong to the port VLAN membership. If disabled, the port VLAN membership rules are enforced on the ingress port. **Note**: When this option is enabled, the `learn` option must be disabled. The default value is `disabled`.
`tageskip e|d` Skips the enforcement of VLAN tag memberships on the egress port. If enabled, tagged packets are allowed to egress the port even if the VLAN tag in the packet does not belong to the port VLAN membership. If disabled, the port VLAN membership rules are enforced on the egress port. The default value is `disabled`.
`arpmp e|d` Enables or disables the ability to forward a copy of ARP request packets on the ingress port to the management processor (MP). If enabled, when an ARP request packet is received on the ingress port, the packet will be flooded out on all ports with the same VLAN membership as the ingress port and a copy is sent to the MP. If disabled, when and an ARP request packet is received on the ingress port, the packet are flooded out on all ports with the same VLAN membership as the ingress port, but a copy is not sent to the MP. The default value is `enabled`.
`retivlan e|d` Enables or disables the ability to retain the ingress VLAN when traffic is steered from external ports to the real server application group. The default value is `disabled`.
`ecmphash sip|dip` Sets the ECMP hashing algorithm to use either the source (`sip`) or destination (`dip`) IP address. The default value is `sip`.

/cfg/port *<port alias or number>*/aclqos
Port ACL Configuration

```
[Port INT2 ACL Menu]
     add      - Add ACL or ACL group to this port
     rem      - Remove ACL or ACL group from this port
     cur      - Display current ACLs for this port
```

Table 12. Port ACL Options (/cfg/port/aclqos)

Command Syntax and Usage
`add acl|grp` *<ACL number or ACL group number>* Adds the specified ACL or ACL Group to the port.
`rem acl|grp` *<ACL number or ACL group number>* Removes the specified ACL or ACL Group from the port.
`cur` Displays current ACL port parameters.

`/cfg/l2/trunk` *<trunk group number>*
Trunk Configuration

```
[Trunk group 1 Menu]
     add      - Add port to trunk group
     rem      - Remove port from trunk group
     slbapp   - Set SLB application group
     ena      - Enable trunk group
     dis      - Disable trunk group
     del      - Delete trunk group
     cur      - Display current Trunk Group configuration
```

Trunk groups can provide super-bandwidth connections between iFlow Directors or other trunk capable devices. A *trunk* is a group of ports that act together, combining their bandwidth to create a single, larger port. Up to 18 trunk groups can be configured on the switch.

The following table describes the trunk configuration commands that are specific to iFlow Director. For more information about configuring trunks, see your *Command Reference*.

Table 13. Trunk Configuration Options (/cfg/l2/trunk)

Command Syntax and Usage
`slbapp` *<0-14>* Sets the SLB application group ID to be managed by this trunk ID reference for load balancing via trunk hashing. The default value is zero (0).

Note: The `slbapp` option is only applicable when `/cfg/slb/method` is set to `trunk`.

/cfg/l2/thash
IP Trunk Hash Configuration

```
[Trunk Hash Menu]
     l2thash  - L2 Trunk Hash Control
     l3thash  - L3 Trunk Hash Control
     ingress  - Enable/disable ingress port hash
     L4port   - Enable/disble L4 port hash
     cur      - Display current Trunk Hash configuration
```

Use the following commands to configure IP trunk hash settings for the switch. The trunk hash settings affect both static trunks and LACP trunks.

Table 14. IP Trunk Hash Options (/cfg/l2/thash)

Command Syntax and Usage
`l2thash` Displays the Layer 2 Trunk Hash Settings menu. To view menu options, see page 89.
`l3thash` Displays the Layer 3 Trunk Hash Settings menu. To view menu options, see page 89.
`ingress enable|disable` Enables or disables trunk hash computation based on the ingress port. The default setting is `disabled`.
`L4port enable|disable` Enables or disables use of Layer 4 service ports (TCP, UDP, and so on) to compute the hash value. The default setting is `disabled`.
`cur` Display current trunk hash configuration.

/cfg/l2/thash/l2thash

Layer 2 Trunk Hash Settings

```
[L2 Trunk Hash Menu]
     smac     - Enable/disable smac hash
     dmac     - Enable/disable dmac hash
     cur      - Display current trunk hash setting for L2 traffic
```

Use the following commands to configure Layer 2 trunk hash parameters for the GbESM.

Table 15. Layer 2 Hash Parameters

<table>
<tr><th>Command Syntax and Usage</th></tr>
<tr><td><code>smac enable|disable</code>
Enables or disables smac hashing.</td></tr>
<tr><td><code>smac enable|disable</code>
Enables or disables smac hashing.</td></tr>
<tr><td><code>cur</code>
Display current Layer 2 trunk hash settings.</td></tr>
</table>

/cfg/l2/thash/l3thash

Layer 3 Trunk Hash Settings

```
     useL2    - Enable/disable L2 hash for IP packet
     sip      - Enable/disable sip hash for IP packet
     dip      - Enable/disable dip hash for IP packet
     cur      - Display current trunk hash setting for L3 traffic
```

Use the following commands to configure Layer 3 trunk hash parameters for the GbESM.

Table 16. Layer 3 Hash Parameters

<table>
<tr><th>Command Syntax and Usage</th></tr>
<tr><td><code>useL2 enable|disable</code>
Enables or disables Layer 2 hashing for IP packets.</td></tr>
<tr><td><code>sip enable|disable</code>
Enables or disables SIP hashing for IP packets.</td></tr>
<tr><td><code>dip enable|disable</code>
Enables or disables DIP hashing for IP packets.</td></tr>
<tr><td><code>cur</code>
Display current Layer 3 trunk hash settings.</td></tr>
</table>

/cfg/l2/failovr
Layer 2 Failover Configuration

```
[Failover Menu]
     trigger  - Trigger Menu
     vlan     - Globally turn VLAN Monitor ON/OFF
     on       - Globally turn Failover ON
     off      - Globally turn Failover OFF
     cur      - Display current Failover configuration
```

Use this menu to configure Layer 2 Failover.

Table 17. Layer 2 Failover Options (/cfg/l2/failovr)

Command Syntax and Usage
`trigger` *<1-8>* Displays the Failover Trigger menu.
`vlan on|off` Globally turns VLAN monitor `on` or `off`. When the VLAN Monitor is `on`, the switch automatically disables only internal ports that belong to the same VLAN as ports in the Failover trigger. The default value is `off`.
`on` Globally turns Layer 2 Failover `on`.
`off` Globally turns Layer 2 Failover `off`.
`cur` Displays current Layer 2 Failover parameters.

/cfg/l3/vrrp/group
Virtual Router Group Configuration

```
[VRRP Virtual Router Group Menu]
     track    - Priority Tracking Menu
     vrid     - Set virtual router ID
     if       - Set interface number
     prio     - Set router priority
     adver    - Set advertisement interval
     garp     - Set gratuitous ARP interval
     predelay - Set preempt-delay interval
     preem    - Enable/disable preemption
     ena      - Enable virtual router
     dis      - Disable virtual router
     del      - Delete virtual router
     cur      - Display current VRRP virtual router configuration
```

The Virtual Router Group menu is used for associating all virtual routers into a single logical virtual router, which forces all virtual routers on the switch to either be master or backup as a group. For more detailed information about VRRP commands, refer to “VRRP Configuration” in the *Command Reference* for your switch.

Table 18. Virtual Router Group Options

Command Syntax and Usage
`if` *<interface number>* `[restricted]` Selects a switch IP interface. The default switch IP interface number is 1. **Restricted**: If you add this option, the switch will send advertisements on this VRRP group interface only. To clear this restriction, enter the command without the option.
`cur` Displays the current configuration information for the virtual router group.

/cfg/l3/vrrp/group/track
Virtual Router Group Priority Tracking Configuration

```
[Virtual Router Group Priority Tracking Menu]
     portlist - Priority Tracking Port List Menu
     ifs      - Enable/disable tracking other interfaces
     ports    - Enable/disable tracking VLAN switch ports
     slb-vrs  - Enable/disable tracking SLB VLAN routers
     cur      - Display current VRRP Group Tracking configuration
```

Table 19. Virtual Router Group Priority Tracking Options

Command Syntax and Usage
`portlist` *<1-2>* Displays the Priority Tracking Port List menu. To view menu options, see page 93.
`ifs disable|enable` When enabled, the priority for this virtual router will be increased for each other IP interface active on this switch. An IP interface is considered active when there is at least one active port on the same VLAN. This helps elect the virtual routers with the most available routes as the master. This command is disabled by default.
`ports disable|enable` When enabled, the priority for this virtual router will be increased for each active port on the same VLAN. A port is considered “active” if it has a link and is forwarding traffic. This helps elect the virtual routers with the most available ports as the master. This command is disabled by default.
`slb-vrs enable|disable` Enables or disables tracking of PBR next hop routers. When enabled, the priority for this virtual router is increased for each active SLB PBR next hop router reachable from this switch. This helps elect the virtual router with the most number of accessible routers as the master. The default setting is disabled.
`cur` Displays the current configuration for priority tracking for this virtual router.

/cfg/l3/vrrp/group/track/portlist

Virtual Router Group Priority Tracking Port List

```
[VRRP Track Port List 1 Menu]
     addport  - Add ports to tracking list
     remport  - Remove ports from tracking list
     del      - Delete port list
     cur      - Display current port list configuration
```

Only the link states of the physical ports are monitored when selective port tracking is configured. If LACP is enabled on a port, the port is considered operational when LACP is up on the port. Ports belonging to static or LACP trunks are treated as a single (aggregated) port, which is considered operational if at least one of the member ports in the trunk is active.

Table 20. VRRP Tracking Port List Options

Command Syntax and Usage
`addport` *<port alias or number>* Adds a physical port or ports to the current port list. These ports will be tracked when ports tracking is enabled. Place all members of a static or LACP trunk group into the same port list. You can add several ports, with each port separated by a comma (,) or a range of ports separated by a dash (-).
`remport` *<port alias or number>* Removes a physical port or ports from the current port list.
`del` Removes all ports from the port list.
`cur` Displays the current configuration for priority tracking for this port list.

/cfg/acl
Access Control List Configuration

```
[ACL Menu]
     acl      - Access Control List Item Config Menu
     acl6     - IPv6 Access Control List Item Config Menu
     group    - Access Control List Group Config Menu
     macl     - Management ACL Config Menu
     cur      - Display current ACL configuration
```

Use this menu to create Access Control Lists and ACL Groups. ACLs define matching criteria used for IP filtering and Quality of Service functions.

Table 21. ACL General Menu options (/cfg/acl)

Command Syntax and Usage
`acl` *<1-508>* Displays the Management Access Control List configuration menu. To view menu options, see page 95.
`acl6` *<1-127>* Displays the IPv6 Access Control List configuration menu. To view menu options, see page 108.
`group` *<1-508>* Displays the ACL Group configuration menu. To view menu options, see page 112.
`macl` *<1-128>* Displays the Management Access Control List configuration menu.
`cur` Displays the current ACL parameters.

/cfg/acl/acl *<ACL number>*
ACL Configuration

```
[ACL 1 Menu]
     mirror   - Mirror Options Menu
     target   - Target Options Menu
     ethernet - Ethernet Header Options Menu
     ipv4     - IP Header Options Menu
     tcpudp   - TCP/UDP Header Options Menu
     meter    - ACL Metering Configuration Menu
     re-mark  - ACL Re-mark Configuration Menu
     pktfmt   - Set to filter specific packet format types
     egrport  - Set to filter for packets egressing this port
     action   - Set filter action
     stats    - Enable/disable statistics for this acl
     reset    - Reset filtering parameters
     cur      - Display current filter configuration
```

These menus and commands allow you to define filtering criteria for each Access Control List (ACL).

Table 22. ACL Options (/cfg/acl/acl x)

Command Syntax and Usage
`mirror` Displays the ACL Mirror Options menu. To view menu options, see page 96.
`target` Displays the ACL Target Options menu. To view menu options, see page 98.
`ethernet` Displays the ACL Ethernet Header menu. To view menu options, see page 101.
`ipv4` Displays the ACL IP Header menu. To view menu options, see page 102.
`tcpudp` Displays the ACL TCP/UDP Header menu. To view menu options, see page 103.
`meter` Displays the ACL Metering menu. To view menu options, see page 104.
`re-mark` Displays the ACL Re-mark menu. To view menu options, see page 105.
`pktfmt` *<packet format>* Displays the ACL Packet Format menu. To view menu options, see page 107.
`egrport` *<port alias or number>* Configures the ACL to function on egress packets.

Table 22. ACL Options (/cfg/acl/acl x) (continued)

Command Syntax and Usage
`action permit|deny|setprio <0-7>|redirect` Configures a filter action for packets that match the ACL definitions. You can choose to permit (pass) or deny (drop) packets, or set the 802.1p priority level (0-7), or redirect traffic to a selected target defined in the Target Options menu.
`stats enable|disable` Enables or disables the statistics collection for the Access Control List.
`reset` Resets the ACL parameters to their default values.
`cur` Displays the current ACL parameters.

/cfg/acl/acl *<ACL number>*/mirror
ACL Mirror Configuration

```
[Mirror Options Menu Menu]
     dest     - Set mirror destination
     slbapp   - Set SLB application as mirror target
     trunk    - Set trunk as mirror target
     port     - Set port as mirror target
     del      - Clear mirror configuration
     cur      - Display current mirror target configuration
```

Note: This menu has been replaced by `cfg/acl/acl` *<number>*`/target/mirror`, You are advised to use that menu (see page 99) instead, as this one will be removed in a future release.

This menu allows you to define mirroring configuration for the selected ACL.

Table 23. ACL Target Mirror Options (/cfg/acl/acl x/target/mirror)

Command Syntax and Usage
`dest slbapp|trunk|port|none` Selects one mirror destination method, as follows: – If `slbapp` is selected, a nonzero value must be entered for the `slbapp` option. – If `trunk` is selected, a nonzero value must be entered for the `trunk` option. – If `port` is selected, a nonzero value must be entered for the `port` option. – If `none` is selected, mirroring is disabled on the ACL. The default value is `none`.
`slbapp` *<0-14>* Sets the SLB application group ID as the target for traffic mirroring if the `dest` option is set to `slbapp`. Otherwise, set this option to 0. The default value is 0 (zero).
`trunk` *<trunk group number>* Sets a single trunk as the target for traffic mirroring if the `dest` option is set to `trunk`. Otherwise, set this option to 0. Traffic will be mirrored to one port selected within the trunk group. The default value is 0 (zero).
`port` *<port alias or number>* Sets a single switch port as the target for traffic mirroring if the `dest` option is set to `port`. Otherwise, set this option to 0. The default value is 0 (zero).
`del` Deletes this ACL Mirroring instance.
`cur` Displays the current parameters for the ACL.

/cfg/acl/acl <ACL number>/target

Target Options Configuration

```
      mirror      - Mirror to Target Menu
      apperr      - ACL Application Error Menu
      dest        - Set target destination
      slbapp      - Set SLB application as target
      nexthop     - Set target nexthop address
      trunk       - Set egress trunk as target
      port        - Set egress port as target
      cluster     - Enable/disable slbapp cluster mode
      del         - Delete target configuration
      cur         - Display current target configuration
```

This menu allows you to define target options for an ACL.

Note: The `cluster` option is only applicable when `/cfg/slb/method` is set to `trunk`.

Table 24. Target Options (/cfg/acl/acl <ACL number>/target)

Command Syntax and Usage
`mirror` Displays the ACL Target Mirror Options menu. To view menu options, see page 99.
`apperr` Displays the ACL Application Error menu. To view menu options, see page 100.
`dest` Sets the target destination.
`slbapp` Sets the SLB application as the target.
`nexthop` Sets the target next hop address for the PBR ACL.
`trunk` Sets the egress trunk as target.
`port` Sets the egress port as target.
`cluster e|d` Defines the forwarding mode when the target destination is specified by slbapp. If enabled (“e”), traffic will be redirected to all active servers within the SLB application group. If disabled (“d”), traffic will be redirected and load balanced to one active server within the SLB application group. The default value is `d` (disabled).

Table 24. Target Options (/cfg/acl/acl <ACL number>/target) (continued)

Command Syntax and Usage
`del` Configures the ACL to function on egress packets.
`cur` Configures a filter action for packets that match the ACL definitions. You can choose to permit (pass) or deny (drop) packets, or set the 802.1p priority level (0-7), or redirect traffic to a selected target defined in the Target Options menu.

/cfg/acl/acl *<ACL number>*/target/mirror

Target Options Mirror Configuration

```
      dest     - Set mirror destination
      slbapp   - Set SLB application as mirror target
      trunk    - Set trunk as mirror target
      port     - Set port as mirror target
      del      - Delete mirror configuration
      cur      - Display current mirror target configuration
```

This menu allows you to define mirroring configuration for the selected ACL target.

Table 25. ACL Target Mirror Options (/cfg/acl/acl x/target/mirror)

Command Syntax and Usage
`dest slbapp|trunk|port|none` Selects one mirror destination method, as follows: – If `slbapp` is selected, a nonzero value must be entered for the `slbapp` option. – If `trunk` is selected, a nonzero value must be entered for the `trunk` option. – If `port` is selected, a nonzero value must be entered for the `port` option. – If `none` is selected, mirroring is disabled on the ACL. The default value is `none`.
`slbapp` *<0-14>* Sets the SLB application group ID as the target for traffic mirroring if the `dest` option is set to `slbapp`. Otherwise, set this option to 0. The default value is 0 (zero).
`trunk` *<trunk group number>* Sets a single trunk as the target for traffic mirroring if the `dest` option is set to `trunk`. Otherwise, set this option to 0. Traffic will be mirrored to one port selected within the trunk group. The default value is 0 (zero).

Table 25. ACL Target Mirror Options (/cfg/acl/acl x/target/mirror)

Command Syntax and Usage
`port` *<port alias or number>* Sets a single switch port as the target for traffic mirroring if the `dest` option is set to `port`. Otherwise, set this option to 0. The default value is 0 (zero).
`del` Deletes this Target Mirroring instance.
`cur` Displays the current parameters for the ACL.

/cfg/acl/acl *<ACL number>*/target/apperr
Target Application Error Configuration

```
[ACL Application Error Menu]
     applist  - Define SLB application dependency list
     ptrunk   - Set egress trunk for pass through
     pport    - Set egress port for pass through
     dest     - Set destination for pass through
     del      - Delete error handling configuration
     cur      - Display current mirror target configuration
```

This menu allows you to define the Application Error options for an ACL.

Table 26. ACL Target Application Error Options (/cfg/acl/acl x/target/apperr)

Command Syntax and Usage
`applist` *<0-14>* Defines the list of SLB application groups for failure monitoring. If any one application group in the list fails, the alternate redirection action specified in the `dest` option will be executed. If `applist` is set to 0, the application error checking is disabled and the parameters set in the `dest` option are ignored. The default value is 0 (zero).
`ptrunk` *<trunk group number>* Sets a single trunk as the pass through path if an application defined in the `applist` fails and the `dest` option is set to trunk. Otherwise, set this option to 0. The default value is 0 (zero).
`pport` *<port alias or number>* Sets a single switch port as the pass through path if an application defined in the `applist` fails and the `dest` option is set to port. Otherwise, set this option to 0. The default value is 0 (zero).

Table 26. ACL Target Application Error Options (/cfg/acl/acl x/target/apperr)

Command Syntax and Usage
`dest deny|permit|trunk|port` Selects an alternate redirection target destination to use if an application defined in `applist` fails. Select one of the following options: – If `trunk` is selected, a nonzero value must be entered for the `ptrunk` option. – If `port` is selected, a nonzero value must be entered for the `pport` option. – If `permit` is selected, traffic for the target destination will be forwarded via Layer 2 for the error path. – If `deny` is selected, traffic will be dropped for the error path. The default value is `deny`.
`del` Deletes this Application Error handling instance.
`cur` Displays the current parameters for the ACL.

`/cfg/acl/acl` *<ACL number>*`/ethernet`
Ethernet Filtering Configuration

```
      smac      - Set to filter on source MAC
      dmac      - Set to filter on destination MAC
      vlan      - Set to filter on VLAN ID
      etype     - Set to filter on ethernet type
      pri       - Set to filter on priority
      reset     - Reset all fields
      cur       - Display current parameters
```

iFlow Director Layer 3 allows selective traffic to be redirected to the real server application group using ACL with Ethernet, IPv, and TCP/UDP filters.

This menu allows you to define Ethernet matching criteria for an ACL.

Table 27. ACL Ethernet Filtering Options (/cfg/acl/acl x/ethernet)

Command Syntax and Usage
`smac` *<MAC address (such as 00:60:cf:40:56:00)> <mask (FF:FF:FF:FF:FF:FF)>* Defines the source MAC address for this ACL.
`dmac` *<MAC address (such as 00:60:cf:40:56:00)> <mask (FF:FF:FF:FF:FF:FF)>* Defines the destination MAC address for this ACL.
`vlan` *<1-4095> <VLAN mask (0xfff)>* Defines a VLAN number and mask for this ACL.
`etype ARP|IP|IPv6|MPLS|RARP|any|none|`*<hex value>* Defines the Ethernet type for this ACL.

Table 27. ACL Ethernet Filtering Options (/cfg/acl/acl x/ethernet)

Command Syntax and Usage
`pri` *<0-7>* Defines the Ethernet priority value for the ACL.
`reset` Resets Ethernet parameters for the ACL to their default values.
`cur` Displays the current Ethernet parameters for the ACL.

/cfg/acl/acl ⟨*ACL number*⟩/ipv4
IP version 4 Filtering Configuration

```
[Filtering IPv4 Menu]
     sip      - Set to filter on source IP address
     dip      - Set to filter on destination IP address
     proto    - Set to filter on prototype
     tos      - Set to filter on TOS
     reset    - Reset all fields
     cur      - Display current parameters
```

This menu allows you to define IPv4 matching criteria for an ACL.

Table 28. ACL IPv4 Filtering Options (/cfg/acl/acl x/ipv4)

Command Syntax and Usage
`sip` *<IP address> <mask (such as 255.255.255.0)>* Defines a source IP address for the ACL. If defined, traffic with this source IP address will match this ACL. Specify an IP address in dotted decimal notation.
`dip` *<IP address> <mask (such as 255.255.255.0)>* Defines a destination IP address for the ACL. If defined, traffic with this destination IP address will match this ACL.
`proto` *<0-255>* Defines an IP protocol for the ACL. If defined, traffic from the specified protocol matches this filter. Specify the protocol number. The following are some of the well-known protocols: **Number** **Name** 1 icmp 2 igmp 6 tcp 17 udp 89 ospf 112 vrrp
`tos` *<0-255>* Defines a Type of Service value for the ACL. For more information on ToS, refer to RFCs 1340 and 1349.

Table 28. ACL IPv4 Filtering Options (/cfg/acl/acl x/ipv4)

Command Syntax and Usage
`reset` Resets the IPv4 parameters for the ACL to their default values.
`cur` Displays the current IPV4 parameters.

/cfg/acl/acl <*ACL number*>/tcpudp
TCP/UDP Filtering Configuration

```
[Filtering TCP/UDP Menu]
     sport     - Set to filter on TCP/UDP source port
     dport     - Set to filter on TCP/UDP destination port
     flags     - Set to filter TCP/UDP flags
     reset     - Reset all fields
     cur       - Display current parameters
```

This menu allows you to define TCP/UDP matching criteria for an ACL.

Table 29. ACL TCP/UDP Filtering Options (/cfg/acl/acl x/tcpudp)

Command Syntax and Usage

`sport` *<source port (1-65535)> <port mask (0x1-0xFFFF)>*
Defines a source port for the ACL. If defined, traffic with the specified TCP or UDP source port will match this ACL. The default port mask is “0xFFFF”. Specify the port number. The following are some of the well-known ports:
Number Name
`20 ftp-data`
`21 ftp`
`22 ssh`
`23 telnet`
`25 smtp`
`37 time`
`42 name`
`43 whois`
`53 domain`
`69 tftp`
`70 gopher`
`79 finger`
`80 http`

`dport` *<destination port (1-65535)> <port mask (0x1-0xFFFF)>* Defines a destination port for the ACL. If defined, traffic with the specified TCP or UDP destination port will match this ACL. The default port mask is “0xFFFF”. Specify the port number as with `sport`.
`flags` *<value (0x0-0x3f)>* Defines a TCP/UDP flag for the ACL.

Table 29. ACL TCP/UDP Filtering Options (/cfg/acl/acl x/tcpudp)

Command Syntax and Usage
`reset` Resets the TCP/UDP parameters for the ACL to their default values.
`cur` Displays the current TCP/UDP Filtering parameters.

/cfg/acl/acl ⟨*ACL number*⟩/meter
ACL Metering Configuration

```
[Metering Menu]
     cir      - Set committed rate in KiloBits/s
     mbsize   - Set maximum burst size in KiloBits
     enable   - Enable/disable port metering
     dpass    - Set to Drop or Pass out of profile traffic
     reset    - Reset meter parameters
     log      - Enable/disable Out of Profile Notification
     cur      - Display current settings
```

This menu defines the metering profile for the selected ACL.

Table 30. ACL Metering Options (/cfg/acl/acl x/meter)

Command Syntax and Usage
`cir` *<64-10000000>* Configures the committed rate, in Kilobits per second. The committed rate must be a multiple of 64.
`mbsize` *<32-4096>* Configures the maximum burst size, in Kilobits. Enter one of the following values for `mbsize`: 32, 64, 128, 256, 512, 1024, 2048, 4096.
`enable e|d` Enables or disables metering on the ACL.
`dpass drop|pass` Configures the ACL Meter to either drop or pass out-of-profile traffic.
`reset` Reset ACL Metering parameters to their default values.
`log e|d` Enables or disables logging out-of-profile notifications.
`cur` Displays current ACL Metering parameters.

/cfg/acl/acl *<ACL number>*/re-mark

ACL Re-Mark Configuration

```
[Re-mark Menu]
     inprof   - In Profile Menu
     outprof  - Out Profile Menu
     reset    - Reset re-mark settings
     cur      - Display current settings
```

You can choose to re-mark IP header data for the selected ACL. You can configure different re-mark values, based on whether packets fall within or out of the ACL Metering profile.

Table 31. ACL Re-mark Options (/cfg/acl/acl x/re-mark)

Command Syntax and Usage
`inprof` Displays the Re-mark In-Profile menu. To view menu options, see page 105.
`outprof` Displays the Re-mark Out-of-Profile menu. To view menu options, see page 107.
`reset` Reset ACL Re-mark parameters to their default values.
`cur` Displays current Re-mark parameters.

/cfg/acl/acl *<ACL number>*/re-mark/inprof

Re-Marking In-Profile Configuration

```
[Re-marking - In Profile Menu]
     up1p     - Set Update User Priority Menu
     updscp   - Set the update DSCP
     reset    - Reset update DSCP settings
     cur      - Display current settings
```

Table 32. ACL Re-Mark In-Profile Options (/cfg/acl/acl x/re-mark/inprof)

Command Syntax and Usage
`up1p` Displays the Re-Mark In-Profile Update User Priority menu.
`updscp` *<0-63>* Sets the DiffServ Code Point (DSCP) of In-Profile packets to the selected value.

Table 32. ACL Re-Mark In-Profile Options (/cfg/acl/acl x/re-mark/inprof) (continued)

Command Syntax and Usage
`reset` Resets the update DSCP parameters to their default values.
`cur` Displays current Re-Mark In-Profile parameters.

/cfg/acl/acl *<ACL number>*/re-mark/inprof/up1p

Update User Priority Configuration

```
[Update User Priority Menu]
     value    - Set the update user priority
     utosp    - Enable/Disable use of TOS precedence
     reset    - Reset in profile up1p settings
     cur      - Display current settings
```

Table 33. ACL Re-Mark User Priority (/cfg/acl/acl x/re-mark/inprof/up1p)

Command Syntax and Usage
`value` *<0-7>* Defines 802.1p value. The value is the priority bits information in the packet structure.
`utosp enable|disable` Enable or disable mapping of TOS (Type of Service) priority to 802.1p priority for In-Profile packets. When enabled, the TOS value is used to set the 802.1p value.
`reset` Resets UP1P settings to their default values.
`cur` Displays current Re-Mark In-Profile User Priority parameters.

/cfg/acl/acl <*ACL number*>/re-mark/outprof

Re-Marking Out-of-Profile Configuration

```
[Re-marking - Out Of Profile Menu]
     updscp   - Set the update DSCP
     reset    - reset update DSCP setting
     cur      - Display current settings
```

Table 34. ACL Re-Mark Out-of-Profile Options (/cfg/acl/acl x/re-mark/outprof)

Command Syntax and Usage
`updscp` *<0-63>* Sets the DiffServ Code Point (DSCP) of Out-of-Profile packets to the selected value. The switch sets the DSCP value on Out-of-Profile packets.
`reset` Resets the update DSCP parameters for Out-of-Profile packets to their default values.
`cur` Displays current Re-Mark Out-of-Profile parameters.

/cfg/acl/acl <*ACL number*>/pktfmt

Packet Format Filtering Configuration

```
[Filtering Packet Format Menu]
     ethfmt   - Set to filter on ethernet format
     tagfmt   - Set to filter on ethernet tagging format
     ipfmt    - Set to filter on IP format
     reset    - Reset all fields
     cur      - Display current parameters
```

This menu allows you to define Packet Format matching criteria for an ACL.

Table 35. ACL Packet Format Filtering Options (/cfg/acl/acl x/pktfmt)

Command Syntax and Usage
`ethfmt eth2|SNAP|LLC|none` Defines the Ethernet format for the ACL.
`tagfmt disabled|any|none|tagged` Defines the tagging format for the ACL.
`ipfmt none|v4|v6` Defines the IP format for the ACL.

Table 35. ACL Packet Format Filtering Options (/cfg/acl/acl x/pktfmt) (continued)

Command Syntax and Usage
`reset` Resets Packet Format parameters for the ACL to their default values.
`cur` Displays the current Packet Format parameters for the ACL.

/cfg/acl/acl6 *<ACL number>*
ACL IPv6 Configuration

```
[ACL6 2 Menu]
     ipv6     - IPv6 Header Options Menu
     tcpudp   - TCP/UDP Header Options Menu
     re-mark  - ACL Re-mark Configuration Menu
     egrport  - Set to filter for packets egressing this port
     action   - Set filter action
     stats    - Enable/disable statistics
     reset    - Reset filtering parameters
     cur      - Display current filter configuration
```

These menus allow you to define filtering criteria for each IPv6 Access Control List (ACL).

Table 36. IPv6 ACL Options

Command Syntax and Usage
`ipv6` Displays the ACL IP Header menu. To view menu options, see page 109.
`tcpudp` Displays the ACL TCP/UDP Header menu. To view menu options, see page 110.
`re-mark` Displays the ACL Re-Mark menu. To view menu options, see page 112.
`egrport` *<port alias or number>* Configures the ACL to function on egress packets.
`action permit|deny|setprio` *<0-7>* Configures a filter action for packets that match the ACL definitions. You can choose to permit (pass) or deny (drop) packets or set the 802.1p priority level (0-7).
`stats e|d` Enables or disables the statistics collection for the Access Control List.

Table 36. IPv6 ACL Options (continued)

Command Syntax and Usage
`reset` Resets the ACL parameters to their default values.
`cur` Displays the current ACL parameters.

/cfg/acl/acl6 *<ACL number>*/ipv6
IP version 6 Filtering Configuration

```
[Filtering IPv6 Menu]
     sip      - Set to filter on source IPv6 address
     dip      - Set to filter on destination IPv6 address
     nexthd   - Set to filter on IPv6 next header
     flabel   - Set to filter on IPv6 flow label
     tclass   - Set to filter on IPv6 traffic class
     reset    - Reset all fields
     cur      - Display current parameters
```

This menu allows you to define IPv6 matching criteria for an ACL.

Table 37. IP version 6 Filtering Options

Command Syntax and Usage
`sip` *<IPv6 address> <prefix length>* Defines a source IPv6 address for the ACL. If defined, traffic with this source IP address will match this ACL.
`dip` *<IPv6 address> <prefix length>* Defines a destination IPv6 address for the ACL. If defined, traffic with this destination IP address will match this ACL.
`nexthd` *<0-255>* Defines the next header value for the ACL. If defined, traffic with this next header value will match this ACL.
`flabel` *<0-0xFFFFF>* Defines the flow label for the ACL. If defined, traffic with this flow label will match this ACL.
`tclass` *<0-255>* Defines the traffic class for the ACL. If defined, traffic with this traffic class will match this ACL.

Table 37. IP version 6 Filtering Options (continued)

Command Syntax and Usage
`reset` Resets the IPv6 parameters for the ACL to their default values.
`cur` Displays the current IPv6 parameters.

`/cfg/acl/acl6` *<ACL number>*`/tcpudp`
IPv6 TCP/UDP Filtering Configuration

```
[Filtering TCP/UDP Menu]
     sport    - Set to filter on TCP/UDP source port
     dport    - Set to filter on TCP/UDP destination port
     flags    - Set to filter TCP/UDP flags
     reset    - Reset all fields
     cur      - Display current parameters
```

This menu allows you to define TCP/UDP matching criteria for an ACL.

Table 38. IPv6 ACL TCP/UDP Filtering Options

Command Syntax and Usage
`sport` *<source port (1-65535)> <mask (0xFFFF)>* Defines a source port for the ACL. If defined, traffic with the specified TCP or UDP source port will match this ACL. Specify the port number. The following are some of the well-known ports: **Number** **Name** 20 ftp-data 21 ftp 22 ssh 23 telnet 25 smtp 37 time 42 name 43 whois 53 domain 69 tftp 70 gopher 79 finger 80 http
`dport` *<destination port (1-65535)> <mask (0xFFFF)>* Defines a destination port for the ACL. If defined, traffic with the specified TCP or UDP destination port will match this ACL. Specify the port number as with `sport`.
`flags` *<value (0x0-0x3f)> <mask (0x0-0x3f)>* Defines a TCP/UDP flag for the ACL.

Table 38. IPv6 ACL TCP/UDP Filtering Options (continued)

Command Syntax and Usage
`reset` Resets the TCP/UDP parameters for the ACL to their default values.
`cur` Displays the current TCP/UDP Filtering parameters.

/cfg/acl/acl6 *<ACL number>*/re-mark

IPv6 Re-Mark Configuration

```
[Re-mark Menu]
     inprof   - In Profile Menu
     reset    - Reset re-mark settings
     cur      - Display current settings
```

You can choose to re-mark IP header data for the selected ACL. You can configure different re-mark values, based on whether packets fall within the ACL metering profile, or out of the ACL metering profile.

Table 39. IPv6 ACL Re-Mark Options

Command Syntax and Usage
`inprof` Displays the Re-Mark In-Profile menu. To view menu options, see page 112.
`reset` Reset ACL re-mark parameters to their default values.
`cur` Displays current re-mark parameters.

/cfg/acl/acl6 *<ACL number>*/re-mark/inprof

IPv6 Re-Marking In-Profile Configuration

```
[Re-marking - In Profile Menu]
     updscp   - Set the update DSCP
     reset    - Reset update DSCP settings
     cur      - Display current settings
```

Table 40. IPv6 ACL Re-Mark In-Profile Options

Command Syntax and Usage
`updscp` *<0-63>* Re-marks the DiffServ Code Point (DSCP) of in-profile packets to the selected value.
`reset` Resets the update DSCP parameters to their default values.
`cur` Displays current re-mark parameters for in-profile packets.

/cfg/acl/group *<ACL Group number>*

ACL Group Configuration

```
[ACL Group 1 Menu]
     add      - Add ACL to group
     rem      - Remove ACL from group
     add6     - Add IPv6 ACL to ACL group
     rem6     - Remove IPv6 ACL from ACL group
     cur      - Display current ACL items in group
```

This menu allows you to configure one or more ACLs to be added to an ACL Group. Once you create an ACL Group, you can assign the ACL Group to one or more ports.

Table 41. ACL Group Options (/cfg/acl/group x)

Command Syntax and Usage
`add acl` *<1-508>* Adds the selected ACL to the ACL Group.
`rem acl` *<1-508>* Removes the selected ACL from the ACL Group.
`add acl6` *<1-127>* Adds the selected IPv6 ACL to the ACL Group.

Table 41. ACL Group Options (/cfg/acl/group x)

Command Syntax and Usage
`rem acl6` *<1-127>* Removes the selected IPv6 ACL from the ACL Group.
`cur` Displays the current ACL group parameters.

/oper
Operations Menu

```
[Operations Menu]
     port     - Operational Port Menu
     slb      - Operational Server Load Balancing Menu
     vrrp     - Operational Virtual Router Redundancy Menu
     ip       - Operational IP Menu
     prm      - Protected Mode Menu
     sys      - Operational System Menu
     passwd   - Change current user password
     clrlog   - Clear syslog messages
     tnetsshc - Close all telnet/SSH connections
     conlog   - Enable/Disable Session Console Logging
     cfgtrk   - Track last config change made
     ntpreq   - Send NTP request
     swkey    - Sofware License Menu
```

Use the Operations Menu to alter operational characteristics without affecting switch configuration.

The following table briefly describes only the operations commands that are specific to iFlow Director. For more information about operations commands, see your *Command Reference*.

Table 42. Operations Menu (/oper)

Command Syntax and Usage
`slb` Displays the Operational Server Load Balancing menu. To view menu options, see page 114.
`cfgtrk` Track the last configuration change made.
`swkey` Displays the Software License Key menu. To view menu options, see page 118.

/oper/slb
Operations-Level SLB Options

```
[Server Load Balancing Operations Menu]
     app      - Application Menu
     real     - Real Server Menu
     cur      - Current Server Load Balancing operational state
```

The operations-level Server Load Balancing options are used for temporarily disabling or enabling real servers.

Note: The `app` option is only applicable when `/cfg/slb/method` is set to `trunk`.

Table 43. Server Load Balancing Operations Options (/oper/slb)

Command Syntax and Usage
`app` *<application number (1-14)>* Displays the Application Operations menu. To view menu options, see page 115.
`real` *<real server number (1-84)>* Displays the Real Server Operations menu. To view menu options, see page 117.
`cur` Displays the current SLB operating parameters.

`/oper/slb/app` *<application number>*
Operations-Level SLB Application Options

```
[Application 1 Menu]
     bucket   - Application Bucket Menu
     update   - Update Application Distribution Matrix
     reset    - Reset Application Distribution Matrix changes
     clear    - Clear Application Distribution Matrix
     cur      - Show current Application state
```

Use this menu to assign switch ports to trunk hash table entries (buckets) that support the application. This feature allows you to fine-tune the hashing to manually control traffic distribution across application ports.

Note: The options in this menu are only applicable when `/cfg/slb/method` is set to `trunk`.

Table 44. Application Operations (/oper/slb/app)

Command Syntax and Usage
`bucket` *<1-64>* Displays the application bucket menu, which allows you to add a real server port to the selected hash bucket. **Note:** Only real server ports that are members of the application can be selected.
`update` Executes changes to the operating hash matrix during runtime operation.
`reset` Resets the hash matrix to its current operating parameters. Use this command to remove changes that have not been updated and revert to the current running matrix.

Table 44. Application Operations (/oper/slb/app) (continued)

Command Syntax and Usage
`clear` Clears the current operating hash matrix. This change is not executed until you issue the `update` command.
`cur` Displays the current operational parameters for the application.

`/oper/slb/app` *<application number>*`/bucket` *<1-64>*
Operations-Level SLB Application Options

```
[Application Bucket 2 Menu]
     port     - Set port in Application Bucket
     cur      - Show current Application Bucket state
```

Use this menu to add a real server port to the selected hash bucket.

Note: The options in this menu are only applicable when `/cfg/slb/method` is set to `trunk`.

Table 45. Application Hash Bucket Options (/oper/slb/app/bucket)

Command Syntax and Usage
`port` *<port alias or number>* Adds the port to the selected hash bucket.
`cur` Displays the current bucket parameters.

/oper/slb/real *<server number>*

Operations-Level SLB Real Server Options

```
[Server Load Balancing Operations Menu]
     restore  - Restore Real Server Menu
     ena      - Enable real server
     dis      - Disable real server
     cur      - Current real server operational state
```

Note: The `restore` option is only applicable when `/cfg/slb/method` is set to `trunk`.

Table 46. Real Server options (/oper/slb/real)

Command Syntax and Usage
`restore` Displays the real server Restore menu.
`ena` Enables the real server. The real server will be returned to its configured operation mode when the switch is reset.
`dis` Disables the real server. The real server will be returned to its configured operation mode when the switch is reset.
`cur` Displays the current real server operational state.

/oper/slb/real *<server number>*/restore

Operations-Level SLB Restore Options

```
[Server Load Balancing Operations Menu]
     replace  - Replace active backup server with primary
     recover  - Recover and remap with primary server
     cur      - Current Server Restoration operational state
```

Use this menu to selectively restore service to a primary server forced into a standby state during runtime operations.

Note: The options in this menu are only applicable when `/cfg/slb/method` is set to `trunk`.

Table 47. Server Restoration Options (/oper/slb/real/restore)

Command Syntax and Usage
`replace` Restores a standby, primary server into active service under the following circumstances. When a backup server is active, a “healthy” primary server maybe forced into a standby mode due to the `preempt dis` (disable pre-emption) option. As a result, the `replace` command can be used to operationally replace the active backup server with the restored primary server without re-mapping the application group. The `replace` command applies only to standby, primary servers with a configured backup and disregards the rules for `preempt dis` on the selected/associated blade(s) from the standby portmap, but still honor the rules for `remap dis` (disable re-mapping) on the associated application group.
`recover` Restores a real server into active service under the following circumstances. When a failed server is restored, it may be added to the active portmap but not in the hashmap due to the `remap dis` (disable re-mapping) option. As a result, the `recover` option can be used to operationally restore the server to active service by injecting and re-mapping the server back into the application group hash map. The `recover` option applies to real servers with or without a configured backup and disregards the rules for the `preempt dis` (disable pre-emption) on the selected/associated blade(s) from the standby portmap and `remap dis` (disable re-mapping) on the associated application group.
`cur` Displays the current server restoration operational state.

/oper/swkey
Software License Key Options

```
[Software License Menu]
     key      - License Key Menu
```

Table 48. Software License Key options (/oper/swkey)

Command Syntax and Usage
`key` Displays the License Key menu.

/oper/swkey/key

License Key Options

```
[License Key Menu]
     enakey   - Enable License Key/Software Feature
     rmkey    - Remove License Key/Software Feature
```

Table 49. License Key Options (/oper/swkey/key)

Command Syntax and Usage
`enakey` *<feature name>* Allows you to unlock iFlow Director. You are prompted to enter the feature name (`ibmiflow`) and the license key code.
`rmkey` *<feature name>* Removes the license key for iFlow Director.

Chapter 5. ISCLI Reference

The following sections include information about iFlow Director ISCLI commands. For more detailed information about switch commands, refer to the *ISCLI Reference* for your switch.

The following topics are discussed in this chapter:

Information Commands

This section describes the information commands that are specific to iFlow Director. For more information about information commands, see your *ISCLI Reference*.

The general information commands are briefly described in the following table.

Table 50. Information Options

Command Syntax and Usage
`show access-control` Displays the current configuration profile for each Access Control List (ACL) and ACL Group. **Command mode:** All
`show software-key` Displays license information for iFlow Director. **Command mode:** All

SLB Information

The following table describes the Server Load Balancing (SLB) information commands.

Table 51. SLB Information Options

Command Syntax and Usage
`show slb real-server` *<real server number (1-84)>* `information` Displays Real server number, real IP address, MAC address, physical switch port, layer where health check is performed, and health check result. **Command mode:** All
`show slb server-group` *<group number (1-14)>* `information` Displays information for the selected real server group. **Command mode:** All
`show slb application` *<application number (1-14)>* `information` Displays information for the selected application. **Command mode:** All
`show slb bind` *<IPv4 address> <mask>* Displays real server port selected by hashing on an IPv4 address. **Command mode:** All
`show slb bind6` *<IPv6 address>* Displays real server port selected by hashing on an IPv6 address. **Command mode:** All

Table 51. SLB Information Options (continued)

Command Syntax and Usage
`show slb ecmp-table` Displays the ECMP hashing table. **Command mode:** All
`show slb information` Displays all SLB information for the switch. **Command mode:** All

Note: The `bind`, `bind6`, and `ecmp-table` options are only applicable when `slb method` is set to `ecmp`.

SLB Statistics Commands

The following section describes the statistics commands that are specific to iFlow Director. For more information about statistics commands, see your *ISCLI Reference*.

The following table describes the Server Load Balancing (SLB) statistics commands.

Table 52. SLB Statistics Options

Command Syntax and Usage
`show slb application {<application number (1-14)>|all} counters` Displays statistics for the selected application. **Command mode:** All
`clear slb application {<application number (1-14)>|all} counters` Clears SLB application statistics. **Command mode:** Privileged EXEC

SLB Configuration

The following sections describe the configuration commands that are specific to iFlow Director. For more information about configuration commands, see your *ISCLI Reference*.

Table 53. Server Load Balancing Configuration Options

Command Syntax and Usage
`slb enable` Globally turns on Server Load Balancing and Application Redirection. **Command mode:** Global configuration
`no slb enable` Globally disables Server Load Balancing. All configuration information will remain in place (if `applied` or `saved`), but the processes will no longer be active in the switch **Command mode:** Global configuration
`slb method ecmp|trunk` Configures the SLB hashing method. The default setting is `ecmp`. **Command mode:** Global configuration
`[no] slb mod` Enables or disables manual hash table distribution. When enabled, you must manually set the hash port matrix. For more information, see "Operations-Level SLB Application Options" on page 158. The default setting is `disabled`. **Note**: When this option is enabled, health checks by the switch are disabled. **Command mode:** Global configuration **Note:** The `slb mod` option is only applicable when `slb method` is set to `trunk`.
`[no] slb synchronization` Enables or disables hash synchronization between two peer switches. **Command mode:** Global configuration **Note:** The `slb synchronization` option is only applicable when `slb method` is set to `ecmp`.
`slb hcwtint` *<0-60>* Sets the health check delay time (in seconds) on failover. The default value is `4`. **Command mode:** Global configuration

Table 53. Server Load Balancing Configuration Options (continued)

Command Syntax and Usage
`slb initarp` *<1-60>* Sets the ARP delay time (in seconds) on failover. The default value is `2`.
`show slb` Displays the current Server Load Balancing configuration. **Command mode:** All

SLB Application Management Configuration

Table 54. SLB Application Management Options

Command Syntax and Usage
`[no] slb application` *<1-14>* `name` *<string, maximum 31 characters>* Defines an alias for each application. **Command mode:** Global configuration
`[no] slb application` *<1-14>* `server-group` *<group number (1-14)>* Sets the SLB group ID to be managed by this application. **Command mode:** Global configuration
`slb application` *<1-14>* `mode {inline|hairpin}` Sets the application operation mode to the following: – inline – application will receive traffic from a local port and forward traffic inline to a remote port. – hairpin – application will receive and forward traffic on its local port. **Command mode:** Global configuration
`[no] slb application` *<1-14>* `aggregation` Enables or disables application aggregation. When enabled, this application is designated as a sub-application that will be aggregated with other sub-applications to form a super-application. The default setting is `disabled`. **Command mode:** Global configuration
`[no] slb application` *<1-14>* `expand` Enables or disables hash expansion for applications that have 8 or fewer real servers. When enabled, the number of trunks required to support the application is increased to 8 and the number of hash table entries is increased to 64. The default setting is `disabled`. **Command mode:** Global configuration

Table 54. SLB Application Management Options (continued)

Command Syntax and Usage
`no slb application` *<1-14>* Deletes this Application Management instance. **Command mode:** Global configuration
`show slb application` *<1-14>* Displays the current configuration parameters for this Application Management instance. **Command mode:** All

Note: The `slb aggregation` and `slb expand` options are only applicable when `slb method` is set to `trunk`.

Application Error Handling Configuration

Table 55. Error Handling Options

Command Syntax and Usage
`slb application` *<1-14>* **error-handling transition-delay** *<0-60>* Sets the application transition delay (in seconds) to define the period of instability while health checking real servers in the application group during the following conditions: – When creating a new SLB application for the first time. – When booting up with an existing SLB application configuration. – When re-populating the application group with an empty port membership. – When changing the real server membership in an SLB application during runtime. – When applying SLB-related configuration changes. The default value is 20 seconds. **Command mode:** Global configuration
`[no] slb application` *<1-14>* **error-handling preemption** Enables or disables pre-emption, which provides the ability to fail back to the primary application servers in the event that service is restored while the corresponding backup server is active. The default setting is `enabled`. **Command mode:** Global configuration

Table 55. Error Handling Options (continued)

Command Syntax and Usage
`[no] slb application` *<1-14>* **error-handling remap** Enables or disables traffic redistribution, or re-mapping of the server group load distribution in the event of server failures within an application group. The default setting is `disabled`. **Command mode:** Global configuration
`show slb application` *<1-14>* **error-handling** Displays the current error handling configuration for this application. **Command mode:** All

Note: The `slb application error-handling` options are only applicable when `slb method` is set to `trunk`.

Real Server SLB Configuration

This menu is used for configuring information about real servers that participate in a server pool for Server Load Balancing or Application Redirection. The required parameters are:

- Real server port number
- Real server enabled (disabled by default)

Table 56. Real Server Configuration Options

Command Syntax and Usage
`[no] slb real-server` *<1-84>* `name` *<1-31 characters>* Defines an alias for each real server. This will enable the network administrator to quickly identify the server by a natural language keyword value. **Command mode:** Global configuration
`[no] slb real-server` *<1-84>* `port` *<port alias or number>* Defines the port to which this real server connects. The default setting is `INT1`. **Command mode:** Global configuration
`slb real-server` *<1-84>* `ip-address` *<real server IP address>* Sets the IP address of the real server, in dotted decimal format to define the IP address for Layer 3 health checks on the specified port number to determine if the server is up. The administrator will be warned if the server does not respond. **Note**: If you define a Real Server IP address, you must also configure the Real Server port. **Command mode:** Global configuration

Table 56. Real Server Configuration Options (continued)

Command Syntax and Usage
`slb real-server` *<1-84>* `interval` *<1-60 seconds, or 400 for milliseconds>* Sets the interval between real server health verification attempts. Determining the health of each real server is a necessary function for SLB. The `inter` option lets you choose the time between health checks. The range is from 1 to 60 seconds, or 400 milliseconds. The default interval is 2 seconds. **Command mode:** Global configuration
`slb real-server` *<1-84>* `retry` *<number of consecutive health checks (2-63)>* Sets the number of failed health check attempts required before declaring this real server inoperative. The default value is 4 attempts. **Command mode:** Global configuration
`slb real-server` *<1-84>* `restore` *<number of consecutive health checks (2-63)>* Sets the number of successful health check attempts required before declaring a server operational. The default value is 8 attempts. **Command mode:** Global configuration
`[no] slb real-server` *<1-84>* `backup` *<real server number (1-84)>* Sets the real server used as the backup server for this real server. To prevent loss of service if a particular real server fails, use this option to assign a backup real server number. Then, if the primary real server becomes inoperative, the switch will activate the backup real server. The same backup server may be assigned to more than one real server within the same application group. **Command mode:** Global configuration
`slb real-server` *<1-84>* `enable` You *must* perform this command to enable this real server for SLB service. **Command mode:** Global configuration
`no slb real-server` *<1-84>* `enable` Disables this real server from SLB service. This option *does not* perform a graceful server shutdown. **Command mode:** Global configuration
`no slb real-server` *<1-84>* Deletes this real server from the SLB configuration. This removes the real server from operation within its real server groups. Use this command with caution, as it will delete any configuration options that have been set for this real server. This option *does not* perform a graceful server shutdown. **Command mode:** Global configuration
`show slb real-server` *<1-84>* Displays the current configuration information for this real server. **Command mode:** All

Real Server Group SLB Configuration

This menu is used for combining real servers into real server groups. Each real server group should consist of all the real servers which provide a specific service for load balancing. Each group must consist of at least one real server. Real server groups are used both for Server Load Balancing and Application Redirection.

Table 57. Real Server Group Configuration Options

<table>
<tr><th>Command Syntax and Usage</th></tr>
<tr><td>slb server-group <i>&lt;1-14&gt;</i> health {link|ping|ips|tcp|http|httphead|smtp|ssl|udpdns|arp|sip-register|script <i>&lt;1-16&gt;</i>}
Sets the type of health checking performed, as follows:
- link
<b>Note:</b> The link option is only applicable when slb method is set to trunk.
- ping
– ips (Intrusion Protection System, a Layer 2 application-specific health check)
— tcp
— http
— httphead
— smtp
— ssl
— udpdns
— arp
— sip-register
- script
The default setting is link.
<b>Note:</b> If the switch is configured with health check smtp, sip-register, or udpdns, you must first change the health-check type to link before changing it to another of the types listed.
<b>Command mode:</b> Global configuration</td></tr>
<tr><td>[no] slb server-group <i>&lt;1-14&gt;</i> name <i>&lt;1-31 characters&gt;</i>
Defines an alias for each Real Server Group. This will enable the network administrator to quickly identify the server group by a natural language keyword value.
<b>Command mode:</b> Global configuration</td></tr>
</table>

Table 57. Real Server Group Configuration Options (continued)

Command Syntax and Usage
`slb server-group` *<1-14>* `real-threshold` *<1-84, 0 for disabled>* Specifies the real server failure threshold for triggering an application bypass action, as defined by the following command: `/cfg/acl/acl #/target/apperr` For more details, see page 146. If the number of operational servers is equal to or less than the threshold, the bypass action is initiated if it is defined. When the number of operational ports rises above the threshold, the application reverts to normal behavior. A SYSLOG message is generated when the # of operational real server ports used within the SLB application group hash matrix have failed down to a non-zero threshold value. A value of 0 (zero) is used to disable the SYSLOG alert. The default threshold value is 0. **Command mode:** Global configuration
`slb server-group` *<1-14>* `service-port` *<0-65534>* Configures the service port to send health-check data on to the server, if necessary. This option is used only for health-check types that require a service port, as follows: – `tcp` - `http` - `httphead` - `smtp` - `ssl` - `udpdns` - `sip-register` The default setting is 0 (zero). **Command mode:** Global configuration
`[no] slb server-group` *<1-14>* `content` *<text string>* Defines supporting information for the selected health-check type, as follows: – `udpdns`: Define a URL, such as `www.ibm.com` (1-127 characters) – `http`: Define a filename or path, such as `index.html` (1-127 characters) – `sip-register`: Define a username, such as `Smith` (1-32 characters) – `smtp`: Define an email address, such as `joe@ibm.com` (1-64 characters) The default value is null (empty). **Command mode:** Global configuration
`slb server-group` *<1-14>* `real-server` *<real server number (1-84)>* Adds a real server to this real server group. **Command mode:** Global configuration
`no slb server-group` *<1-14>* `real-server` *<real server number (1-84)>* Removes a real server from this real server group. **Command mode:** Global configuration

Table 57. Real Server Group Configuration Options (continued)

Command Syntax and Usage
`no slb server-group` *<1-14>* Deletes this real server group from the SLB configuration. **Command mode:** Global configuration
`show slb server-group` *<1-14>* Displays the current configuration parameters for this real server group. **Command mode:** All

Scriptable Health Check Configuration

The maximum size of a script is defined by the total number of characters allowed in the script
(6144 characters).

Table 58. Scriptable Health-Check Options

Command Syntax and Usage

`slb script` *<1-16>* `open` *<0-65534>*
Initiates a TCP three-way handshake to the selected service port. Some well known port numbers/names are listed here:
`20 = ftp-data`
`21 = ftp`
`23 = telnet`
`25 = smtp`
`53 = dns`
`69 = tftp`
`80 = http`
`443 = https`
`1812 = radius`
Command mode: Global configuration

`slb script` *<1-16>* `send` *<1-502 characters>* Sends a text string to the server port. **Command mode:** Global configuration
`slb script` *<1-16>* `expect` *<1-502 characters>* Waits for a string from the server and checks if the argument matches the first part of the string received. This command fails if the expected string is not found in the server response. If a match is found, the next script command is executed. If a match is not found, the system waits for the next packet to arrive. If no matching string arrives by the next health check interval, the script health check fails and the current connection is terminated. **Command mode:** Global configuration
`slb script` *<1-16>* `close` Closes an open TCP connection. **Command mode:** Global configuration

Table 58. Scriptable Health-Check Options

Command Syntax and Usage
`no slb script` *<1-16>* `last-command` Removes the last line from the script. **Command mode:** Global configuration
`no slb script` *<1-16>* Deletes the script. **Command mode:** Global configuration
`show slb script` *<1-16>* Displays the current health-check script parameters. **Command mode:** All

Policy-Based Routing Health Check Configuration

Table 59. Policy-Based Router Mapping Options

Command Syntax and Usage
`slb pbr interval` *<1-60>* Configures the number of seconds between health checks. The default value is 30 seconds. **Command mode:** Global configuration
`slb pbr retry` *<1-63>* Configures the number of consecutive failed health checks required to determine that the link is down. The default value is 4. **Command mode:** Global configuration
`slb pbr restore` *<1-63>* Configures the number of consecutive positive health checks required to determine that the link is up. The default value is 2. **Command mode:** Global configuration
`show slb pbr` Displays the current policy-based routing configuration. **Command mode:** All

Port iFlow Configuration

The following table describes the port configuration commands that are specific to iFlow Director. For more information about configuring ports, see your *ISCLI Reference*.

Table 60. Port Commands

Command Syntax and Usage
`interface port` *<port alias or number>* Enter Interface port mode. **Command mode:** Global configuration
`[no] tag-pvid` Disables or enables VLAN tag persistence. When disabled, the VLAN tag is removed from packets whose VLAN tag matches the port PVID. The default setting is `disabled`. **Command mode:** Interface port
`[no] tagpvid-ingress` Enables or disables the option to insert a VLAN tag to packets that ingress a port using the PVID. When enabled, packets (untagged or single tagged) that ingress a port have an additional VLAN tag inserted, with the VLAN ID set to the port PVID. The default setting is `disabled`. **Command mode:** Interface port
`[no] tagskip-ingress` Skips the enforcement of VLAN tag memberships on the ingress port. If enabled, tagged packets are allowed to ingress the port if the VLAN tag in the packet does not belong to the port VLAN membership. If disabled, the port VLAN membership rules are enforced on the ingress port. The default setting is `disabled`. **Note**: When this option is enabled, the **learn** option must be disabled. **Command mode:** Interface port
`[no] tagskip-egress` Skips the enforcement of VLAN tag memberships on the egress port. If enabled, tagged packets are allowed to egress the port even if the VLAN tag in the packet does not belong to the port VLAN membership. If disabled, the port VLAN membership rules are enforced on the egress port. The default setting is `disabled`. **Command mode:** Interface port

Table 60. Port Commands

Command Syntax and Usage
`[no] copy-arp-to-mp` Enables or disables the ability to forward a copy of ARP request packets on the ingress port to the management processor (MP). If enabled, when an ARP request packet is received on the ingress port, the packet will be flooded out on all ports with the same VLAN membership as the ingress port and a copy is sent to the MP. If disabled, when and an ARP request packet is received on the ingress port, the packet are flooded out on all ports with the same VLAN membership as the ingress port, but a copy is not sent to the MP. The default setting is `enabled`. **Command mode:** Interface port
`[no] retain-ingress-vlan` Enables or disables the ability to retain the ingress VLAN when traffic is steered from external servers to the blade servers. The default value is `disabled` (no). **Command mode:** Interface port
`ecmphash sip|dip` Sets the ECMP hashing algorithm to filter on the source (`sip`) or destination (`dip`) IP address. The default value is `sip`. **Command mode:** Interface port

Port ACL Configuration

Table 61. Port ACL Options

Command Syntax and Usage
`access-control list` *<ACL number>* Adds the specified ACL to the port. You can add multiple ACL lists to a port. **Command mode:** Interface port
`no access-control list` *<ACL number>* Deletes the specified ACL from the port. **Command mode:** Interface port
`access-control group` *<ACL group number>* Adds the specified ACL group to the port. You can add multiple ACL groups to a port. **Command mode:** Interface port

Table 61. Port ACL Options

Command Syntax and Usage
`no access-control group` *<ACL group number>* Removes the specified ACL group from the port. **Command mode:** Interface port
`show interface port {`*<port alias or number>*`} access-control` Displays current ACL port parameters. **Command mode:** All

Trunk Configuration

Trunk groups can provide super-bandwidth connections between iFlow Directors or other trunk capable devices. A *trunk* is a group of ports that act together, combining their bandwidth to create a single, larger port. Up to 18 trunk groups can be configured on the switch.

The following table describes the trunk configuration commands that are specific to iFlow Director. For more information about configuring trunks, see your *ISCLI Reference*.

Note: The trunk configurations options are only applicable when `slb method` is set to `trunk`.

Table 62. Trunk Configuration Options

Command Syntax and Usage
`[no] portchannel` *<1-18>* `slb-application` *<1-14>* Sets the SLB application group ID to be managed by this trunk ID reference for load balancing via trunk hashing. The default value is zero (0). **Command mode:** Global configuration

Failover Configuration

The following Layer 2 Failover commands are specific to iFlow Director.

Table 63. Layer 2 Failover Configuration Options

Command Syntax and Usage
`[no] failover vlan` Globally turns VLAN monitor `on` or `off`. When the VLAN Monitor is `on`, the switch automatically disables only internal ports that belong to the same VLAN as ports in the Failover trigger. The default value is `off`. **Command mode:** Global configuration
`failover enable` Globally turns Layer 2 Failover `on`. **Command mode:** Global configuration
`no failover enable` Globally turns Layer 2 Failover `off`. **Command mode:** Global configuration

Failover Trigger Configuration

Table 64. Failover Trigger Configuration Options

Command Syntax and Usage
`[no] failover trigger {<1-8>} enable` Enables or disables the Failover trigger. **Command mode:** Global configuration
`failover trigger {<1-8>} limit <0-1024>` Configures the minimum number of operational links allowed within each trigger before the trigger initiates a failover event. If you enter a value of zero (0), the switch triggers a failover event only when no links in the trigger are operational. **Command mode:** Global configuration
`show failover trigger {<1-8>}` Displays the current L2 Failover trigger settings. **Command mode:** All

Auto Monitor Configuration

Table 65. Auto Monitor Configuration Options

Command Syntax and Usage
`failover trigger {<1-8>} amon portchannel <trunk group number>` Adds a trunk group to the Auto Monitor. **Command mode:** Global configuration
`no failover trigger {<1-8>} amon portchannel <trunk group number>` Removes a trunk group from the Auto Monitor. **Command mode:** Global configuration
`failover trigger {<1-8>} amon adminkey <1-65535>` Adds a LACP *admin key* to the Auto Monitor. LACP trunks formed with this admin key will be included in the Auto Monitor. **Command mode:** Global configuration
`no failover trigger {<1-8>} amon adminkey <1-65535>` Removes a LACP *admin key* from the Auto Monitor. **Command mode:** Global configuration

Failover Manual Monitor - Monitor Configuration

Use this menu to define the port links to monitor. The Manual Monitor - Monitor configuration accepts only external uplink ports.

Table 66. Failover Manual Monitor - Monitor Options

Command Syntax and Usage
`failover trigger {`*<1-8>*`} mmon monitor member` *<port alias or number>* Adds the selected port to the Manual Monitor - Monitor configuration. **Command mode:** Global configuration
`no failover trigger {`*<1-8>*`} mmon monitor member` *<port alias or number>* Removes the selected port from the Manual Monitor - Monitor configuration. **Command mode:** Global configuration
`failover trigger` *<1-8>* `mmon monitor portchannel` *<trunk number>* Adds the selected trunk group to the Manual Monitor - Monitor configuration. **Command mode:** Global configuration
`no failover trigger` *<1-8>* `mmon monitor portchannel` *<trunk number>* Removes the selected trunk group from the Manual Monitor - Monitor configuration. **Command mode:** Global configuration
`failover trigger` *<1-8>* `mmon monitor adminkey` *<1-65535>* Adds an LACP *admin key* to the Manual Monitor - Monitor configuration. LACP trunks formed with this *admin key* will be included in the Manual Monitor - Monitor configuration. **Command mode:** Global configuration
`no failover trigger` *<1-8>* `mmon monitor adminkey` *<1-65535>* Removes an LACP *admin key* from the Manual Monitor - Monitor configuration. **Command mode:** Global configuration
`show failover trigger` *<1-8>* Displays the current Layer 2 Failover settings. **Command mode:** All

Failover Manual Monitor - Control Configuration

Use this menu to define the port links to control.

The Manual Monitor - Control configuration accepts internal and external ports, but not management ports.

Table 67. Failover Manual Monitor - Control Options

Command Syntax and Usage
`failover trigger {`<*1-8*>`} mmon control member` <*port alias or number*> Adds the selected port to the Manual Monitor - Control configuration. **Command mode:** Global configuration
`no failover trigger {`<*1-8*>`} mmon control member` <*port alias or number*> Removes the selected port from the Manual Monitor - Control configuration. **Command mode:** Global configuration
`failover trigger` <*1-8*> `mmon control portchannel` <*trunk number*> Adds the selected trunk group to the Manual Monitor - Control configuration. **Command mode:** Global configuration
`no failover trigger` <*1-8*> `mmon control portchannel` <*trunk number*> Removes the selected trunk group from the Manual Monitor - Control configuration. **Command mode:** Global configuration
`failover trigger` <*1-8*> `mmon control adminkey` <*1-65535*> Adds an LACP *admin key* to the Manual Monitor - Control configuration. LACP trunks formed with this *admin key* will be included in the Manual Monitor - Control configuration. **Command mode:** Global configuration
`no failover trigger` <*1-8*> `mmon control adminkey` <*1-65535*> Removes an LACP *admin key* from the Manual Monitor - Control configuration. **Command mode:** Global configuration
`show failover trigger` <*1-8*> Displays the current Failover settings. **Command mode:** All

Virtual Router Group Configuration

Virtual Router Group commands are used for associating all virtual routers into a single logical virtual router, which forces all virtual routers on the switch to either be master or backup as a group. For more detailed information about VRRP commands, refer to "VRRP Configuration" in the *Command Reference* for your switch.

Table 68. VRRP Virtual Router Group Configuration Options

Command Syntax and Usage
`group interface` *<interface number>* `[restricted]` Selects a switch IP interface. The default switch IP interface number is 1. **Restricted**: If you add this option, the switch will send advertisements on this VRRP group interface only. To clear this restriction, enter the command without the option. **Command mode:** Router VRRP
`show ip vrrp group` Displays the current configuration information for the virtual router group. **Command mode:** All

Virtual Router Group Priority Tracking Configuration

Table 69. Virtual Router Group Priority Tracking Options

Command Syntax and Usage
`[no] group track interfaces` When enabled, the priority for this virtual router will be increased for each other IP interface active on this switch. An IP interface is considered active when there is at least one active port on the same VLAN. This helps elect the virtual routers with the most available routes as the master. This command is disabled by default. **Command mode:** Router VRRP
`[no] group track ports` When enabled, the priority for this virtual router will be increased for each active port on the same VLAN. A port is considered "active" if it has a link and is forwarding traffic. This helps elect the virtual routers with the most available ports as the master. This command is disabled by default. **Command mode:** Router VRRP

Table 69. Virtual Router Group Priority Tracking Options

Command Syntax and Usage
`[no] group track slb-vrouter` Enables or disables tracking of VLAN Routers. When enabled, the priority for this virtual router is increased for each active SLB VLAN router reachable from this switch. This helps elect the virtual router with the most number of accessible routers as the master. The default setting is disabled. **Command mode:** Router VRRP
`show ip vrrp group track` Displays the current configuration for priority tracking for this virtual router. **Command mode:** All

Virtual Router Group Priority Tracking Port List

Only the link states of the physical ports are monitored when selective port tracking is configured. If LACP is enabled on a port, the port is considered operational when LACP is up on the port. Ports belonging to static or LACP trunks are treated as a single (aggregated) port, which is considered operational if at least one of the member ports in the trunk is active.

Table 70. VRRP Tracking Port List Options

Command Syntax and Usage
`group track portlist {`*<1-2>*`}` *<port alias or number>* Adds a physical port or ports to the current port list. These ports will be tracked when ports tracking is enabled. Place all members of a static or LACP trunk group into the same port list. You can add several ports, with each port separated by a comma (,) or a range of ports, separated by a dash (-). **Command mode:** Router VRRP
`no group track portlist {`*<1-2>*`}` *<port alias or number>* Removes a physical port or ports from the current port list. **Command mode:** Router VRRP
`no group track portlist {`*<1-2>*`}` Removes all ports from the port list. **Command mode:** Router VRRP
`show ip vrrp group track` Displays the current configuration for VRRP group priority tracking. **Command mode:** All

Access Control Configuration

Use these commands to create Access Control Lists and ACL Groups. ACLs define matching criteria used for IP filtering and Quality of Service functions.

Table 71. General ACL Configuration Options

Command Syntax and Usage
`[no] access-control list` *<1-508>* Configures an IPv4 Access Control List. To view command options, see page 143. **Command mode:** Global configuration
`[no] access-control list6` *<1-127>* Configures an IPv6 Access Control List. To view command options, see page 153. **Command mode:** Global configuration
`[no] access-control group` *<1-508>* Configures an ACL Group. To view command options, see page 156. **Command mode:** Global configuration
`show access-control` Displays the current ACL parameters. **Command mode:** All

ACL IPv4 Configuration

These commands allow you to define filtering criteria for each IPv4 Access Control List (ACL).

Table 72. ACL Configuration Commands

Command Syntax and Usage
`access-control list` *<1-508>* `egress-port port` *<port alias or number>* Configures the ACL to function on egress packets. **Command mode:** Global configuration
`no access-control list` *<1-508>* `egress-port` Removes the ACL from egress port functions. **Command mode:** Global configuration
`access-control list` *<1-508>* `action {permit|deny|set-priority` *<0-7>*`|` `redirect}` Configures a filter action for packets that match the ACL definitions. You can choose to permit (pass) or deny (drop) packets, or set the 802.1p priority level (0-7), or redirect traffic to a selected target. **Command mode:** Global configuration

Table 72. ACL Configuration Commands

Command Syntax and Usage
`[no] access-control list` *<1-508>* `counters` Enables or disables the statistics collection for the Access Control List. **Command mode:** Global configuration
`default access-control list` *<1-508>* Resets the ACL parameters to their default values. **Command mode:** Global configuration
`show access-control list` *<1-508>* Displays the current ACL parameters. **Command mode:** All

ACL Mirror Configuration

This menu allows you to define mirroring configuration for the selected ACL target.

Table 73. ACL Mirror Options

Command Syntax and Usage
`access-control list` *<1-508>* `mirror port` *<port alias or number>* Sets a single switch port as the target for traffic mirroring. **Command mode:** Global configuration
`no access-control list` *<1-508>* `mirror` Deletes this Target Mirroring instance. **Command mode:** Global configuration
`show access-control list` *<1-508>* `mirror` Displays the current traffic mirroring parameters for the ACL. **Command mode:** All

ACL Target Options Configuration

This menu allows you to define target options for an ACL.

Table 74. ACL Target Options

Command Syntax and Usage
`access-control list` *<1-508>* `target destination permit` Selects one redirection target destination method to permit. Traffic for the target destination will be forwarded via Layer 2. **Command mode:** Global configuration
`access-control list` *<1-508>* `target destination slb-application` *<1-14>* Sets the SLB application group ID as the target for traffic redirection. **Command mode:** Global configuration

Table 74. ACL Target Options (continued)

Command Syntax and Usage
`access-control list <1-508> target destination trunk <1-18>` Sets a single trunk as the target for traffic redirection. Traffic will be forwarded out on one port selected within the trunk group. **Command mode:** Global configuration
`access-control list <1-508> target destination port <port alias or number>` Sets a single switch port as the target for traffic redirection. **Command mode:** Global configuration
`[no] access-control list <1-508> cluster` Defines the forwarding mode when an SLB application group is specified. If enabled, traffic is redirected to all active servers within the SLB application group. If disabled, traffic is redirected and load balanced to one active server within the SLB application group. The default value is `disabled`. **Command mode:** Global configuration **Note:** The `cluster` option is only applicable when `slb method` is set to `trunk`.
`[no] access-control list <1-508> target` Deletes this ACL target. **Command mode:** Global configuration
`show access-control list <1-508> target` Displays the current parameters for the ACL. **Command mode:** All

Target Options Mirror Configuration

This menu allows you to define mirroring configuration for the selected ACL target.

Table 75. ACL Target Mirror Options

Command Syntax and Usage
`access-control list <1-508> target mirror destination slb-application <1-14>` Sets the SLB application group ID as the target for traffic mirroring. **Command mode:** Global configuration
`access-control list <1-508> target mirror destination trunk <1-18>` Sets a single trunk as the target for traffic mirroring. Traffic will be mirrored to one port selected within the trunk group. **Command mode:** Global configuration

Table 75. ACL Target Mirror Options (continued)

Command Syntax and Usage
`access-control list` *<1-508>* `target mirror destination port` *<port alias or number>* Sets a single switch port as the target for traffic mirroring. **Command mode:** Global configuration
`no access-control list` *<1-508>* `target mirror` Deletes this Target Mirroring instance. **Command mode:** Global configuration
`show access-control list` *<1-508>* `target mirror` Displays the current parameters for the ACL. **Command mode:** All

Target Application Error Configuration

This menu allows you to define the Application Error options for an ACL.

Table 76. Target Application Error Options

Command Syntax and Usage
`access-control list` *<1-508>* `target application-error application-list` *<0-14>* Defines the list of SLB application groups for failure monitoring. If any one application group in the list fails, the alternate redirection action will be executed. If `application-list` is set to 0, the application error checking is disabled. The default value is 0 (zero). **Command mode:** Global configuration
`access-control list` *<1-508>* `target application-error destination port` *<port alias or number>* Sets a single switch port as the pass through path if an application defined in the `application-list` fails. **Command mode:** Global configuration
`access-control list` *<1-508>* `target application-error destination trunk` *<1-18>* Sets a single trunk as the pass through path if an application defined in the `application-list` fails. **Command mode:** Global configuration
`access-control list` *<1-508>* `target application-error destination permit` Selects an alternate redirection target destination to use if an application defined in `application-list` fails. When `permit` is selected, traffic for the target destination will be forwarded via Layer 2 for the error path. **Command mode:** Global configuration

Table 76. Target Application Error Options (continued)

Command Syntax and Usage
`no access-control list <1-508> target application-error` Deletes this Application Error handling instance. **Command mode:** Global configuration
`show access-control list <1-508> target application-error` Displays the current parameters for the ACL. **Command mode:** All

Ethernet Filtering Configuration

These commands allow you to define Ethernet matching criteria for an ACL.

Table 77. Ethernet Filtering Configuration Options

Command Syntax and Usage
`[no] access-control list {<1-508>} ethernet source-mac-address {<MAC address>} {<MAC mask>}` Defines the source MAC address for this ACL. **Command mode:** Global configuration
`[no] access-control list {<1-508>} ethernet destination-mac-address {<MAC address>} {<MAC mask>}` Defines the destination MAC address for this ACL. **Command mode:** Global configuration
`[no] access-control list {<1-508>} ethernet vlan {<VLAN ID>} {<VLAN mask>}` Defines a VLAN number and mask for this ACL. **Command mode:** Global configuration
`[no] access-control list {<1-508>} ethernet ethernet-type {arp|ip|ipv6|mpls|rarp|any|0xXXXX}` Defines the Ethernet type for this ACL. **Command mode:** Global configuration
`[no] access-control list {<1-508>} ethernet priority <0-7>` Defines the Ethernet priority value for the ACL. **Command mode:** Global configuration
`default access-control list {<1-508>} ethernet` Resets Ethernet parameters for the ACL to their default values. **Command mode:** Global configuration
`show access-control list {<1-508>} ethernet` Displays the current Ethernet parameters for the ACL. **Command mode:** All

IP version 4 Filtering Configuration

These commands allow you to define IPv4 matching criteria for an ACL.

Table 78. IP version 4 Filtering Configuration Options

Command Syntax and Usage
`[no] access-control list {`*<1-508>*`} ipv4 source-ip-address` *<IP address> {<IP mask>* Defines a source IP address for the ACL. If defined, traffic with this source IP address will match this ACL. Specify an IP address in dotted decimal notation. **Command mode:** Global configuration
`[no] access-control list {`*<1-508>*`}ipv4 destination-ip-address` *<IP address> <IP mask>* Defines a destination IP address for the ACL. If defined, traffic with this destination IP address will match this ACL. **Command mode:** Global configuration
`[no] access-control list {`*<1-508>*`} ipv4 protocol` *<0-255>* Defines an IP protocol for the ACL. If defined, traffic from the specified protocol matches this filter. Specify the protocol number. The following are some of the well-known protocols. Number Name 1 icmp 2 igmp 6 tcp 17 udp 89 ospf 112 vrrp **Command mode:** Global configuration
`[no] access-control list {`*<1-508>*`} ipv4 type-of-service` *<0-255>* Defines a Type of Service value for the ACL. For more information on ToS, refer to RFC 1340 and 1349. **Command mode:** Global configuration
`default access-control list {`*<1-508>*`} ipv4` Resets the IPv4 parameters for the ACL to their default values. **Command mode:** Global configuration
`show access-control list {`*<1-508>*`} ipv4` Displays the current IPV4 parameters. **Command mode:** All

TCP/UDP Filtering Configuration

These commands allow you to define TCP/UDP matching criteria for an ACL.

Table 79. TCP/UDP Filtering Configuration Options

<table>
<tr><th>Command Syntax and Usage</th></tr>
<tr><td><code>[no] access-control list {<i><1-508></i>} tcp-udp source-port <i><1-65535> <port mask (0x1-0xFFFF)></i></code>
Defines a source port for the ACL. If defined, traffic with the specified TCP or UDP source port will match this ACL. Specify the port number. The following are some of the well-known ports:
<table>
<tr><th>Number</th><th>Name</th></tr>
<tr><td><b>20</b></td><td>ftp-data</td></tr>
<tr><td><b>21</b></td><td>ftp</td></tr>
<tr><td><b>22</b></td><td>ssh</td></tr>
<tr><td><b>23</b></td><td>telnet</td></tr>
<tr><td><b>25</b></td><td>smtp</td></tr>
<tr><td><b>37</b></td><td>time</td></tr>
<tr><td><b>42</b></td><td>name</td></tr>
<tr><td><b>43</b></td><td>whois</td></tr>
<tr><td><b>53</b></td><td>domain</td></tr>
<tr><td><b>69</b></td><td>tftp</td></tr>
<tr><td><b>70</b></td><td>gopher</td></tr>
<tr><td><b>79</b></td><td>finger</td></tr>
<tr><td><b>80</b></td><td>http</td></tr>
</table>
<b>Command mode:</b> Global configuration</td></tr>
<tr><td><code>[no] access-control list {<i><1-508></i>} tcp-udp destination-port <i><1-65535> <port mask (0x1-0xFFFF)></i></code>
Defines a destination port for the ACL. If defined, traffic with the specified TCP or UDP destination port will match this ACL. Specify the port number as with <code>sport</code>.
<b>Command mode:</b> Global configuration</td></tr>
<tr><td><code>[no] access-control list {<i><1-508></i>} tcp-udp flags <i><flag (0x0-0x3f)></i></code>
Defines a TCP/UDP flag for the ACL.
<b>Command mode:</b> Global configuration</td></tr>
<tr><td><code>default access-control list {<i><1-508></i>} tcp-udp</code>
Resets the TCP/UDP parameters for the ACL to their default values.
<b>Command mode:</b> Global configuration</td></tr>
<tr><td><code>show access-control list {<i><1-508></i>} tcp-udp</code>
Displays the current TCP/UDP Filtering parameters.
<b>Command mode:</b> All</td></tr>
</table>

Packet Format Filtering Configuration

These commands allow you to define Packet Format matching criteria for an ACL.

Table 80. Packet Format Filtering Configuration Options

Command Syntax and Usage
`[no] access-control list {<1-508>} packet-format ethernet {ethernet-type2|snap|llc}` Defines the Ethernet format for the ACL. **Command mode:** Global configuration
`[no] access-control list {<1-508>} packet-format tagging {any|none|tagged}` Defines the tagging format for the ACL. **Command mode:** Global configuration
`[no] access-control list {<1-508>} packet-format ip {ipv4|ipv6}` Defines the IP format for the ACL. **Command mode:** Global configuration
`default access-control list {<1-508>} packet-format` Resets Packet Format parameters for the ACL to their default values. **Command mode:** Global configuration
`show access-control list {<1-508>} packet-format` Displays the current Packet Format parameters for the ACL. **Command mode:** All

ACL Metering Configuration

These commands define the Access Control profile for the selected ACL or ACL Group.

Table 81. ACL Metering Configuration Options

Command Syntax and Usage
`access-control list {<1-508>} meter committed-rate <64-10000000>` Configures the committed rate, in Kilobits per second. The committed rate must be a multiple of 64. **Command mode:** Global configuration
`access-control list {<1-508>} meter maximum-burst-size <32-4096>` Configures the maximum burst size, in Kilobits. Enter one of the following values for `mbsize`: 32, 64, 128, 256, 512, 1024, 2048, 4096. **Command mode:** Global configuration
`[no] access-control list {<1-508>} meter enable` Enables or disables ACL Metering. **Command mode:** Global configuration

Table 81. ACL Metering Configuration Options (continued)

Command Syntax and Usage
`access-control list {<1-508>} meter action {drop|pass}` Configures the ACL Meter to either drop or pass out-of-profile traffic. **Command mode:** Global configuration
`[no] access-control list {<1-508>} meter log` Enables or disables logging out-of-profile notifications. **Command mode:** Global configuration
`show access-control list {<1-508>} meter` Displays current ACL Metering parameters. **Command mode:** All

ACL Re-Mark Configuration

You can choose to re-mark IP header data for the selected ACL or ACL Group. You can configure different re-mark values, based on whether packets fall within the ACL Metering profile, or out of the ACL Metering profile.

Re-Marking In-Profile Configuration

Table 82. ACL Re-Mark Configuration Options

Command Syntax and Usage
`access-control list {<1-508>} re-mark in-profile dscp <0-63>` Sets the DiffServ Code Point (DSCP) of In-Profile packets to the selected value. **Command mode:** Global configuration
`show access-control list {<1-508>} re-mark` Displays current Re-Mark parameters. **Command mode:** All

Update User Priority Configuration

Table 83. ACL User Priority Configuration Options

Command Syntax and Usage
`access-control list {<1-508>} re-mark in-profile dot1p <0-7>` Defines 802.1p value. The value is the priority bits information in the packet structure. **Command mode:** Global configuration
`[no] access-control list {<1-508>} re-mark in-profile use-tos-precedence` Enable or disable mapping of TOS (Type of Service) priority to 802.1p priority for In-Profile packets. When enabled, the TOS value is used to set the 802.1p value. **Command mode:** Global configuration
`show access-control list {<1-508>} re-mark` Displays current Re-Mark parameters. **Command mode:** All

Re-Marking Out-of-Profile Configuration

Table 84. ACL Out-of-Profile Configuration Options

Command Syntax and Usage
`access-control list {<1-508>} re-mark out-profile dscp <0-63>` Sets the DiffServ Code Point (DSCP) of Out-of-Profile packets to the selected value. The switch sets the DSCP value on Out-of-Profile packets. **Command mode:** Global configuration
`show access-control list {<1-508>} re-mark` Displays current Re-Mark parameters. **Command mode:** All

ACL IPv6 Configuration

These commands allow you to define filtering criteria for each IPv6 Access Control List (ACL).

Table 85. IPv6 ACL Options

Command Syntax and Usage
[no] access-control list6 <*1-127*> egress-port port <*port alias or number*> Configures the ACL to function on egress packets. **Command mode:** Global configuration
access-control list6 <*1-127*> action {permit|deny|set-priority <*0-7*>} Configures a filter action for packets that match the ACL definitions. You can choose to permit (pass) or deny (drop) packets, or set the 802.1p priority level (0-7). **Command mode:** Global configuration
[no] access-control list6 <*1-127*> statistics Enables or disables the statistics collection for the Access Control List. **Command mode:** Global configuration
default access-control list6 <*1-127*> Resets the ACL parameters to their default values. **Command mode:** Global configuration
show access-control list <*1-127*> Displays the current ACL parameters. **Command mode:** All

IP version 6 Filtering Configuration

These commands allow you to define IPv6 matching criteria for an ACL.

Table 86. IP version 6 Filtering Options

Command Syntax and Usage
[no] access-control list6 <*1-127*> ipv6 source-address <*IPv6 address*> <*prefix length (1-128)*> Defines a source IPv6 address for the ACL. If defined, traffic with this source address will match this ACL. **Command mode:** Global configuration
[no] access-control list6 <*1-127*> ipv6 destination-address <*IPv6 address*> <*prefix length (1-128)*> Defines a destination IPv6 address for the ACL. If defined, traffic with this destination address will match this ACL. **Command mode:** Global configuration

Table 86. IP version 6 Filtering Options (continued)

Command Syntax and Usage
`[no] access-control list6 <1-127> ipv6 next-header <0-255>` Defines the next header value for the ACL. If defined, traffic with this next header value will match this ACL.
`[no] access-control list6 <1-127> ipv6 flow-label <0-0xFFFFF>` Defines the flow label for the ACL. If defined, traffic with this flow label will match this ACL.
`[no] access-control list6 <1-127> ipv6 traffic-class <0-255>` Defines the traffic class for the ACL. If defined, traffic with this traffic class will match this ACL.
`default access-control list6 <1-127> ipv6` Resets the IPv6 parameters for the ACL to their default values. **Command mode:** Global configuration
`show access-control list6 <1-127> ipv6` Displays the current IPv6 parameters. **Command mode:** All

IPv6 TCP/UDP Filtering Configuration

These commands allows you to define TCP/UDP matching criteria for an ACL.

Table 87. IPv6 ACL TCP/UDP Filtering Options

Command Syntax and Usage
`[no] access-control list6` *<1-127>* `tcp-udp source-port` *<1-65535> <port mask (0x0001-0xFFFF)>* Defines a source port for the ACL. If defined, traffic with the specified TCP or UDP source port will match this ACL. Specify the port number. The following are some of the well-known ports: **Number** **Name** 20 `ftp-data` 21 `ftp` 22 `ssh` 23 `telnet` 25 `smtp` 37 `time` 42 `name` 43 `whois` 53 `domain` 69 `tftp` 70 `gopher` 79 `finger` 80 `http` **Command mode:** Global configuration
`[no] access-control list6` *<1-127>* `tcp-udp destination-port` *<1-65535> <port mask (0x0001-0xFFFF)>* Defines a destination port for the ACL. If defined, traffic with the specified TCP or UDP destination port will match this ACL. Specify the port number as with `sport`. **Command mode:** Global configuration
`[no] access-control list6` *<1-127>* `tcp-udp flags` *<flags (0x0-0x3F)> <flags mask (0x0-0x3F)>>* Defines a TCP/UDP flag for the ACL. **Command mode:** Global configuration
`default access-control list6` *<1-127>* `tcp-udp` Resets the TCP/UDP parameters for the ACL to their default values. **Command mode:** Global configuration
`show access-control list6` *<1-127>* `tcp-udp` Displays the current TCP/UDP Filtering parameters. **Command mode:** All

IPv6 Re-Mark Configuration

You can choose to re-mark IP header data for the selected ACL. You can configure different re-mark values, based on whether packets fall within the ACL metering profile, or out of the ACL metering profile.

IPv6 Re-Marking In-Profile Configuration

Table 88. IPv6 Re-Marking In-Profile Options

Command Syntax and Usage
`[no] access-control list6 <1-127> re-mark in-profile dot1p <0-7>` Re-marks the 802.1p value. The value is the priority bits information in the packet structure. **Command mode:** Global configuration
`[no] no access-control list6 <1-127> re-mark in-profile dscp <0-63>` Re-marks the DSCP value for in-profile traffic. **Command mode:** Global configuration
`[no] no access-control list6 <1-127> re-mark in-profile use-tos-precedence` Enables or disables mapping of TOS (Type of Service) priority to 802.1p priority for in-profile packets. When enabled, the TOS value is used to set the 802.1p value. **Command mode:** Global configuration
`default access-control list6 <1-127> re-mark` Sets the ACL re-mark parameters to their default values. **Command mode:** Global configuration
`show access-control list6 <1-127> re-mark` Displays current re-mark parameters. **Command mode:** All

ACL Group Configuration

These commands allow you to configure one or more ACLs to be added to an ACL Group. Once you create an ACL Group, you can assign the ACL Group to one or more ports.

Table 89. ACL Group Configuration Options

Command Syntax and Usage
`access-control group {<1-508>} list <1-508>` Adds the selected ACL to the ACL Group. **Command mode:** Global configuration
`no access-control group {<1-508>} list <1-508>` Removes the selected ACL from the ACL Group. **Command mode:** Global configuration

Table 89. ACL Group Configuration Options (continued)

Command Syntax and Usage
`access-control group {`*<1-508>*`} list6` *<1-127>* Adds the selected IPv6 ACL to the ACL Group. **Command mode:** Global configuration
`no access-control group {`*<1-508>*`} list6` *<1-127>* Removes the selected IPv6 ACL from the ACL Group. **Command mode:** Global configuration
`show access-control group` *<1-508>* Displays the current ACL group parameters. **Command mode:** All

Operations Commands

The Operations commands allow you to alter switch operational characteristics without affecting switch configuration. This section describes the operations commands that are specific to iFlow Director. For more information about operations commands, see your *ISCLI Reference*.

The following topics are discussed in this section:

- "Operations-Level SLB Options" on page 158
- "Software License Key Options" on page 161

Operations-Level SLB Options

The operations-level Server Load Balancing options are used for temporarily disabling or enabling real servers.

Operations-Level SLB Application Options

Use this menu to assign switch ports to trunk hash buckets that support the application. This feature allows you to fine-tune the hashing to produce better traffic distribution across application ports.

Note: The operations-level SLB applications options are only applicable when `slb method` is set to `trunk`.

Table 90. Application Operations

Command Syntax and Usage
`slb application` *<1-14>* `bucket` *<1-64>* `port` *<port alias or number>* Adds a switch port to the selected hash bucket. **Command mode:** Privileged EXEC
`slb application` *<1-14>* `update` Executes changes to the operating hash matrix during runtime operation. **Command mode:** Privileged EXEC
`slb application` *<1-14>* `reset` Resets the hash matrix to its current operating parameters. Use this command to remove changes that have not been updated, and revert to the current running matrix. **Command mode:** Privileged EXEC
`slb application` *<1-14>* `clear` Clears the current operating hash matrix. This change is not executed until you issue the `update` command. **Command mode:** Privileged EXEC
`show slb application` *<1-14>* `operational-state` Displays the current operational parameters for the application. **Command mode:** All

Operations-Level SLB Real Server Options

Table 91. Real Server Operations Options162

Command Syntax and Usage
`slb real-server` *<1-84>* `enable` Enables the real server. The real server will be returned to its configured operation mode when the switch is reset. **Command mode:** Privileged EXEC
`no slb real-server` *<1-84>* `enable` Disables the real server. The real server will be returned to its configured operation mode when the switch is reset. **Command mode:** Privileged EXEC
`show slb real-server` *<1-84>* `operational-state` Displays the current real server operational state. **Command mode:** All

Operations-Level SLB Restore Options

Use this menu to selectively restore service to a primary server forced into a standby state during runtime operations.

Note: The `replace` and `recover` options are only applicable when `slb method` is set to `trunk`.

Table 92. Server Restoration Options

Command Syntax and Usage
`slb real-server` *<1-84>* `replace` Restores a standby, primary server into active service under the following circumstances. When a backup server is active, a "healthy" primary server maybe forced into a standby mode due to the `preempt dis` (disable pre-emption) option. As a result, the `replace` command can be used to operationally replace the active backup server with the restored primary server without re-mapping the application group. The `replace` command applies only to standby, primary servers with a configured backup and disregards the rules for `preempt dis` on the selected/associated blade(s) from the standby portmap, but still honor the rules for `remap dis` (disable re-mapping) on the associated application group. **Command mode:** Privileged EXEC
`slb real-server` *<1-84>* `recover` Restores a real server into active service under the following circumstances. When a failed server is restored, it may be added to the active portmap but not in the hashmap due to the `remap dis` (disable re-mapping) option. As a result, the `recover` option can be used to operationally restore the server to active service by injecting and re-mapping the server back into the application group hashmap. The `recover` option applies to real servers with or without a configured backup and disregards the rules for the `preempt dis` (disable pre-emption) on the selected/associated blade(s) from the standby portmap and `remap dis` (disable re-mapping) on the associated application group. **Command mode:** Privileged EXEC
`show slb real-server` *<1-84>* `operational-state` Displays the current real server operational state. **Command mode:** All

Software License Key Options

Table 93. License Key Options

Command Syntax and Usage
`software-key` *<feature name> <key code>* Allows you to unlock iFlow Director. You are prompted to enter the feature name (`ibmiflow`), and the license key code. **Command mode:** Privileged EXEC
`no software-key` *<feature name>* Removes the license key for iFlow Director. **Command mode:** Privileged EXEC

Appendix A. Getting help and technical assistance

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your system, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure that you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure that they are connected.
- Check the power switches to make sure that the system and any optional devices are turned on.
- Use the troubleshooting information in your system documentation, and use the diagnostic tools that come with your system. Information about diagnostic tools is in the *Problem Determination and Service Guide* on the IBM *Documentation* CD that comes with your system.
- Go to the IBM support website at http://www.ibm.com/systems/support/ to check for technical information, hints, tips, and new device drivers or to submit a request for information.

You can solve many problems without outside assistance by following the troubleshooting procedures that IBM provides in the online help or in the documentation that is provided with your IBM product. The documentation that comes with IBM systems also describes the diagnostic tests that you can perform. Most systems, operating systems, and programs come with documentation that contains troubleshooting procedures and explanations of error messages and error codes. If you suspect a software problem, see the documentation for the operating system or program.

Using the documentation

Information about your IBM system and pre-installed software, if any, or optional device is available in the documentation that comes with the product. That documentation can include printed documents, online documents, ReadMe files, and Help files. See the troubleshooting information in your system documentation for instructions for using the diagnostic programs. The troubleshooting information or the diagnostic programs might tell you that you need additional or updated device drivers or other software. IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. To access these pages, go to http://www.ibm.com/systems/support/ and follow the instructions. Also, some documents are available through the IBM Publications Center at http://www.ibm.com/shop/publications/order/.

Getting help and information on the World Wide Web

On the World Wide Web, the IBM website has up-to-date information about IBM systems, optional devices, services, and support. The address for IBM System x® and xSeries® information is http://www.ibm.com/systems/x/. The address for IBM BladeCenter information is http://www.ibm.com/systems/bladecenter/. The address for IBM IntelliStation® information is http://www.ibm.com/intellistation/.

You can find service information for IBM systems and optional devices at http://www.ibm.com/systems/support/.

Software service and support

Through IBM Support Line, you can get telephone assistance, for a fee, with usage, configuration, and software problems with System x and x Series servers, BladeCenter products, IntelliStation workstations, and appliances. For information about which products are supported by Support Line in your country or region, see http://www.ibm.com/services/sl/products/.

For more information about Support Line and other IBM services, see http://www.ibm.com/services/, or see http://www.ibm.com/planetwide/ for support telephone numbers. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

Hardware service and support

You can receive hardware service through your IBM reseller or IBM Services. To locate a reseller authorized by IBM to provide warranty service, go to http://www.ibm.com/partnerworld/ and click **Find Business Partners** on the right side of the page. For IBM support telephone numbers, see http://www.ibm.com/planetwide/. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

In the U.S. and Canada, hardware service and support is available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9 a.m. to 6 p.m.

IBM Taiwan product service

台灣IBM產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

IBM Taiwan product service contact information:

IBM Taiwan Corporation
3F, No 7, Song Ren Rd.
Taipei, Taiwan
Telephone: 0800-016-888

Index