

53-1002805-04
31 October 2013



Brocade MLX Series and NetIron Family

Documentation Updates

Supporting Multi-Service IronWare R05.4.00e

BROCADE

Copyright © 2013 Brocade Communications Systems, Inc. All Rights Reserved.

Brocade, Brocade Assurance, the B-wing symbol, BigIron, DCX, Fabric OS, FastIron, MLX, NetIron, SAN Health, ServerIron, Turbolron, VCS, and VDX are registered trademarks, and AnyIO, Brocade One, CloudPlex, Effortless Networking, ICX, NET Health, OpenScript, and The Effortless Network are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned may be trademarks of their respective owners.

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

The authors and Brocade Communications Systems, Inc. shall have no liability or responsibility to any person or entity with respect to any loss, cost, liability, or damages arising from the information contained in this book or the computer programs that accompany it.

The product described by this document may contain "open source" software covered by the GNU General Public License or other open source license agreements. To find out which open source software is included in Brocade products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit <http://www.brocade.com/support/oscd>.

Brocade Communications Systems, Incorporated

Corporate and Latin American Headquarters
Brocade Communications Systems, Inc.
130 Holger Way
San Jose, CA 95134
Tel: 1-408-333-8000
Fax: 1-408-333-8101
E-mail: info@brocade.com

European Headquarters
Brocade Communications Switzerland Sàrl
Centre Swissair
Tour B - 4ème étage
29, Route de l'Aéroport
Case Postale 105
CH-1215 Genève 15
Switzerland
Tel: +41 22 799 5640
Fax: +41 22 799 5641
E-mail: emea-info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Systems China HK, Ltd.
No. 1 Guanghua Road
Chao Yang District
Units 2718 and 2818
Beijing 100020, China
Tel: +8610 6588 8888
Fax: +8610 6588 9999
E-mail: china-info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Systems Co., Ltd. (Shenzhen WFOE)
Citic Plaza
No. 233 Tian He Road North
Unit 1308 - 13th Floor
Guangzhou, China
Tel: +8620 3891 2000
Fax: +8620 3891 2111
E-mail: china-info@brocade.com

Document History

Title	Publication number	Summary of changes	Date
<i>Brocade MLX Series and NetIron Family Documentation Updates</i>	53-1002805-01	New document	19 December 2012
<i>Brocade MLX Series and NetIron Family Documentation Updates</i>	53-1002805-02	NetIron 05.4.00c Release updates.	26 March 2013
<i>Brocade MLX Series and NetIron Family Documentation Updates</i>	53-1002805-03	NetIron 05.4.00d Release updates.	9 July 2013
<i>Brocade MLX Series and NetIron Family Documentation Updates</i>	53-1002805-04	NetIron 05.4.00e Release updates. Update to the switch fabric module LED status.	31 October 2013

Contents

About This Document

In this chapter	ix
How this document is organized	ix
Brocade resources	ix
Getting technical help.....	ix
Document feedback	x

Chapter 1

Documentation Updates for the Brocade MLX Series and Netron Family Configuration Guide

In this chapter	1
Support for IPv6 anycast addresses	3
New LAG formation rule	3
Deleting CSPF groups	3
Deleting a CSPF group.....	3
IPv6 Traceroute over an MPLS network.....	5
IPv6 VRRP-E short path forwarding for MCT	10
IPv6 VRRP-E short-path forwarding and revertible option	10
IPv6 VRRP-E short-path forwarding delay.....	11
VRRP and VRRP-E support for routing over VPLS on CES and CER.....	14
Hybrid port mode OpenFlow.....	15
Hybrid port mode operation	16
Configuring hybrid port mode OpenFlow	16
Feature information.....	16
Limitations and prerequisites	16
sFlow null0 sampling	20
Support matrix for MSTP and STP global/STP/RSTP	22
Aggregated TM VOQ statistics collection.....	23
Supported modules.....	23
Displaying TM statistics from one queue or all queues	23
Displaying TM statistics from the multicast queue	25
Displaying QoS packet and byte counters.....	26
Multi-Chassis Trunk (MCT) client-interfaces delay	27
Deletion of ACLs bound to an interface.....	28
Configuring an encrypted syslog server.....	30
Displaying the configured server connections.....	31

Global ACL command to delete ACLs bound to an interface	32
Changing the router ID.	32
Show lag.	34
OpenFlow Hybrid Port Mode for IPv6.	35
Bypass LSP Liberal Path Selection	36
Current algorithm.	36
New algorithm	37
Show command enhancements.	38
Max Queue Depth and Buffer Utilization CLI enhancements	39
Displaying Traffic Manager max queue depth summary.	39
Displaying Traffic Manager maximum buffer utilization	40
Transparent forwarding of L2 and L3 protocols on a VLL for CES and CER	
41	
Forward Error Correction mode	43
Manual deletion of an OpenFlow rule	44
Show tech enhancement for OpenFlow.	44
Root Guard.	45
Discontinuing FID updates	46
Change the max-response-time value	46
Clearing the QoS packet and byte counters	46
IP assignment within a LAG	47
Update to Chapter 17 of the NetIron 5.4.00a Configuration Guide	47
STP feature configuration.	48
Fast port span	48
Fast Uplink Span	50
Protecting against UDP attacks	54
ACL accounting on Brocade NetIron CES and Brocade NetIron CER	
devices.	54
Displaying VLAN information	55
Sflow sampling on Brocade NetIron CES and	
Brocade NetIron CER devices.	55
LACP Enhancement.	56
LACP flap counters.	56
CSPF limitation	58
Fabric Auto Tuning SNMP and syslog enhancement	58
TM CLI command changes	58
FE command changes	60
Default global metric for ISIS	62
Configuration steps	62
ISIS Show command	63
Configuring Secure Shell and Secure Copy.	64
Configuring DSA or RSA public key authentication	64

Data Integrity Protection for Metro	64
Configuring Data Integrity Protection for Metro	65
New configuration commands	65
New show commands	65
Syslog messages	66
Management module redundancy overview	66
Globally changing the IP MTU	66
Tuning multicast parameters	67
Displaying a traffic engineering path to a destination	69
Egress port and priority based rate shaping	71
Multicast queue size, flow control, rate shaping and egress buffer threshold	71
Ingress traffic shaping per multicast stream	73
Implementation considerations	73
Configuring multicast traffic policy maps	74
Binding multicast traffic policy maps	75
Configuration example for rate shaping IPTV multicast stream	76
Tuning multicast parameters	77
Applying traffic policing parameters directly to a port	78
Deprecated MPLS Commands	79
OSPF with MCT	79
Site-Local IPv6 Addresses	79
IPv6 ND Router Advertisement Control	79
Client-role-revertible-delay timer	80
Port-based Rate Limiting	80
MSTP support for PBB	80
RPF	81
Configuration considerations for RPF	81
Keep-alive VLAN	81
Configuring FDP	81
Enabling interception of CDP packets globally	82
Configuring VPLS endpoint over FDP/CDP interface	83
Configuring VLL endpoint over FDP/CDP enabled interface	84
PIM over MCT	86
MCT feature interaction	86
Show lag_ecmp_port command	87
Match command	93
Match command options	93
Command examples	93

Chapter 2 Documentation updates for Multi-Service IronWare Diagnostic Guide

Management module diagnostics	95
Running management module diagnostics	95

Chapter 3	Documentation updates for Unified IP MIB Reference	
	RFC 4293: Management Information Base for the Internet Protocol (IP)	99
	Fabric drop count	103
	brcdNPCSRAMErrorTable (to query for NP CSRAM errors)	104
	brcdNPLPMRAMErrorTable (to query for NP LPM-RAM errors) . . .	105
	Traps	106
	Agent board table	107
 Chapter 4	 Documentation Updates for the MLX Series and NetIron XMR Series Hardware Installation Guide	
	In this chapter	109
	100xGbE 2-port interface module	110
	Cooling system and fans.	110
	High-speed switch fabric modules	110
	Switch fabric modules	111
	Managing Routers and Modules	111
	Maintenance and Field Replacement	111
	Cable specifications	112
	Enabling and disabling management module CPU usage calculations 112	
	Installing NIBI-16-FAN-EXH-A fan assemblies	112
	Module identification correction	115
	Brocade MLX Series and NetIron XMR supplemental upgrade procedures 116	
	Upgrading MBRIDGE or MBRIDGE32 images on management modules	116

About This Document

In this chapter

- “How this document is organized” on page ix
- “Brocade resources” on page ix
- “Getting technical help” on page ix
- “Document feedback” on page x

How this document is organized

This document contains updates to the Multi-Service IronWare R05.4.00b product manuals. These updates include document fixes and changes covering new features. [Table 1](#) below list the most recently released Multi-Service IronWare R05.4.00b product manuals.

TABLE 1 Documentation supporting Multi-Service IronWare R05.4.00b

Publication Title	Fabric OS Release	Page Number	Publication Date
<i>Brocade MLX Series and NetIron Family Configuration Guide</i>	R05.4.00a and later	Updates on page 1 .	September 2012
<i>Brocade MLX Series and NetIron XMR Hardware Installation Guide</i>	R05.4.00a and later	No Updates	September 2012
<i>Brocade NetIron CES Series and NetIron CER Series Hardware Installation Guide</i>	R05.4.00a and later	No Updates	September 2012
<i>Multi-Service IronWare Software Upgrade Guide</i>	R05.4.00a and later	No Updates	September 2012
<i>Brocade MLX Series and NetIron XMR Diagnostics Guide</i>	R05.4.00a and later	No Updates	September 2012
<i>Unified IP MIB Reference</i>	R05.4.00a and later	No Updates	September 2012
<i>Brocade MLX Series and NetIron XMR YANG Guide</i>	R05.4.00a and later	No Updates	September 2012

Brocade resources

For the latest documentation, go to <http://www.brocade.com/ethernetproducts>

Getting technical help

For the latest Technical Support contact information including e-mail and telephone contact information, go to <http://www.brocade.com/services-support/index.page>.

Document feedback

Quality is our first concern at Brocade and we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you. Forward your feedback by email to:

`documentation@brocade.com`

Provide the title and version number of the document and as much detail as possible about your comment, including the topic heading and page number and your suggestions for improvement.

Documentation Updates for the Brocade MLX Series and NetIron Family Configuration Guide

In this chapter

The updates in this chapter are for the *Brocade MLX Series and NetIron Family Configuration Guide*, publication number 53-1002544-02, published September 2012.

The following features were added or modified as part of the 5.4.00b release.

- “Support for IPv6 anycast addresses” on page 3
- “New LAG formation rule” on page 3
- “Deleting CSPF groups” on page 3
- “IPv6 Traceroute over an MPLS network” on page 5
- “IPv6 VRRP-E short path forwarding for MCT” on page 10
- “VRRP and VRRP-E support for routing over VPLS on CES and CER” on page 14
- “Hybrid port mode OpenFlow” on page 15
- “sFlow null0 sampling” on page 20
- “Support matrix for MSTP and STP global/STP/RSTP” on page 22
- “Aggregated TM VOQ statistics collection” on page 23
- “Displaying QoS packet and byte counters” on page 26
- “Multi-Chassis Trunk (MCT) client-interfaces delay” on page 27
- “Deletion of ACLs bound to an interface” on page 28
- “Configuring an encrypted syslog server” on page 30

The following features were added or modified as part of the 5.4.00c release.

- “OpenFlow Hybrid Port Mode for IPv6” on page 35
- “Bypass LSP Liberal Path Selection” on page 36
- “Max Queue Depth and Buffer Utilization CLI enhancements” on page 39
- “Transparent forwarding of L2 and L3 protocols on a VLL for CES and CER” on page 41
- “Forward Error Correction mode” on page 43
- “Manual deletion of an OpenFlow rule” on page 44
- “Show tech enhancement for OpenFlow” on page 44
- “Root Guard” on page 45

In this chapter

The following features were added or modified as part of the 5.4.00d release.

- IEEE 802.1ag Connectivity Fault Management (CFM), the MEP mep-id range parameter was updated to 1-8191.
- “LACP Enhancement” on page 56
- “CSPF limitation” on page 58
- “Fabric Auto Tuning SNMP and syslog enhancement” on page 58
- “Default global metric for ISIS” on page 62
- “Data Integrity Protection for Metro” on page 64
- “Fabric drop count” on page 103
- “100xGbE 2-port interface module” on page 110
- “Tuning multicast parameters” on page 67

The following features were added or modified as part of the 5.4.00e release.

- “Applying traffic policing parameters directly to a port” on page 78
- “Deprecated MPLS Commands” on page 79
- “OSPF with MCT” on page 79
- “Site-Local IPv6 Addresses” on page 79
- “IPv6 ND Router Advertisement Control” on page 79
- “Client-role-reversible-delay timer” on page 80
- “Configuring FDP” on page 81
- “Enabling interception of CDP packets globally” on page 82
- “Configuring VPLS endpoint over FDP/CDP interface” on page 83
- “Configuring VLL endpoint over FDP/CDP enabled interface” on page 84
- “Show lag_ecmp_port command” on page 87

Support for IPv6 anycast addresses

In the NetIron 5.4.00a Configuration Guide, the list of unsupported features for Brocade MLX series and Brocade NetIron XMR devices incorrectly includes IPv6 anycast address.

Brocade MLX series and Brocade NetIron XMR devices support IPv6 anycast addresses starting in NetIron 5.4.00b.

New LAG formation rule

The 10Gx24-DM module ports can only be part of LAGs exclusively consisting of 24x10G ports. A LAG cannot have a mix of 24x10G module ports and any other 10G module ports.

Deleting CSPF groups

TABLE 1 Supported platforms

Features supported	Brocade NetIron XMR Series	Brocade MLX Series	Brocade NetIron CES 2000 Series BASE package	Brocade NetIron CES 2000 Series ME_PREM package	Brocade NetIron CES 2000 Series L3_PREM package	Brocade NetIron CER 2000 Series Base package	Brocade NetIron CER 2000 Series Advanced Services package
Deleting a MPLS CSPF group	Yes	Yes	No	Yes	No	No	Yes

This feature is an enhancement to all Brocade devices running MPLS, enabling users to delete all the CSPF fate-share groups using a single command. Users are required to confirm execution with a warning message. Previous implementations required users to delete each group individually. The enhancement is backward compatible so the earlier command continues to be supported on all Brocade devices running MPLS.

For additional MPLS CSPF fate-sharing group information, refer to the latest *Brocade MLXe and NetIron Family Configuration Guide* located at www.brocade.com.

Deleting a CSPF group

In this example, group3 has already been set up as a fate-sharing CSPF group. To delete this CSPF fate-sharing group, enter the following command in router MPLS mode.

```
Brocade(config-mpls)# no cspf-group group3
```

Syntax: [no] cspf-group group-name

The *group-name* variable specifies the name of the fate-sharing group and can be up to 128 characters. The objects that can be specified for a fate-sharing group are interface, point-to-point link, node, and subnet. The maximum number of CSPF fate-sharing groups that can be configured on a device is 1000. To delete each configuration group individually, enter the above command with the relevant value for the <group-name> argument.

Deleting CSPF groups

This feature enhancement allows you to delete all configured groups at once. Use a single **no cspf-group** command. This command is only available at the router-mpls level and takes no arguments.

Sample configuration

These are the commands for use with the feature.

```
Brocade (config) #router mpls
Brocade (config-mpls) #no cspf-group
This will delete all the CSPF groups
Do you want to continue? (enter 'y' or 'n'): y
Brocade (config-mpls)#
```

All the CSPF groups are deleted at once at this point.

NOTE

If there are no cspf-groups to delete, the system generates an error message.

```
Brocade (config-mpls) #no cspf-group
This will delete all the CSPF groups
Do you want to continue? (enter 'y' or 'n'): y
No CSPF-groups to delete
Brocade (config-mpls)#
```

IPv6 Traceroute over an MPLS network

TABLE 2 Supported platforms

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series Base package	Brocade Netron CER 2000 Series Advanced Services package
IPv6 Traceroute over an MPLS network	Yes	Yes	No	Yes	Yes	No	Yes

NOTE

IPv6 MPLS traceroute not supported on the BR-MLX-10Gx24-DM 24-port 10GbE module.

IPv6 traceroute behavior is similar to IPv4 traceroute. However, unlike IPv4 traceroute, IPv6 traceroute has a new 6PE label added during each hop across the MPLS cloud. Based on the IP header value, the node devices differentiate if the Internet Control Message Protocol version 6 (ICMPv6) echo request is from an IPv6 or IPv4 source device.

When the traceroute sends ICMPv6 echo request packets with a TTL value (hop limit) value of 1, the first router in the path replies with the *ttl-exceeded* error message to the source. The next packet has a TTL (hop limit) value of 2 and the second router replies with the *ttl-exceeded* error message. This process continues till the destination host receives the packets and returns an ICMPv6 Echo Reply message.

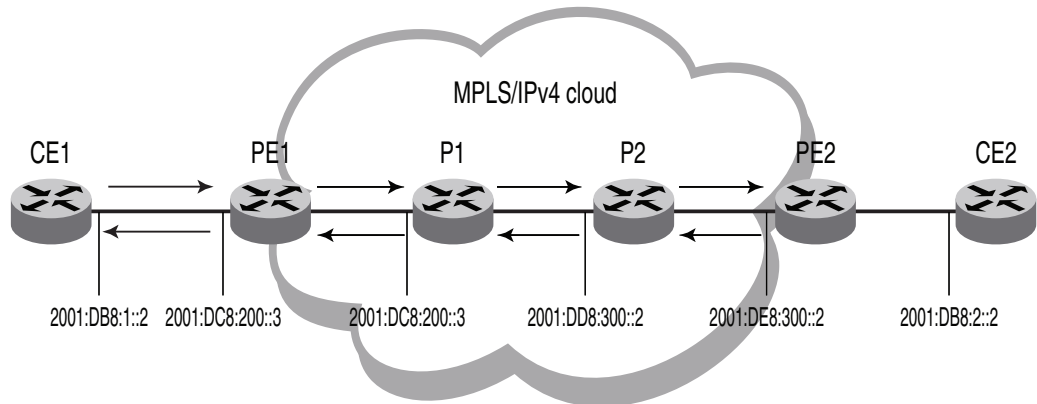
Based on the *ttl-exceeded* messages or the ICMPv6 Echo Reply messages received during the traceroute operation, the source device obtains details such as the hop sequence, total hops taken to complete the path, and the IPv4 or IPv6 addresses of devices that it passed during the path. For each hop, the traceroute gathers information about the hop number, best hop time, and the TTL value.

IPv6 Traceroute over an MPLS network

Tracing an IPv6 route through an MPLS domain

Figure 1 shows an MPLS-enabled provider network consisting of four LSRs. PE1 is the ingress PE Label Edge Router (LER), P1 and P2 are transit LSRs, and PE2 is the egress provider edge LER. CE1 and CE2 are CE devices located in different geographical locations.

FIGURE 1 IPv6 Traceroute in an MPLS cloud



To understand the IPv6 traceroute behavior in an MPLS domain, assume the following:

- Customer traffic is tunneled through a MPLS VPN network, and traffic within the MPLS core is forwarded by label-switching only.
- The CE1 router sends UDP packets from CE1 router towards the CE2 router.
- Traceroute is configured to generate ICMPv6 messages per ICMP extensions and to use LSPs to forward these messages. Refer to “[Configuring IPv6 Traceroute over MPLS](#)” on page 8 for more information.
- The PE routers are aware of the source and destination IPv6 addresses while the transit LSRs have no such knowledge.
- The **traceroute** command is issued from CE1 to CE2 and reports the following information:

```
Brocade# traceroute ipv6 2001:DB8:2::2
Type Control-c to abort
Tracing the route to IPv6 node 2001:DB8:2::2 from 1 to 30 hops

 1  <1 ms  <1 ms  <1 ms  2001:DB8:1::2
 2  <1 ms  <1 ms  <1 ms  2001:DC8:200::3
    MPLS Label=1026 Exp=0 TTL=1 S=0
    MPLS Label=794624 Exp=0 TTL=1 S=1
 3  <1 ms  <1 ms  <1 ms  2001:DD8:300::2
    MPLS Label=1029 Exp=0 TTL=1 S=0
    MPLS Label=794624 Exp=0 TTL=2 S=1
 4  <1 ms  <1 ms  <1 ms  2001:DE8:300::2
 5  <1 ms  <1 ms  <1 ms  2001:DB8:2::2
```

NOTE

The traceroute output reports information on a traceroute packet only when its TTL equals 1. Label stack information associated with subsequent routing of the ICMP message along the LSPs to the destination and back to the source is not displayed.

In the [Figure 1](#) scenario, the traceroute operation can be described as follows:

1. CE1 sends a traceroute probe with a TTL of 1 to its peer, CE2, with the destination IP address of 2001:DB8:2::2. PE1 decrements the packet's TTL by one and drops the expired packet. It generates a *ttl-exceeded* ICMPv6 message, and sends it back to CE1 with the source IPv6 address embedded in the IPv6 header of the expired packet. Traceroute reports the PE1 IPv6 address at hop 1, but there is no label information.

```
1. <1 ms <1 ms <1 ms 2001:DB8:1::2
```

2. CE1 sends a second traceroute probe to CE2, with an incremented TTL value of 2. PE1 decrements the TTL value to 1, and adds the 6PE label and the Label Distribution Protocol (LDP) label onto the packet to route it to CE2 by way of the transit router P1. PE1 also copies the TTL value from the IP header into the TTL field of the labels (recall that TTL propagation must be enabled on the ingress PE).

The transit router P1 decrements the TTL, drops the expired packet since the TTL value is 0, and generates a *ttl-exceeded* ICMPv6 message. Before dropping the packet, and using the ICMPv6 extension mechanism, P1 copies the packet's label stack plus its IP header and appends both to the ICMPv6 message. Though the message destination is CE1, P1 cannot return the ICMPv6 message directly to CE1. It uses label-switching to forward the encapsulated ICMP response in the direction of the original traceroute probe along the configured LSPs and back to CE1. P1 sets the maximum TTL value of 255 to ensure that the message can reach its destination before it times out.

Traceroute reports the IP address of P1, plus the label stack that was pushed onto the traceroute packet by PE1 and received by P1 when the packet's TTL was 1.

```
2 <1 ms <1 ms <1 ms 2001:DC8:200::3
   MPLS Label=1026 Exp=0 TTL=1 S=0
   MPLS Label=794624 Exp=0 TTL=1 S=1
```

3. The third traceroute probe (TTL=3) is forwarded until it expires at the transit router P2. P2 (the Penultimate Hop Popping (PHP) LSR) generates the ICMPv6 message, appends the label stack from the expired traceroute packet, and passes it on to PE2 without imposing a label. PE2 forwards the ICMPv6 message back to CE1 along the return LSP.

Traceroute reports the IP address of P2, plus the label stack which P2 received with the traceroute packet from P1 when the packet's TTL was 1.

```
3 <1 ms <1 ms <1 ms 2001:DD8:300::2
   MPLS Label=1029 Exp=0 TTL=1 S=0
   MPLS Label=794624 Exp=0 TTL=2 S=1
```

4. The fourth traceroute probe (TTL=4) is forwarded until it expires at the egress provider edge device PE2. PE2 drops the packet and generates a *ttl-exceeded* ICMPv6 message without label stack extension since there is no label stack to report.

Traceroute reports only the IP address of PE2. The transit router P2 popped the outer label before passing the traceroute packet on to the egress PE2 and PE2 pops the VPN label before sending the ICMPv6 message back to the customer source device CE1.

```
4 <1 ms <1 ms <1 ms 2001:DE8:300::2
5 <1 ms <1 ms <1 ms 2001:DB8:2::2
```

IPv6 Traceroute over an MPLS network

- The fifth traceroute probe (TTL=5) has a TTL large enough for the packets to reach the customer destination device CE2. CE2 generates an ICMPv6 *port unreachable* message, which CE2 sends back to CE1.

Traceroute reports only the IP address of the destination device CE2. No label extension is added because the received packet is not labeled. The *port unreachable* message is label-switched back to the customer source device CE1, as a normal data packet.

```
5    <1 ms    <1 ms    <1 ms  2001:DB8:2::2
```

Configuring IPv6 Traceroute over MPLS

The **ipv6 icmp mpls-response** command configures the behavior of the traceroute operation by controlling both the ICMPv6 message format (use ICMPv6 label stack extensions or not) and the manner in which the ICMPv6 messages are forwarded through an MPLS domain (by way of IP routing table lookup or through label-switching using LSPs).

MPLS response is enabled by default. To enable the MPLS response after it was disabled, enter the following command:

```
Brocade(config)# ipv6 icmp mpls-response
```

You can use this version of the command if the traceroute is over an IPv6-aware MPLS core. In such a case, IPv6 traceroute uses the default option of using the routing tables to forward packets. The IPv6 link local addresses should not be used to send the ICMPv6 packet. At the same time, you can still use the **ipv6 icmp mpls-response use-lsp** command to use the configured LSPs.

To specify using LSP to forward the ICMPv6 messages with MPLS label extensions, enter the following command:

```
Brocade(config)# ipv6 icmp mpls-response use-lsp
```

Use this version of the command if the MPLS core is non IPv6-aware, because the IPv6 forwarding will not work.

To specify generating ICMPv6 messages without MPLS label extensions, enter the following command:

```
Brocade(config)# ipv6 icmp mpls-response no-label-extensions
```

To disable the IPv6 Traceroute over MPLS feature, enter the following command:

```
Brocade(config)# no ipv6 icmp mpls-response
```

Syntax: [no] **ipv6 icmp mpls-response** [**use-lsp**] [**no-label-extension**]

The **mpls-response** parameter enables the ICMPv6 traceroute response in default mode. The feature is enabled by default and configured to use IP routing to forward ICMP messages.

The **use-lsp** parameter enables forwarding of ICMPv6 error messages along the LSPs configured for the MPLS domain. By default, using configured LSPs use is disabled.

The **no-label-extension** parameter disables the use of label stack information in the ICMPv6 error messages.

The **no** option disables the ICMPv6 traceroute response configuration. When the ICMP traceroute feature is disabled, standard traceroute using IPv6 forwarding is used to trace a traffic path through an MPLS domain.

NOTE

The **ipv6 icmp mpls-response** command supports TTL expiry for IPv6 packets only.

The output of the **show ipv6 traffic** command displays counts for ICMPv6 *ttl-exceeded* error reply packets.

IPv6 VRRP-E short path forwarding for MCT

TABLE 3 Supported devices for IPv6 VRRP-E short path forwarding for MCT.

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series BASE package	Brocade Netron CER 2000 Series Advanced Services package
IPv6 VRRP-E short path forwarding for MCT	Yes	Yes	Yes	Yes	Yes	Yes	Yes

For additional Multi-Chassis Trunking (MCT) information, refer to the latest *Brocade MLXe and Netron Family Configuration Guide* located at www.brocade.com.

IPv6 VRRP-E short-path forwarding and revertible option

Short-path forwarding enables the short path forwarding on an IPV6 VRRP-E device. It will revert back to standard behavior (no short-path forwarding) temporarily even if short-path forwarding is configured.

Configuration considerations

- VRRP-E virtual MAC will be synced and learned on ICL ports on backup routers through the ICL.
- ICL ports must be member ports of VLANs that CCEP ports are members of.
- VRRP or VRRP-E master router will be broadcast hello packets to all VLAN member ports including ICL ports. Normal VLAN FID will be used for broadcasting.
- VRRP or VRRP-E backup routers will not be flood back hello packets received from ICL ports to ICL ports, but will be flooded to other non- ICL ports.
- MCT switches must have complete routing information using static routes for L3 forwarding.
- For MCT switches configured with VRRP or VRRP-E, track-port features can be enabled to track the link status to the core switches so the VRRP or VRRP-E failover can be triggered.

NOTE

Brocade recommends disabling ICMP redirect globally to avoid unintended CPU forwarding of traffic when VRRP or VRRP-E is configured.

L3 traffic forwarding behaviors

When one MCT switch act as VRRP or VRRP-E master router and the other MCT switch is VRRP or VRRP-E backup, the following behavior will be seen:

- Packets sent to VRRP-E virtual IPv6 address will be L2 switched to the VRRP-E master router for forwarding.
- The VRRP-E MAC will be learned by the other MCT switch that acts as backup router.
- Both data traffic and VRRP-E control traffic will need to travel through ICL unless the short-path forwarding feature is enabled.

When both MCT devices act as the VRRP or VRRP-E backup routers, the following behavior will be seen:

- Packets sent to VRRP-E virtual IPv6 address will be L2 switched to the VRRP-E master router for forwarding.
- VRRP-E MAC will be learned by both MCT switches acting as backup routers.
- Both data traffic and VRRP-E control traffic will need to travel through ICL unless the short-path forwarding feature is enabled.

Under the IPv6 VRRP-E VRID configuration level, use the **short-path-forwarding** command. If the revertible option is not enabled, short path forwarding will be disabled if the VRRP-E router priority is below the revert-priority configured value. Use the following command to enable short path forwarding.

```
Brocade(config-if-e1000-vrid-2)# short-path-forwarding revert-priority 60
```

Syntax: **[no] short-path-forwarding [revert-priority value]**

Use the supplied priority value as a threshold to determine if the short-path-forwarding behavior should be effective or not. If one or more ports tracked by the track-port command go down, the current priority of IPv6 VRRP-E will be lowered by a specific amount configured in the track-port command for each port that goes down.

Once the current-priority is lower than the threshold, the short-path-forwarding will be temporarily suspended and revert back to the regular VRRP-E forwarding behavior without short-path-forwarding enabled.

The reverting behavior is only temporary. If one or more of the already down ports tracked by the track-port command come back, it is possible that the current priority of VRRP-E will be higher than the threshold again and the short-path-forwarding behavior will be resumed.

IPv6 VRRP-E short-path forwarding delay

Use IPv6 VRRP-e short-path forwarding delay to configure the time delay required to enable short path forwarding after reloading the backup router. When configured, short path forwarding will be enabled only after the configured delay time after the MP initialization is completed (from the time all modules in the system are UP). Default value is set to 0 seconds.

This is global IPv6 VRRP-E configuration will effect all IPv6 VRRP-E instances.

```
Brocade(config)# [no] short-path-forwarding-delay 100
```

Syntax: **short-path-forwarding-delay seconds**

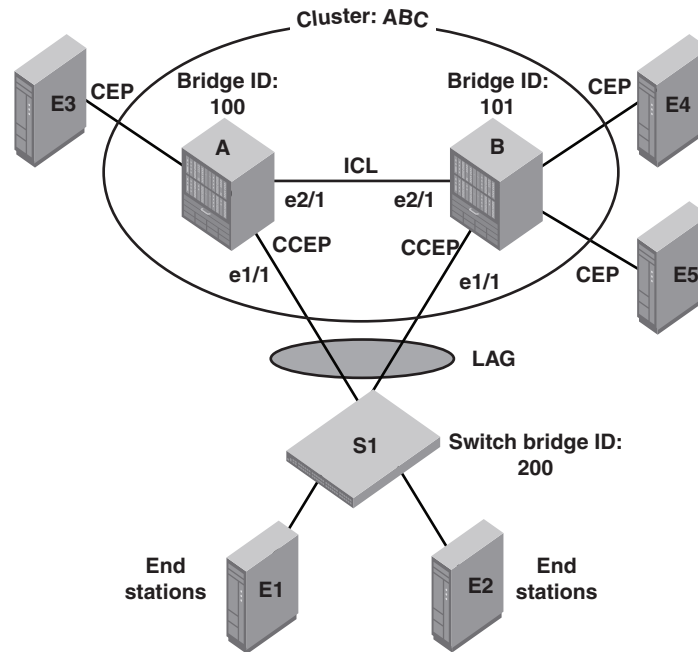
Sample configurations

```
Brocade(config)#short-path-forwarding-delay 100
Brocade(config)#ipv6 router vrrp-extended
Brocade(config-ipv6-vrrpe-router)#interface ve 10
Brocade(config-vif-10)# ipv6 address 2003::10:11/64
Brocade(config-vif-10)#ipv6 vrrp-extended vrid 10
Brocade(config-vif-10-ipv6-vrid-10)#backup priority 50
Brocade(config-vif-10-ipv6-vrid-10)#ipv6-address 2003::11:50
Brocade(config-vif-10-ipv6-vrid-10)#short-path-forwarding revert-priority 120
```

1

IPv6 VRRP-E short path forwarding for MCT

Sample MCT Configuration



Switch A:

```
vlan 4090
  tagged ethe 2/1
  router-interface ve 1
!
```

interface ve 1

```
ip address 192.168.1.1/24
!
```

cluster ABC

```
rbridge-id 100
session-vlan 4090
member-vlan 100 to 300
icl icl_a_b ethernet 2/1
peer 10.10.20.2 rbridge-id 101 icl icl_a_b
deploy
client switch_s1
  rbridge-id 200
  client-interface ethernet 1/1
  deploy
exit
!
```

IPv6 VRRP Configuration

```
vlan 200
  tagged ethe 1/1 ethe 2/1
  router-interface ve 10
!
```

Ipv6 router vrrp

```
interface ve 10
  ipv6 address 10::1/64
  ipv6 vrrp vrid 10
```

```

    backup priority 50
    ipv6-address 10::100
    activate
!
Switch B:
vlan 4090
    tagged ethe 2/1
    router-interface ve 1
!
interface ve 1
ip address 192.168.1.2/24
!
cluster ABC
rbridge-id 101
    session-vlan 4090
    member-vlan 100 to 300
    icl icl_a_b ethernet 2/1
peer 10.10.20.1 rbridge-id 100 icl icl_a_b
deploy
client switch_s1
    rbridge-id 200
    client-interface ethernet 1/1
deploy
exit
!
IPv6 VRRP Configuration
vlan 200
    tagged ethe 1/1 ethe 2/1
    router-interface ve 10
!
ipv6 router vrrp
interface ve 10
    ipv6 address 10::2/64
    ipv6 vrrp vrid 10
    backup priority 50
    ipv6-address 10::100
    activate
!

```

NOTE

Cluster client-rbridge-id on both switch A and B have to be same value for a given MCT.

Switch S1:

```

lag "mct_s1" static id 1
    ports ethernet 7/1 to 7/2
    primary-port 7/1
    deploy
!
vlan 200
    tagged ethe 7/1
    router-interface ve 10
!
interface ve 10
    ipv6 address 10::99/64

```

1 VRRP and VRRP-E support for routing over VPLS on CES and CER

VRRP and VRRP-E support for routing over VPLS on CES and CER

TABLE 4 VRRP and VRRP-E support on devices for routing over VPLS on CES and CER.

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series BASE package	Brocade Netron CER 2000 Series Advanced Services package
VRRP and VRRP-E support for routing over VPLS on CES and CER	Yes	Yes	No	Yes	No	No	Yes

Routing over VPLS was introduced in R05.4.00 for the Brocade MLX series and Netron XMR series routers. This release adds support for VRRP and VRRPE support for routing over VPLS on CES and CER.

For additional Routing over VPLS information, refer to the latest *Brocade MLXe and NetIron Family Configuration Guide* located at www.brocade.com.

Hybrid port mode OpenFlow


TABLE 5 Supported devices for hybrid port mode OpenFlow

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series BASE package	Brocade Netron CER 2000 Series Advanced Services package
Hybrid port mode OpenFlow	Yes	Yes	No	No	No	No	No

OpenFlow hybrid-enabled ports support both OpenFlow traffic forwarding and normal routing traffic forwarding. OpenFlow hybrid-enabled ports support "protected VLANs" and "unprotected VLANs". Protected VLANs are not subject to defined OpenFlow flows on the OpenFlow hybrid port mode ports. OpenFlow flows on a hybrid port mode port will not match any traffic on protected VLANs. Unprotected VLANs are subject to defined OpenFlow flows on the OpenFlow hybrid-port-mode port. OpenFlow flows on a hybrid-port-mode port are allowed to match on the traffic of unprotected VLANs.

Figure 2 shows a topology in which port 1/1 on Device-1 and port 1/4 on Device-2 are hybrid-port-mode OpenFlow ports with VLAN 10 as a configured protected VLAN. By configuring a virtual ethernet on a protected VLAN 10 and assigning an address to route the traffic of the nodes, you are able to send protected VLAN traffic between the nodes and route the traffic as per the routing table. Traffic flowing on other VEs created on top of other VLANs (the unprotected VLANs) treated as unprotected VLAN traffic and is subject to OpenFlow rules lookup. OpenFlow traffic can be forwarded through this port.

FIGURE 2 Hybrid port mode OpenFlow topology



Hybrid port mode operation

Consider Device-1 in [Figure 2](#). Ingress traffic on VLAN 10 on hybrid port 1/1 will be processed for normal routing. Traffic on other VLANs will be processed against OpenFlow flows on port 1/1 and switched accordingly. A preconfigured number of protected VLANs can be supported for normal routing. The Spanning tree protocols (STP) state of these routing VLANs will be set to forwarding, as the Layer 2 protocol is not supported. Normal routing is not supported on unprotected VLANs.

Configuring hybrid port mode OpenFlow

1. Enable OpenFlow at the global configuration level.
2. Configure OpenFlow controller configurations.
3. Configure the system maximum configuration for the maximum OpenFlow entries. (The default is 0.)
4. Configure the maximum OpenFlow flow-protected VLAN entries. (The default is 0.)

NOTE

System reload is required once you change the system maximum values.

5. Configure protected VLANs on the port. A maximum of 40 protected VLANs can be configured on an OpenFlow port.
6. Enable OpenFlow hybrid port mode on the desired interfaces.
7. Configure a VE for the interface by specifying the protected VLAN and add routing entries.

Feature information

- Switchover and HLOS are not supported. When the active management processor (MP) goes down, communication with the controller is brought down and the flow tables on the MP and all line processors (LP) are cleared. The connection with the controller is re-established after switchover.
- When LP is reset, the flow table on the LP is restored once the LP comes up and flows specific to that LP are maintained in the MP.
- When an OpenFlow enabled port goes up or down, no rules are removed. The addition or deletion of rules depends solely on the controller.
- 4K OpenFlow content-addressable memory (CAM) entries in OpenFlow CAM for normal OpenFlow entries are supported.
- 2K protected VLANs and ports combinations are supported.

Limitations and prerequisites

Brocade devices support version 1.0.0 of the OpenFlow protocol. The following limitations and prerequisites apply to the configuration of OpenFlow hybrid port mode.

- Normal IPv4 and IPv6 routing are not supported on unprotected VLANs.
- IPv4 routing is only supported on OpenFlow ports in this release.
- Layer 2 or L2VPN forwarding is not supported on ports in hybrid port mode ports because MAC learning is disabled on these ports.

- A port can be enabled for hybrid port mode only if the port is untagged in the default VLAN.
- Ports in OpenFlow hybrid port mode cannot be added as untagged ports to regular VLANs or L2VPN because this can cause a problem with topology discovery.
- As routing is enabled on a port in hybrid port mode, OpenFlow traffic or unprotected VLAN traffic sent with destination MAC address as the port's MAC address and matching IP route entries on the port can potentially find its VLAN and MAC address modified unless the OpenFlow rules explicitly set the VLAN and destination MAC address in the outgoing packet.
- Inbound normal ACL configuration is not supported on the port in hybrid port mode.
- Any port with the default VLAN not equal to the system default VLAN ID cannot be enabled for hybrid port mode.
- Policy based routing (PBR) is not supported.
- Protected VLAN traffic that does not have matching IP route entries will be dropped.
- Multiple interfaces cannot be part of a VE created on a port in hybrid port mode with a protected VLAN.
- The following are supported on protected VLANs:
 - IPv4 packets.
 - BGP, OSPF and IS-IS protocols.

NOTE

Layer 2 or L2VPN, VRF are not supported.

- When protected VLANs are configured but the port is not part of the VLAN, the traffic coming on the port with the protected VLAN will be dropped.
- Port in hybrid port mode OpenFlow doesn't support MPLS running on the same port.

These are the CLI commands for configuring different features for a hybrid port mode interface.

Enabling OpenFlow hybrid port mode

The purpose of the OpenFlow Enable command is to enable or disable hybrid port mode on the port and the port becomes a normal port on an interface. The **no** version of the command disables the hybrid port mode on the port and the port becomes a normal port.

```
Brocade(config-if-e10000-2/5)# openflow enable layer2 hybrid-mode
```

Syntax: [no] openflow enable layer2 | layer3 [hybrid-mode]

Adding or Deleting protected VLANs

The purpose of the OpenFlow protected -vlans command is to add or delete protected VLANs on a hybrid port-mode interface. The **no** form of the command is used to delete the configured protected VLANs from the hybrid port-mode port.

```
Brocade(config-if-e10000-2/5)# openflow protected-vlans 10
```

Syntax: [no] openflow protected-vlans id1 id2 ...idn

VLANs can be configured individually.

NOTE

A VLAN range is not specified for this command .

Setting the system maximum

The **system-max openflow-pvlan-entries** command sets the CAM size of OpenFlow protected VLAN entries for the device. By default, this value is set to 0.

```
Brocade(config)# system-max openflow-pvlan-entries 2000
```

Syntax: **system-max openflow-pvlan-entries** *value*

The value variable represents the number of port and protected VLAN combination entries that can be configured in the system. The range is from 0 to 2048. Once this command is used, you must reload the system.

Displaying OpenFlow configuration information

The show OpenFlow command displays the configuration for OpenFlow.

```
Brocade(config)# show openflow
Number of Controllers:      2
```

Controller 1:

```
Connection Mode:          passive, TCP,
Listening Address:        0.0.0.0
Connection Port:          6633
Connection Status:
SSL Connection             :False
No TCP connection found.
```

Controller 2:

```
Connection Mode:          active, TCP,
Controller Address:       10.20.101.199
Connection Port:          23
Connection Status:
Local IP address:port <-> Remote IP address:port TCP state   RcvQue  RxBuffe
SendQue  TxBuffe
10.20.178.73  8807      10.20.101.199  23    ESTABLISHED  0       0       0
0
SSL Connection             :False
Match Capabilty:
L2: Port, Source MAC, Destination MAC, Ether type, Vlan, Vlan PCP
L3: Port, Vlan, Vlan PCP, Source IP, Destination IP, IP Protocol, IP TOS, IP Src
Port, IP Dst Port
```

```
Normal Openflow Enabled Ports:      e2/1
Hybrid Mode ports                    Protected Vlan-IDs
e4/1                                  (100,101,102,103)
e7/2                                  (200)
Default action: drop
Maximum number of flows allowed: 4096
Maximum number of Protected Vlans allowed: 2048
```

Displaying Ethernet slots and ports

The show interface command gives the number of ports and their slots for the ethernet interface.

```
brocade(config-if-e10000-2/5)# show in ethernet 2/5
10GigabitEthernet2/5 is up, line protocol is up
  STP Root Guard is disabled, STP BPDU Guard is disabled
  Hardware is 10GigabitEthernet, address is 000c.dbf5.bd00 (bia 000c.dbf5.bd34)
  Configured speed 10Gbit, actual 10Gbit, configured duplex fdx, actual fdx
  Member of VLAN 11 (untagged), 1 L2 VLANS (tagged), port is in dual mode, port
state is
.....
.....
Openflow: Enabled, Openflow Index 53, Flow Type Layer2
  Openflow: Enabled, Openflow Index 53, Flow Type Layer2
  Openflow: Hybrid Mode  Openflow: Protected Vlans : 10 20 30
.....
```

sFlow null0 sampling

TABLE 6 Supported devices for sFlow null0 sampling

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series BASE package	Brocade Netron CER 2000 Series Advanced Services package
sFlow null0 sampling	Yes	Yes	No	No	No	No	No

This feature allows Brocade devices to sample null0 dropped packets. You will find this useful in cases such as DOS attack on a particular route.

Configuring steps

1. Enable sFlow.
2. Enable null0 sampling .
3. Configure null0 routes.

NOTE

Above commands can be performed in any order.

Feature highlights

- By default, null0 sFlow sampling feature is disabled.
- IPv4, IPv4-VPN, IPv6 null0 routes can be sFlow sampled.
- Only explicitly configured null0 routes can be sFlow sampled. Implicit null0 drops cannot be sFlow sampled.

Limitations

- When this feature is enabled, due to sampling of more packets (discarded packets) than the usual number till now, the actual sampling rate for regular streams will be reduced.
- This feature does not support PBR related null0 drops.
- This feature does not support default null0 route drops.

Backward compatibility

The current sFlow functionalities and ACL based sFlow functionalities will co-exist with this feature. As the dropped packets hit the traffic manager (TM), if mirroring is enabled on that port, these dropped packets will also get mirrored.

Enabling or disabling the null0 sFlow sampling

These commands include the enabling and disabling of the null0 sampling.

Enter the following command to enable sFlow sampling for null0 routes.

```
Brocade(config)#sflow null0-sampling
```

To disable null0 sampling, enter the following command.

```
Brocade(config)#no sflow null0-sampling
```

Syntax: [no] sflow null0-sampling

Configuring a null0 route

For configuring a route for null0 sampling, use the following command.

```
Brocade(config)#ip route 10.10.10.100/32 null0
```

Syntax: [no] [ip|ipv6] route *ip-addr* null0

Displaying sFlow show command

This command will display the configuration for sFlow.

```
Brocade(config)#show sflow
```

```
sFlow services are enabled.
sFlow management VRF is enabled.
sFlow management VRF name is default-vrf.
sFlow agent IP address: 55.55.55.56
sFlow agent IPV6 address: unspecified
sFlow source IP address: unspecified, UDP 8888
sFlow source IPv6 address: unspecified, UDP 8888
Collector IP 77.7.7.2, UDP 6343
Polling interval is 20 seconds.
Configured default sampling rate: 1 per 2048 packets.
0 UDP packets exported
124 sFlow samples collected.
133 sFlow management-vrf UDP packets dropped
0 ACL sFlow samples collected.
sFlow ports      Global Sample Rate   Port Sample Rate   Hardware Sample Rate
          1/5                2048                2048   port_down
          1/8                2048                2048   2048
sFlow Null-0 Sampling is Enabled.
```

1

Support matrix for MSTP and STP global/STP/RSTP

Support matrix for MSTP and STP global/STP/RSTP

Table 7 provides the MSTP compatibility matrix as of NetIron 5.4.00b.

TABLE 7 MSTP compatibility matrix

xSTP Protocol	Global STP	Single STP	Single RSTP	Per Vlan STP	Per Vlan RSTP
MSTP	Not supported	Not supported	Not supported	Supported	Supported

Aggregated TM VOQ statistics collection

The following is an update to the list of modules that support Aggregated TM VOQ statistics collection.

Supported modules

Traffic Manager queue statistics are only reported on the following interface modules:

- BR-MLX-10Gx8-X, NI-MLX-10Gx8-M, and NI-MLX-10Gx8-D
- BR-MLX-100Gx2-X and BR-MLX-100Gx1-X
- NI-MLX-48-T-A
- BR-MLX-24x1GF-X-ML, BR-MLX-24x1GC-X-ML, BR-MLX-24x1GF-X, and BR-MLX-24x1GC-X
- **BR-MLX-10Gx24-DM** (Added in NetIron 5.4.00b)

NOTE

The following modules are not supported NI-X-OC192x2, NI-X-OC48x8, NI-X-OC48x4, and NI-X-OC48x2.

Displaying TM statistics from one queue or all queues

Use the following command to display traffic manager statistics for ethernet.

```
Brocade# show tm-voq-stat src_port eth 2/1 dst_port ethernet
-----ethernet 2/2 - 1/4-----
EnQue Pkt Count          4168645330
  EnQue Bytes Count      1010575722
  DeQue Pkt Count         0
  DeQue Bytes Count       0
  Total Discard Pkt Count 2084322665
  Total Discard Bytes Count 505287857
  Oldest Discard Pkt Count 0
  Oldest Discard Bytes Count 0
  WRED Dropped Pkt Count 1594822490
  WRED Dropped Bytes Count 126321962
  Current Queue Depth    0
  Maximum Queue Depth since Last read 0
```

Use the following command to display traffic manager statistics for all priorities.

```
Brocade# show tm-voq-stat src_port p1/1 dst_port p1/2
----- Ports 1/1 - 1/4 -----
Priority = 0
  EnQue Pkt Count          81581531
  EnQue Bytes Count        2692190523
  DeQue Pkt Count          81581531
  DeQue Bytes Count        2692190523
  Total Discard Pkt Count   0
  Total Discard Bytes Count 0
  Oldest Discard Pkt Count 0
  Oldest Discard Bytes Count 0
  WRED Dropped Pkt Count   0
  WRED Dropped Bytes Count 0
```

Aggregated TM VOQ statistics collection

```

    Current Queue Depth                0
    Maximum Queue Depth since Last read 2310
Priority = 1
    EnQue Pkt Count                    0
    EnQue Bytes Count                  0
    DeQue Pkt Count                     62
    DeQue Bytes Count                  1302
    Total Discard Pkt Count             0
    Total Discard Bytes Count           0
    Oldest Discard Pkt Count            0
    Oldest Discard Bytes Count          0
    WRED Dropped Pkt Count              0
    WRED Dropped Bytes Count            21
    Current Queue Depth                 0
    Maximum Queue Depth since Last read 0
Priority = 2
....

```

Syntax: `show tm-voq-stat src_port source-port dst_port ethernet destination-port priority`

Specification of a **source-port** and **destination-port** is required.

You can optionally specify a **priority** to limit the display to a single priority.

The output from the TM Q statistics is available only if the src card type is a module listed in the supported modules list

You can optionally specify a **priority** to limit the display to a single priority or use the **all** parameter to display all priorities.

TABLE 8 Traffic Manager statistics

This field...	Displays...
EnQue Pkt Count	A count of all packets entering ingress queues on this traffic manager.
EnQue Byte Count	A count of all bytes entering ingress queues on this traffic manager.
DeQue Pkt Count	A count of all packets dequeued from ingress queues and forwarded on this traffic manager.
DeQue Byte Count	A count of all bytes dequeued from ingress queues and forwarded on this traffic manager.
TotalQue Discard Pkt Count	A count of all packets failing to enter ingress queues on this traffic manager. This may be due to: <ul style="list-style-type: none"> the queue reaching its maximum depth, WRED, or other reasons. the network processor deciding to drop packets for reasons including: an unknown Layer-3 route, RPF, or segment filtering.
TotalQue Discard Byte Count	A count of all bytes failing to enter ingress queues on this traffic manager. This may be due to: <ul style="list-style-type: none"> the queue reaching its maximum depth, WRED, or other reasons. the network processor deciding to drop packets for reasons including: an unknown Layer-3 route, RPF, or segment filtering.
Oldest Discard Pkt Count	A count of all packets entering ingress queues on this traffic manager, but deleted afterwards due to buffer full.
Oldest Discard Byte Count	A count of all bytes entering ingress queues on this traffic manager, but deleted afterwards due to buffer full.

TABLE 8 Traffic Manager statistics (Continued)

This field...	Displays...
WRED Dropped Pkt Count	A count of all packets entering ingress queues on this traffic manager but dropped due to WRED.
WRED Dropped Bytes Count	A count of all bytes entering ingress queues on this traffic manager but dropped due to WRED.
Maximum Queue Depth since Last read	The maximum queue depth since last access to read.

Displaying TM statistics from the multicast queue

Use the following command to display traffic manager statistics from the Multicast queue for priority 1 on a module.

```
Brocade# show tm-voq-stat src_port eth 4/1 multicast 1
Priority = 0/1
  EnQue Pkt Count                0
  EnQue Bytes Count              0
  DeQue Pkt Count                0
  DeQue Bytes Count              0
  Total Discard Pkt Count        0
  Total Discard Bytes Count      0
  Oldest Discard Pkt Count       0
  Oldest Discard Bytes Count     0
  WRED Dropped Pkt Count         0
  WRED Dropped Bytes Count       0
  Current Queue Depth            0
  Maximum Queue Depth since Last read 0
```

Syntax: `show tm-voq-stat src_port source-port multicast priority | all`

Specification of a **source-port** is required.

You can optionally specify a **priority** to limit the display to a single priority or use the **all** parameter to display all priorities.

Displaying QoS packet and byte counters

You can enable the collection of statistics for Ingress and Egress packet priorities using the **enable-qos-statistics** command. Once the collection of statistics is enabled, the **show np statistics** command can be used to display a count of the packet priorities of Ingress and Egress packets as shown in the following.

```

Brocade# show np statistics
TD: Traffic Despritor. Each TD has size of 512 Bytes

MODULE # 0 PPCR # 0 :
Ingress Counters :
Received packets                               = 5172
Discarded packets                             = 0
Received TDs on traffic class 0                = 0
Received TDs on traffic class 1                = 0
Received TDs on traffic class 2                = 0
Received TDs on traffic class 3                = 0
Received TDs on traffic class 4                = 0
Received TDs on traffic class 5                = 0
Received TDs on traffic class 6                = 0
Received TDs on traffic class 7                = 10344

Egress Counters :
Transmitted unicast packets                    = 0
Transmitted multicast packets                  = 0
Transmitted broadcast packets                  = 0
Filtered packets due to VLAN spanning tree    = 0
Tail dropped packets                           = 0
Control packets                               = 10344
Packets filtered due to egress forward restrictions = 0
Packets dropped due to full multicast egress queue = 91459

TD: Traffic Despritor. Each TD has size of 512 Bytes

MODULE # 1 PPCR # 0 :
Ingress Counters :
Received packets                               = 47809289718
Discarded packets                             = 0
Received TDs on traffic class 0                = 47809289569
Received TDs on traffic class 1                = 0
Received TDs on traffic class 2                = 0
Received TDs on traffic class 3                = 0
Received TDs on traffic class 4                = 0
Received TDs on traffic class 5                = 0
Received TDs on traffic class 6                = 0
Received TDs on traffic class 7                = 0

Egress Counters :
Transmitted unicast packets                    = 18561287821
Transmitted multicast packets                  = 0
Transmitted broadcast packets                  = 0
Filtered packets due to VLAN spanning tree    = 0
Tail dropped packets                           = 5910551222
Control packets                               = 0
Packets filtered due to egress forward restrictions = 0
Packets dropped due to full multicast egress queue = 0

```

Multi-Chassis Trunk (MCT) client-interfaces delay

Use the `client-interfaces delay` command to set the delay before bringing up the CCEP port. This command is used to set the delay, so that after a node is reloaded, with just L2vpn peer alone, the delay to bring up the CCEP port will be the designated value.

```
Brocade(config-cluster-TOR)#client-interfaces delay 60
```

Syntax: `[no] client-interfaces delay` *time in sec*

The default value for delay is 30 seconds. The acceptable values range between 20 to 600 seconds.

NOTE

Client-interface delay is only applied with just L2 VPN. It does not support L2+L2VPN.

Deletion of ACLs bound to an interface

TABLE 9 Supported platforms

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series BASE package	Brocade Netron CER 2000 Series Advanced Services package
Deletion of ACLs bound to an interface	Yes	Yes	No	No	No	No	No

To delete an ACL bound to an interface, use the **force-delete-bound-acl** command. The **no** form of **force-delete-bound-acl** command does not allow the deletion of a bound ACL. Initially **force-delete-bound-acl** is disabled.

```
Brocade(config)#acl-policy
```

The **force-delete-bound-acl** command allows the ACLs bound to an interface to be deleted.

```
Brocade(config-acl-policy)# force-delete-bound-acl
```

The **no force-delete-bound-acl** command does not allow the ACLs bound to an interface to be deleted.

```
Brocade(config-acl-policy)# no force-delete-bound-acl
```

Syntax: [no] force-delete-bound-acl

When **force-delete-bound-acl** is enabled, it allows deletion of ACLs bound to one or more interfaces. After enabling this command for the deletion of the ACLs, however the binding of the ACL to an interface still remains. On rebinding this will be an empty ACL and will have no affect on traffic forwarding. On rebinding the CAM entries are reprogrammed appropriately, so no ACL filtering takes place after the ACL is deleted. This command is available as a sub-command of **acl-policy** command. However like any other ACL modification the CAM is only reprogrammed during **rebind**. Without a **rebind** the old filters are still present in the CAM.

NOTE

When this command is enabled, an ACL can be deleted even if it is bound to one or more interfaces. However, the interface binding to the ACL remains. This will be an empty ACL and will have no affect on traffic forwarding. In case of subnet broadcast ACL bindings, when an empty ACL is bound to an interface, implicit deny entries are programmed to the CAM and will have effect on traffic forwarding.

An example of the command is as below.

```
Brocade(config-acl-policy)# force-delete-bound-acl
Brocade(config-acl-policy)# exit
Brocade(config)# show access-list all
ACL configuration:
!
mac access-list SampleACL
  permit any any 10 etype any
!
Brocade(config)# show access-list bindings
L4 configuration:
!
```

Deletion of ACLs bound to an interface

1

```
interface ethe 2/1
  mac access-group SampleACL in
!
Brocade(config)#show cam l2acl
  SLOT/PORT  Interface number
Brocade(config)# sh cam l2acl 2/1
LP Index  VLAN Src MAC          Dest MAC          Port  Action  PRAM
  (Hex)
2 0a3800 10 0000.0000.0000 0000.0000.0000 0    Pass   0009c
2 0a3802 0 0000.0000.0000 0000.0000.0000 0    Drop   0009d
Brocade(config)#
Brocade(config)#no mac acc SampleACL
Brocade(config)#sh cam l2acl 2/1
LP Index  VLAN Src MAC          Dest MAC          Port  Action  PRAM
  (Hex)
Brocade(config)#show access-list all ACL configuration:
!
Brocade(config)#show access-list bindings
L4 configuration:
!
!
interface ethe 2/1 mac access-group SampleACL in
!
Brocade(config)#
```

NOTE

Rebinding of an ACL is explicitly required for IPv4 and IPv6 ACLs.

Configuring an encrypted syslog server

You can configure up to six encrypted syslog servers, but only one is active at any time, with the other servers acting as standby. When you add an encrypted syslog server, if there is no active syslog server, a session is established with the configured server. If a new connection is added when an active session exists, a new session with another encrypted syslog server is not attempted.

A new syslog server session is attempted in the following scenarios:

- Current active encrypted syslog server configuration is removed or the SSL connection to the active syslog server is closed
- During a device reload
- During switch over of the management module
- No active syslog server is found when the device sends syslog messages

Attempts to connect to a new syslog server starts with the first configured syslog server. The device attempts to establish an SSL connection with a server until a successful SSL connection is established. During this interval, the trap hold down timer is started and all the syslog messages are queued. When the timer expires, the device sends queued log messages to the connected syslog server.

Configuring encrypted syslog servers requires two steps:

- Installing the SSL Client certificate from a remote machine
- Adding encrypted syslog servers

Installing the SSL client certificate

Before you can configure an encrypted syslog server for the device, you must install the SSL client certificate. Do one of the following to install the SSL client certificate.

Using TFTP:

1. Use TFTP to copy the SSL Client Certificate and private key from the remote machine if TFTP is enabled on the device. Enter the following commands in sequence in any order:

```
Brocade# copy tftp flash 10.25.101.121 cert.p12 client-certificate
Brocade# copy tftp flash 10.25.101.121 privkeyfile client-private-key
```

Syntax: copy tftp flash <remote_ip> <cert_file> client-certificate

and

Syntax: copy tftp flash <remote_ip> <priv_key_file> client-private-key

The **remote_ip** keyword specifies the IP address of the remote host where the SSL Client certificate and private key are present. The **cert_file** keyword specifies the filename of the SSL Client Certificate, and the **priv_key_file** keyword specifies the filename of the private key.

Using SCP

1. Use SCP to copy the SSL Client Certificate and private key from the remote machine. Enter the following commands in sequence in any order at the remote host where the SSL Client Certificate and private key are present:

```
Host# scp cert.p12 user@10.25.105.121:sslclientcert
Host# scp privkeyfile user@10.25.105.121:sslclientprivkey
```

Syntax: scp <cert_file> user@<remote_ip>:sslclientcert

and

Syntax: scp <priv_key_file> user@<remote_ip>:sslclientprivkey

The **remote_ip** keyword specifies the IP address of the device. The **cert_file** keyword specifies the filename of the SSL Client Certificate, and the **priv_key_file** keyword specifies the filename of the private key.

Adding an encrypted syslog server

To configure an encrypted server connection, enter the following command:

```
Brocade (config)# logging host 10.25.105.201 ssl-port 60514
```

Syntax: logging host [ipv6] <ip-address> | <ipv6-address> ssl-port <port>

The **ip-address** keyword specifies the syslog server. The **ssl-port** keyword specifies the SSL port that will be used to connect to the specified syslog server.

NOTE

You can configure an encrypted syslog server connection only after the device has been placed in the Common Criteria mode. While you can configure these when the device is in the Administrative mode, the configuration takes effect only after the device is put in the Common Criteria Operational mode.

Displaying the configured server connections

You can display the active encrypted syslog server connection with the **show ip ssl** command:

```
Brocade# show ip ssl
Session Source IP      Source Port  Remote IP    Remote Port
0        10.25.105.80    633         10.25.105.201 60514
```

In addition, you can use the show logging command to display the active SSL-encrypted syslog server along with the logging level information.

```
Brocade# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 27 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning
Current active SSL syslog server: 10.25.105.201:60514
```

Global ACL command to delete ACLs bound to an interface

The access-list command now allows the deletion of Access Control Lists (ACLs) bound to an interface. The following examples show the application of an ACL and the deletion of the ACL that has been bound to an interface.

1. ACL configuration

```
Brocade(config)# access-list 102 permit ip any any
```

2. Application of the ACL to interfaces

```
Brocade(config)# int eth 1/2
Brocade(config-if-e10000-1/2)# ip access-group 102 in
Brocade(config-if-e10000-1/2)# exit

Brocade(config)# int eth 4/3
Brocade(config-if-e10000-4/3)# ip access-group 102 out
Brocade(config-if-e10000-1/2)# exit
```

3. Deleting an ACL definition

```
Brocade(config)# no access-list 102 permit ip any any
Brocade(config)# ACL 102 is in use. Would you like to delete it? (Y/N)
```

Changing the router ID

In most configurations, a Brocade device has multiple IP addresses, usually configured on different interfaces. As a result, a Brocade device's identity to other devices varies depending on the interface to which the other device is attached. Some routing protocols, including OSPF and BGP4, identify a Brocade device by just one of the IP addresses configured on the Brocade device, regardless of the interfaces that connect the Brocade devices. This IP address is the router ID.

NOTE

RIP does not use the router ID.

NOTE

If you change the router ID, all current BGP4 sessions are cleared.

By default, the router ID on a Brocade device is one of the following:

- If the device has loopback interfaces, the default router ID is the IP address configured on the lowest numbered loopback interface configured on the Brocade device. For example, if you configure loopback interfaces 1, 2, and 3 as follows, the default router ID is 10.9.9.9/24:
 - Loopback interface 1, 10.9.9.9/24
 - Loopback interface 2, 10.4.4.4/24
 - Loopback interface 3, 10.1.1.1/24
- If the IP address from loopback1 interface (lowest numbered loopback interface) is removed, the next lowest loopback interface IP address is selected as router-id.
- If a loopback interface is not configured, then the lowest IP address configured over the physical interface is selected as the router ID.

If you prefer, you can explicitly set the router ID to any valid IP address. The IP address should not be in use on another device in the network.

You can set a router ID for a specific VRF as described within this section. In order to make the route ID calculation more deterministic, the device calculates the router-id value during bootup and does not calculate or change the router-id value unless the IP address used for the router-id value on the device is deleted, or the **clear router-id** command is issued. Additionally, setting a router-id value overrides the existing router-id value and takes effect immediately. Once a router-id value set by a user is removed using the **no ip router-id x.x.x.x** command, the device will again recalculate the router-id value based on current information.

NOTE

The Brocade device uses the same router ID for both OSPF and BGP4. If the device is already configured for OSPF, you may want to use the router ID that is already in use on the device rather than set a new one. To display the router ID, enter the **show ip** command at any CLI level.

To change the router ID, enter a command such as the following.

```
Brocade(config)# ip router-id 10.157.22.26
```

Syntax: **[no] ip router-id ip-addr**

The *ip-addr* can be any valid, unique IP address.

To set the router ID within a VRF, enter a command such as the following.

```
Brocade(config)# vrf blue
Brocade(config-vrf-blue)# ip router-id 10.157.22.26
```

Syntax: **[no] ip router-id ip-addr**

NOTE

The command for setting the router ID for a specified VRF is exactly the same as for the default VRF. The only difference is that when setting it for a specific VRF, the **ip router-id** command is configured within the VRF as shown in the example.

NOTE

You can specify an IP address used for an interface, but do not specify an IP address in use by another device.

Show lag

The show lag command has been updated to display the number of available LAGs, including 100g LAGs.

```

Brocade #
Brocade #show lag b
Total number of LAGs      : 2, 100g : 2
Total number of deployed LAGs : 2, 100g : 2
Total number of trunks created : 2 (254 total available), 100g : 2 (14 total
available)
LACP System Priority / ID      :1 / 0024.3883.3600
LACP Long timeout             :90, default: 90
LACP Short timeout           :3, default: 3

LAG          Type      Deploy Trunk Primary      Port List
100g_lag    static     Y      1      3/1          e 3/1
10g_lag     static     Y      2      2/1          e 2/1
1g_lag      static     Y      3      1/21         e 1/21
lag2        dynamic    Y      4      3/2          e 3/2

```

NOTE

The update is to the output only.

OpenFlow Hybrid Port Mode for IPv6

This feature enables an OpenFlow enabled port to support normal IPv6 routing on protected VLANs.

Bypass LSP Liberal Path Selection

TABLE 10 Supported platforms

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series Base package	Brocade Netron CER 2000 Series Advanced Services package
Liberal bypass LSP selection	Yes	Yes	No		No	No	Yes

Current algorithm

The current algorithm to select a Bypass LSP is very conservative. The restrictive algorithm can run into situations where no backup path can be established. This can occur when bypass LSPs cannot qualify under those restrictions, for example, under a less meshed topology such as single ring topology.

The following process explains the steps of facility backup path computation, which involves selecting the best qualified bypass LSP.

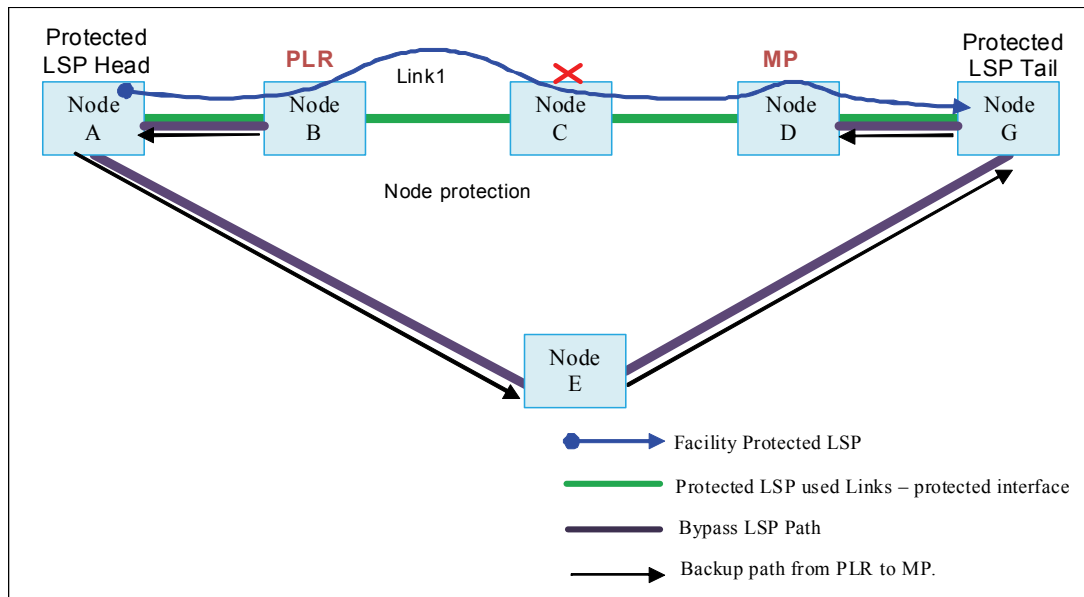
- **Merge point selection:** The PLR backup query process first selects in the order of preferred merge point based on ingress signaled property. A merge point is selected in the order of preference, from available bypass LSPs reaching this merge point. If no bypass LSPs qualify to serve, move on to next preferred merge point. The merge point preference order depends on the ingress signal with the node protection desired flag.
 - a. If node protection desired flag is present, PLR will go through merge point in the order of next-next-hop (if present, to achieve node protection), next-hop (link protection), hops after next-next-hop in sequence of traverse if any present.
 - b. If node protection desired flag is not present, the next-hop as the only merge point is selected.
- **Bypass LSP qualification:** Bypass LSP cannot traverse any link attached to nodes traversed by protected session between PLR and egress of LSP.
- If there is more than one bypass LSP qualified to serve for backup path, the lowest LSP cost metric is considered. If more than one bypass LSP is available with lowest cost, the one with the lowest number of riding backup sessions is selected.

For example, in [Figure 3](#) on page 37, the FRR LSP path is: “A” to “B” to “C” to “D” to “G”. A Bypass LSP is created at router B and the new path is: “B” to “E” to “G” to “D”.

In this example, the PLR is node B and the merge point (MP) is node D.

As per the current algorithm, the Bypass LSP traverses the link attached to LSP's egress node G. Therefore, the Bypass LSP does not qualify to provide protection for the FRR LSP per the current restrictive mode of the Bypass LSP selection/qualification.

FIGURE 3 Bypass LSP



New algorithm

The new facility backup computation mode applies an algorithm between an extremely conservative approach and an extremely liberal approach. The changes only affect how a bypass LSP is qualified. There are no changes regarding the merge point selection or capability to disable providing node protection or tie breaker from multiple qualified bypass LSPs.

For a bypass LSP to qualify for the backup path, it must pass following tests:

- Bypass LSP cannot traverse any nodes between PLR and merge point if there are any

Using the example above (Figure 3 on page 37), the new algorithm mandates that the Bypass LSP cannot traverse node C.

Therefore, with the new Bypass LSP selection mode turned on, the above Bypass LSP qualifies and can be used to provide protection for the FRR LSP.

Enabling bypass LSP liberal path selection

The **cspf-computation-mode** command enables or disables backup query algorithm using minimum restrictions to qualify the bypass LSP. This command can be executed on the fly and bypass LSP selection process will use the restricted or liberal mode, depending upon the current configuration. Changing the computation mode on the fly will not impact the already selected bypass LSPs. This configuration is equally applicable to dynamic bypass LSP selection as well.

To enable the liberal bypass LSP selection, enter the following commands.

```
Brocade(config)# router mpls
Brocade(config-mpls)#policy
Brocade(config-mpls-policy)#cspf-computation-mode use-bypass-liberal
```

Syntax: [no] **cspf-computation-mode** [**use-bypass-liberal** | use-bypass-metric]

By default, backup query uses full restrictions to qualify bypass LSP during backup query.

Bypass LSP Liberal Path Selection

This command can be executed at any time. The bypass LSP selection process will use the restricted or liberal mode depending upon the current configuration. Changing the computation mode will not impact the already selected bypass LSPs.

Show command enhancements

The configuration which enables liberal mode will be displayed as part of the following commands. Note that by default, this option is disabled and not shown. This is different from the command of `ospf-computation-mode use-bypass-metric`, which will show disabled when not enabled.

Show mpls config

```
Brocade(config-mpls-policy)#show mpls config
router mpls
  policy
    cspf-computation-mode use-bypass-metric
    cspf-computation-mode use-bypass-liberal
```

Show mpls policy

```
Brocade(config-mpls-policy)#show mpls policy
Current MPLS policy settings:
  CSPP interface constraint: enabled
  CSPP-Group computation-mode: disabled
  CSPP computation-mode: Use bypass metric
  CSPP computation-mode: Use bypass liberal
  TTL propagation for MPLS label: disabled, IPVPN: disabled, IP over MPLS: enabled
  Inter-AS route filtering: enabled, Intra-AS iBGP route filtering: disabled
  Ingress tunnel accounting: enabled
  Polling interval for MPLS LSP traffic statistics: 300 seconds
  Advertise TE parameters via: ISIS level-2
  Handle neighbor down event - ISIS: Yes OSPF: No
  LSP rapid retry: disabled, maximum number of retries: no limit
  LSP periodic retry time: 30 seconds
  FRR backup/detour retry time: 30 seconds
  Admin group:
    blue, group number: 1
    yellow, group number: 2
    red, group number: 6
  green, group number: 8
```

Show mpls lsp

```
Brocade# show mpls lsp frr_lsp
LSP frr_lsp, to 7.7.7.2
  From: 7.7.7.1, admin: UP, status: UP, tunnel interface(primary path): tn10
  Times primary LSP goes up since enabled: 1
  Metric: 0, Adaptive
  ...
  Recorded routes:
    Protection codes: P: Local N: Node B: Bandwidth I: InUse
    7.1.1.0 (P) -> 7.1.13.1
  Fast Reroute: facility backup desired
  Backup LSP: UP, out-label: 3, outbound interface: e4/9 bypass_lsp: byp1
  Path cspf-group computation-mode: add-penalty, cspf-computation-mode: use-
  bypass-metric, use-bypass-liberal, cost: 0
  Global revertiveness enabled with hold time 5 secs
  FRR Forwarding State: Pri(active), Backup(up)
```


Max Queue Depth and Buffer Utilization CLI enhancements

This section describes two commands which summarize the buffer utilization and maximum queue depth across all queues on a per slot basis.

TABLE 11 Supported platforms

Features supported	Brocade Netron XMR Series	Brocade MLX Series and Brocade MLXe series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series Base package	Brocade Netron CER 2000 Series Advanced Services package
Max queue depth and buffer utilization enhancements	Yes	Yes	No	No	No	No	No

TABLE 12 Interface card support

Features supported	NI-X-100 Gx2	BR-MLX-10Gx24	BR-MLX-40Gx4-X	NI-MLX-10 Gx8-M NI-MLX-10 Gx8-X	NI-MLX-10G x8-M BR-MLX-10 Gx8-X	NI-XMR-10Gx4 NI-MLX-10Gx4 BR-MLX-10Gx4-X NI-XMR-10Gx20-GC NI-XMR-10Gx20-SFP NI-MLX-10Gx20-GC NI-MLX-10Gx20-SFP	BR-MLX-10Gx24-X BR-MLX-10Gx24-X NI-MLX-10Gx48-T NI-MLX-10Gx48-T-A
Max queue depth and buffer utilization enhancements	Yes	Yes	No	Yes	Yes	No	Yes on max-queue-depth CLI. No on max buffer utilization CLI.

Displaying Traffic Manager max queue depth summary

Use the following command to display the traffic manager (TM) max queue depth summary from specified ports on a module.

```
Brocade# show tm-voq-stat max-queue-depth slot 3
```

```
----- Ports 3/1 - 3/24 -----
QType   Max Depth   Max Util   Destination Port
0       1013804     96%       3/1
1       1013848     96%       3/1
2       1013666     96%       3/4
3       1013794     96%       3/1
4       1013564     96%       3/1
5       538         0%        2/7
```

1

Max Queue Depth and Buffer Utilization CLI enhancements

```
6          532          0%          2/7
7          0          0%          NA
----- Ports 3/25 - 3/48 -----
QType    Max Depth    Max Util    Destination Port
0         0          0%         NA
1         0          0%         NA
2         0          0%         NA
3         0          0%         NA
4         0          0%         NA
5         0          0%         NA
6         0          0%         NA
7         0          0%         NA
```

TABLE 13

Field	Explanation
QType	Queue priority
Max Depth	Maximum queue depth of any queue with Qtype in bytes
Max Util	Percentage of max queue util (max-queue-depth / max-queue-size)
Destination Port	Destination port of queue that had highest max queue depth

Syntax: `<show | clear> tm-voq-stat max-queue-depth slot <slotnum>`

You can clear the max queue depth report using the **clear** command.

Displaying Traffic Manager maximum buffer utilization

Use the following command to monitor Traffic Manager maximum buffer utilization.

```
Brocade# show tm buffer-pool-stats slot 3
----- Ports 3/1 - 3/4 -----
Maximum Buffer Size:          0 (0%)
Maximum Occupied Buffer Descriptors:  0 (0%)

----- Ports 3/5 - 3/8 -----
Maximum Buffer Size:          0 (0%)
Maximum Occupied Buffer Descriptors:  0 (0%)
```

TABLE 14

Field	Explanation
Maximum Buffer Size	High watermark of buffer size in bytes (for both Gold and Bronze traffic) since last read. Also shows percentage of buffer used out of max packet buffer. Clear on read.
Maximum Occupied Buffer Descriptors	High watermark of descriptors (buffer pointers) used (for both Gold and Bronze traffic) since last read. Also shows percentage of descriptors used out of total descriptors. Clear on read.

Syntax: `show tm buffer-pool-stats slot <slotnum>`

You can optionally clear the buffer pool statistics report using the **clear tm buffer-pool-stats** command.

Transparent forwarding of L2 and L3 protocols on a VLL for CES and CER

TABLE 15 Feature support Table

Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_Prem package	Brocade Netron CER 2000 Series BASE package	Brocade Netron CER 2000 Series Advanced Services package
N	N	Y	Y	Y	Y	Y

The command **forward-all-control** has been implemented in NetIron 5.4.00c. This command adds per port Layer 2 and Layer 3 (L2/L3) protocols ACL filters for the VLL end point port. The command **no forward-all-control** will remove the L2/L3 protocols ACL filters for the VLL end point port.

NOTE

The **forward-all-control** command is only applicable to the Brocade NetIron CER and Brocade NetIron CES.

To implement per port Layer 2 and Layer 3 (L2/L3) protocols ACL filters, enter the following command.

```
Brocade(config)# int eth 1/1
Brocade (config-if-e1000-1/1)# forward-all-control
```

Syntax: [no] **forward-all-control**

The command **no forward-all-control** will delete VLL end point port L2/L3 protocols ACL filters. For LAG, only the primary port needs to be configured.

NOTE

The **forward-all-control** command will let L2/L3 protocols on the port go with hardware forwarding without going to the CPU. If the **no forward-all-control** command is executed, the L2/L3 functions may be impacted.

The **show interfaces ethernet slot/port** command will display the configuration status of the **forward-all-control** command.

The following output example shows the **show interfaces ethernet slot/port** command with the **forward-all-control** command disabled.

```
Brocade#show interfaces ethernet 1/1
GigabitEthernet1/1 is up, line protocol is up
  STP Root Guard is disabled, STP BPDU Guard is disabled
  Hardware is GigabitEthernet, address is 001b.eda3.f841 (bia 001b.eda3.f841)
  Configured speed auto, actual 1Gbit, configured duplex fdx, actual fdx
```

1 Transparent forwarding of L2 and L3 protocols on a VLL for CES and CER

```
Member of 1 L2 VLAN(S) (tagged), port is in tagged mode, port state is
Forwarding
STP configured to ON, Priority is level0, flow control enabled
Priority force disabled, Drop precedence level 0, Drop precedence force disabled
dhcp-snooping-trust configured to OFF
mirror disabled, monitor disabled
LACP BPDU Forwarding:Disabled
LLDP BPDU Forwarding:Disabled
L2L3 protocols Forwarding:Disabled
Not member of any active trunks
```

...

The following output example shows the **show interfaces ethernet slot/port** command with the **forward-all-control** command enabled.

```
Brocade(config-if-e1000-1/1)#forward-all-protocol
Brocade(config-if-e1000-1/1)#show interfaces ethernet 1/1
GigabitEthernet1/1 is up, line protocol is up
STP Root Guard is disabled, STP BPDU Guard is disabled
Hardware is GigabitEthernet, address is 001b.eda3.f841 (bia 001b.eda3.f841)
Configured speed auto, actual 1Gbit, configured duplex fdx, actual fdx
Member of 1 L2 VLAN(S) (tagged), port is in tagged mode, port state is
Forwarding
STP configured to ON, Priority is level0, flow control enabled
Priority force disabled, Drop precedence level 0, Drop precedence force disabled
dhcp-snooping-trust configured to OFF
mirror disabled, monitor disabled
LACP BPDU Forwarding:Disabled
LLDP BPDU Forwarding:Disabled
L2L3 protocols Forwarding:Enabled
Not member of any active trunks
```

...

Forward Error Correction mode

TABLE 16 Feature support Table

Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_Prem package	Brocade Netron CER 2000 Series BASE package	Brocade Netron CER 2000 Series Advanced Services package
Y	Y	N	N	N	N	N

Using Forward Error Correction (FEC) mode enabled modules on a Brocade MLXe series chassis will reduce packet drops due to CRC errors. FEC will automatically be enabled on supported line cards and fabric links in a Brocade MLXe series chassis.

Forward Error Correction (FEC) mode is applicable for the Brocade MLXe series platforms. It will be operational only on the 32Ke chassis for the following cards:

- 2x100G
- 24x10G
- hSFMs (FE600 based SFMs)

FEC mode is applied on a per link basis. Both sides of the link (TM side and FE side) must be in the same mode. In a Brocade MLXe series chassis, the following applies:

- All fabric facing links on the 2x100 and 24x10 TMs will have FEC enabled
- hSFM links connected to 2x100 and 24x10 will have FEC enabled

Manual deletion of an OpenFlow rule

TABLE 17 Feature support Table

Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_Prem package	Brocade Netron CER 2000 Series BASE package	Brocade Netron CER 2000 Series Advanced Services package
Y	Y	Y	Y	Y	Y	Y

The command **clear openflow flowid** *Flow ID* has been implemented in Netron 5.4.00c. This command adds an enhancement to delete a single OpenFlow rule based on a Flow ID.

```
Brocade# clear openflow flowid 6
```

Syntax: **clear openflow flowid** [Flow ID]

The command will delete a single OpenFlow rule with the given [Flow ID]. The command will delete the rule irrespective of the state it is in (ACTIVE, PENDING_ADD, PENDING_MODIFY, PENDING_DELETE). The same rule can be added again later from the controller if needed. However, the flow-id of the deleted rule cannot be reused again.

NOTE

The **clear openflow** command cannot be used simultaneously when there is deletion from a controller already in progress. The **clear** command will exit with the following message: "Deletion from Controller in progress..Try again later !"

Show tech enhancement for OpenFlow

The **show openflow tech-support** command has changed to **show tech-support openflow**. The changes ensure that the OpenFlow feature is in line with all other feature commands.

```
Brocade# show tech-support openflow.
```

Syntax: **show tech-support openflow**

This command will now capture the output of the following commands:

- show openflow datapath-id
- show openflow controller
- show openflow interface
- show openflow flows
- show versions
- show interfaces
- show statistics
- show running-config
- show logging
- show save

Root Guard

NOTE

This enhancement is to synchronize the “root protect CCEP” ports states to the peer MCT.

In NetIron 05.4.00c, a new security feature has been added that allows a CCEP port to run STP, but not allow the connected device to become the Root. The Root Guard feature provides a way to enforce the root bridge placement in the network and allows STP to interoperate with user network bridges while still maintaining the bridged network topology that the administrator requires. Errors are triggered if any change from the root bridge placement is detected.

NOTE

The feature is also available for RSTP.

When Root Guard is enabled on a port, it keeps the port in designated FORWARDING state. If the port receives a superior BPDU, which is a Root Guard violation, it sets the port into BLOCKING state and triggers a Syslog message and an SNMP trap. No further traffic will be forwarded on this port. This allows the bridge to prevent traffic from being forwarded on ports connected to rogue or misconfigured STP or RSTP bridges.

NOTE

Root protect should be configured on CCEP ports of both the peers to sync the state properly.

Root Guard should be configured on all ports where the root bridge should not appear. In this way, the core bridged network can be cut off from the user network by establishing a protective perimeter around it.

Once the port stops receiving superior BPDUs, Root Guard will automatically set the port back to a FORWARDING state after the timeout period has expired.

NOTE

Root Guard may prevent network connectivity if improperly configured. It needs to be configured on the perimeter of the network rather than the core. Also, Root Guard should be configured only on the primary port of a LAG. If a port configured with Root Guard is made a secondary port, the LAG deployment will be vetoed.

Enabling Root Guard

Root Guard is configured on a per interfaces basis. To enable Root Guard, enter a command such as the following.

```
Brocade(config)# interface ethernet 5/5
Brocade(config-if-e10000-5/5) spanning-tree root-protect
```

Syntax: [no] spanning-tree root-protect

Enter the **no** form of the command to disable Root Guard on the port.

Refer to the *Root Guard* section of the NetIron 5.4.00 Configuration Guide for information and procedures including:

- Setting the Root Guard timeout period
- Checking if Root Guard is configured
- Displaying the Root Guard state
- Reconfiguring the timeout period
- Checking for Syslog messages

Discontinuing FID updates

When the following command is enabled, FID updates will not be sent to the line card. A new FID will be created on the management card and sent to the line card. This command will create a new FID before breaking the old FID, thereby avoiding traffic loss.

```
Brocade# ip multicast no-fid-update
```

Syntax: ip-multicast no-fid-update

Change the max-response-time value

Use the following command to change the maximum response time value set in the IGMP Group Specific Query and IGMP Group Source Specific message.

NOTE

This command will not change the max-response-time value set in the General Query message.

```
Brocade# ip multicast max-response-time
```

Syntax: ip-multicast max-response-time

Clearing the QoS packet and byte counters

You can clear the QoS counters whose display is generated using the **show np statistics** command as shown in the following.

Syntax: clear np statistics

```
Brocade# clear np statistics
```


IP assignment within a LAG

Layer 3 static or dynamic LAG support IP assignment. All the configurations has to be done on the primary port of the LAG.

The following is a sample configuration:

```
lag lag_dist_a_1 dynamic id 15
ports ethe 1/1 to 1/12
primary-port 1/1
deploy
!
router vrrp
!
interface ethe 1/1
ip address 192.168.10.1 255.255.255.0
ip vrrp vrid 1
backup priority 50 track-priority 10
ip-address 192.168.1.10
activate
```

Update to Chapter 17 of the NetIron 5.4.00a Configuration Guide

Chapter 17 Configuring Quality of Service (QoS) for the Brocade NetIron CES and Brocade NetIron CER Series has been updated. The **show np statistics** command is not applicable on the Brocade NetIron CER and Brocade NetIron CES.

STP feature configuration

TABLE 18 Feature support table

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series BASE package	Brocade Netron CER 2000 Series Advanced Services package
Enhanced support for Fast Port Span, Fast Uplink Span, and Single-instance Span	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Spanning Tree Protocol (STP) features extend the operation of standard STP, enabling you to fine-tune standard STP and avoid some of its limitations.

This section describes how to configure these parameters using the CLI.

Fast port span

When STP is running on a device, message forwarding is delayed during the spanning tree recalculation period following a topology change. The STP forward delay parameter specifies the period of time a bridge waits before forwarding data packets. The forward delay controls the listening and learning periods of STP reconvergence. You can configure the forward delay to a value from 4 – 30 seconds. The default is 15 seconds. Thus, using the standard forward delay, convergence requires 30 seconds (15 seconds for listening and an additional 15 seconds for learning) when the default value is used.

This slow convergence is undesirable and unnecessary in some circumstances. The Fast Port Span feature allows certain ports to enter the forwarding state in four seconds. Specifically, Fast Port Span allows faster convergence on ports that are attached to end stations and thus do not present the potential to cause Layer 2 forwarding loops. Because the end stations cannot cause forwarding loops, they can safely go through the STP state changes (blocking to listening to learning to forwarding) more quickly than is allowed by the standard STP convergence time. Fast Port Span performs the convergence on these ports in four seconds (two seconds for listening and two seconds for learning).

In addition, Fast Port Span enhances overall network performance in the following ways:

- Fast Port Span reduces the number of STP topology change notifications on the network. When an end station attached to a Fast Span port comes up or down, the Brocade device does not generate a topology change notification for the port. In this situation, the notification is unnecessary since a change in the state of the host does not affect the network topology.

- Fast Port Span eliminates unnecessary MAC cache aging that can be caused by topology change notifications. Bridging devices age out the learned MAC addresses in their MAC caches if the addresses are unrefreshed for a given period of time, sometimes called the MAC aging interval. When STP sends a topology change notification, devices that receive the notification use the value of the STP forward delay to quickly age out their MAC caches. For example, if a device normal MAC aging interval is 5 minutes, the aging interval changes temporarily to the value of the forward delay (for example, 15 seconds) in response to an STP topology change. In normal STP, the accelerated cache aging occurs even when a single host goes up or down. Because Fast Port Span does not send a topology change notification when a host on a Fast Port Span port goes up or down, the unnecessary cache aging that can occur in these circumstances under normal STP is eliminated.

Fast Port Span is a system-wide parameter and is enabled by default. Thus, when you boot a device, all the ports that are attached only to end stations run Fast Port Span. For ports that are not eligible for Fast Port Span, such as ports connected to other networking devices, the device automatically uses the normal STP settings. If a port matches any of the following criteria, the port is ineligible for Fast Port Span and uses normal STP instead:

- The port is 802.1Q tagged
- The port is a member of a trunk group
- The port has learned more than one active MAC address
- An STP Configuration BPDU has been received on the port, thus indicating the presence of another bridge on the port.

You also can explicitly exclude individual ports from Fast Port Span if needed. For example, if the only uplink ports for a wiring closet switch are Gbps ports, you can exclude the ports from Fast Port Span.

Disabling and re-enabling fast port span

Fast Port Span is a system-wide parameter and is enabled by default. Therefore, all ports that are eligible for Fast Port Span use it.

To disable or re-enable Fast Port Span, enter the following commands.

```
Brocade(config)#no fast port-span
Brocade(config)#write memory
```

Syntax: [no] fast port-span

NOTE

The **fast port-span** command has additional parameters that let you exclude specific ports. These parameters are shown in the following section.

To re-enable Fast Port Span, enter the following commands.

```
Brocade(config)#fast port-span
Brocade(config)#write memory
```

Excluding specific ports from fast port span

To exclude a port from Fast Port Span while leaving Fast Port Span enabled globally, enter commands such as the following.

```
Brocade(config)#fast port-span exclude ethernet 1
```

STP feature configuration

```
Brocade(config)#write memory
```

To exclude a set of ports from Fast Port Span, enter commands such as the following.

```
Brocade(config)#fast port-span exclude ethernet 1 ethernet 2 ethernet 3
Brocade(config)#write memory
```

To exclude a contiguous (unbroken) range of ports from Fast Span, enter commands such as the following.

```
Brocade(config)#fast port-span exclude ethernet 1 to 24
Brocade(config)#write memory
```

Syntax: [no] fast port-span [exclude ethernet <port> [ethernet <port>] | to [<port>]]

Specify the port variable in one of the following formats:

- FWS, FCX, and ICX stackable switches – stack-unit/slotnum/portnum
- FSX 800 and FSX 1600 chassis devices – slotnum/portnum
- ICX devices – slotnum/portnum
- FESX compact switches – portnum

To re-enable Fast Port Span on a port, enter a command such as the following.

```
Brocade(config)#no fast port-span exclude ethernet 1
Brocade(config)#write memory
```

This command re-enables Fast Port Span on port 1 only and does not re-enable Fast Port Span on other excluded ports. You also can re-enable Fast Port Span on a list or range of ports using the syntax shown above this example.

To re-enable Fast Port Span on all excluded ports, disable and then re-enable Fast Port Span by entering the following commands.

```
Brocade(config)#no fast port-span
Brocade(config)#fast port-span
Brocade(config)#write memory
```

Disabling and then re-enabling Fast Port Span clears the exclude settings and thus enables Fast Port Span on all eligible ports. To make sure Fast Port Span remains enabled on the ports following a system reset, save the configuration changes to the startup-config file after you re-enable Fast Port Span. Otherwise, when the system resets, those ports will again be excluded from Fast Port Span.

Fast Uplink Span

The Fast Port Span feature described in the previous section enhances STP performance for end stations. The Fast Uplink Span feature enhances STP performance for wiring closet switches with redundant uplinks. Using the default value for the standard STP forward delay, convergence following a transition from an active link to a redundant link can take 30 seconds (15 seconds for listening and an additional 15 seconds for learning).

You can use the Fast Uplink Span feature on a Brocade device deployed as a wiring closet switch to decrease the convergence time for the uplink ports to another device to just one second. The new Uplink port directly goes to forward mode (bypassing listening and learning modes). The wiring closet switch must be a Brocade device but the device at the other end of the link can be a Brocade device or another vendor's switch.

Configuration of the Fast Uplink Span feature takes place entirely on the Brocade device. To configure the Fast Uplink Span feature, specify a group of ports that have redundant uplinks on the wiring closet switch (Brocade device). If the active link becomes unavailable, the Fast Uplink Span feature transitions the forwarding to one of the other redundant uplink ports in just one second. All Fast Uplink Span-enabled ports are members of a single Fast Uplink Span group.

NOTE

To avoid the potential for temporary bridging loops, Brocade recommends that you use the Fast Uplink feature only for wiring closet switches (switches at the edge of the network cloud). In addition, enable the feature only on a group of ports intended for redundancy, so that at any given time only one of the ports is expected to be in the forwarding state.

NOTE

When the wiring closet switch (Brocade device) first comes up or when STP is first enabled, the uplink ports still must go through the standard STP state transition without any acceleration. This behavior guards against temporary routing loops as the switch tries to determine the states for all the ports. Fast Uplink Span acceleration applies only when a working uplink becomes unavailable.

Active uplink port failure

The active uplink port is the port elected as the root port using the standard STP rules. All other ports in the group are redundant uplink ports. If an active uplink port becomes unavailable, Fast Uplink Span transitions the forwarding of traffic to one of the redundant ports in the Fast Uplink Span group in one second bypassing listening and learning port states.

Switchover to the active uplink port

When a failed active uplink port becomes available again, switchover from the redundant port to the active uplink port is delayed by 30 seconds. The delay allows the remote port to transition to forwarding mode using the standard STP rules. After 30 seconds, the blocked active uplink port begins forwarding in just one second and the redundant port is blocked.

NOTE

Use caution when changing the spanning tree priority. If the switch becomes the root bridge, Fast Uplink Span will be disabled automatically.

Fast Uplink Span Rules for Trunk Groups

If you add a port to a Fast Uplink Span group that is a member of a trunk group, the following rules apply:

- If you add the primary port of a trunk group to the Fast Uplink Span group, all other ports in the trunk group are automatically included in the group. Similarly, if you remove the primary port in a trunk group from the Fast Uplink Span group, the other ports in the trunk group are automatically removed from the Fast Uplink Span group.

STP feature configuration

- You cannot add a subset of the ports in a trunk group to the Fast Uplink Span group. All ports in a trunk group have the same Fast Uplink Span property, as they do for other port properties.
- If the working trunk group is partially down but not completely down, no switch-over to the backup occurs. This behavior is the same as in the standard STP feature.
- If the working trunk group is completely down, a backup trunk group can go through an accelerated transition only if the following are true:
 - The trunk group is included in the fast uplink group.
 - All other ports except those in this trunk group are either disabled or blocked. The accelerated transition applies to all ports in this trunk group.

When the original working trunk group comes back (partially or fully), the transition back to the original topology is accelerated if the conditions listed above are met.

Configuring a Fast Uplink Port Group

To configure a group of ports for Fast Uplink Span, enter the following commands:

```
Brocade(config)# fast uplink-span ethernet 4/1 to 4/4
Brocade(config)# write memory
```

Syntax: [no] fast uplink-span [ethernet <port> [ethernet <port>... | to <port>]]

Specify the port variable in one of the following formats:

- FWS, FCX, and ICX stackable switches – stack-unit/slotnum/portnum
- FSX 800 and FSX 1600 chassis devices – slotnum/portnum
- ICX devices – slotnum/portnum
- FESX compact switches – portnum

This example configures four ports, 4/1 – 4/4, as a Fast Uplink Span group. In this example, all four ports are connected to a wiring closet switch. Only one of the links is expected to be active at any time. The other links are redundant. For example, if the link on port 4/1 is the active link on the wiring closet switch but becomes unavailable, one of the other links takes over. Because the ports are configured in a Fast Uplink Span group, the STP convergence takes one second instead of taking at least 30 seconds using the standard STP forward delay.

You can add ports to a Fast Uplink Span group by entering the fast uplink-span command additional times with additional ports. The device can have only one Fast Uplink Span group, so all the ports you identify as Fast Uplink Span ports are members of the same group.

To remove a Fast Uplink Span group or to remove individual ports from a group, use “no” in front of the appropriate fast uplink-span command. For example, to remove ports 4/3 and 4/4 from the Fast Uplink Span group configured above, enter the following commands:

```
Brocade(config)# no fast uplink-span ethernet 4/3 to 4/4
Brocade(config)# write memory
```

To check the status of ports with Fast Uplink Span enabled.

```
Brocade(config)# show span fast-uplink-span
```

```
STP instance owned by VLAN 1
```

```
Global STP (IEEE 802.1D) Parameters:
```

```
VLAN Root          Root Root   Prio Max He- Ho- Fwd Last   Chg Bridge
```

STP feature configuration

ID	ID	Cost	Port	Priority	Age	Hello	Hold	Delay	Change	Count	Address
				Hex	sec	sec	sec	sec	sec		
1	000000c100000001	2	1/3/1	8000	20	2	1	15	65	15	000011111111

Port STP Parameters:

Port Num	Priority	Path Cost	State	Fwd Trans	Design Cost	Designated Root	Designated Bridge
	Hex						
1/1/2	80	0	DISABLED	0	0	0000000000000000	0000000000000000
1/1/3	80	0	DISABLED	0	0	0000000000000000	0000000000000000
1/1/4	80	4	FORWARDING	1	2	000000c100000001	8000000011111111
1/1/5	80	0	DISABLED	0	0	0000000000000000	0000000000000000
1/1/6	80	0	DISABLED	0	0	0000000000000000	0000000000000000
1/1/7	80	0	DISABLED	0	0	0000000000000000	0000000000000000
1/1/8	80	0	DISABLED	0	0	0000000000000000	0000000000000000
1/1/9	80	0	DISABLED	0	0	0000000000000000	0000000000000000

Syntax: show span fast-uplink-span

Configuring Fast Uplink Span within a VLAN

You can also configure Fast Uplink Span on the interfaces within a VLAN.

To configure Fast Uplink Span for a VLAN, enter command such as the following.

```
Brocade(config)#vlan 10
Brocade(config-vlan-10)#untag ethernet 8/1 to 8/2
Brocade(config-vlan-10)#fast uplink-span ethernet 8/1 to 8/2
```

Syntax: [no] fast uplink-span ethernet <port-no>

To check the status of Fast Uplink Span for a specified VLAN.

```
Brocade(config-vlan-2)#show span vlan 2 fast-uplink-span
```

```
STP instance owned by VLAN 2
```

Global STP (IEEE 802.1D) Parameters:

VLAN ID	Root ID	Root Cost	Root Port	Priority	Max Age	Hello	Hold	Fwd Delay	Last Change	Chg Count	Bridge Address
				Hex	sec	sec	sec	sec	sec		
2	8000000011111111	0	Root	8000	20	2	1	15	29596	0	000011111111

Port STP Parameters:

Port Num	Priority	Path Cost	State	Fwd Trans	Design Cost	Designated Root	Designated Bridge
	Hex						
1/1/1	80	4	LISTENING	0	0	8000000011111111	8000000011111111

Syntax: show span vlan <vlan-id> fast-uplink-span

The VLAN <vlan-id> parameter displays Fast Uplink Span information for the specified VLAN.

Protecting against UDP attacks

The following section is an update to the NetIron Configuration Guide chapter titled **Protecting against Denial of Service Attacks**, specifically to the section titles **Protecting against UDP attacks**.

Limitation

In the Brocade NetIron CES and Brocade NetIron CER, UDP rate-limiting will only work in the following scenarios:

- When sending 1 % of 1G traffic with packet size of 64 bytes to the device for configured Burst-max value (up to 8000)
- When sending 10 % of 1G traffic with packet size of 64 bytes to the device for configured Burst-max value (up to 1500)
- When sending 100 % of 1G traffic with packet size of 64 bytes to the device for configured Burst-max value (up to 500).

ACL accounting on Brocade NetIron CES and Brocade NetIron CER devices

The following update has been applied to the
Enabling ACL accounting on Brocade NetIron CES and Brocade NetIron CER devices

NOTE

ACL accounting on Brocade NetIron CES and Brocade NetIron CER devices is applicable only on the outbound counter, not the inbound counter.

Displaying VLAN information

The following change is an update to Chapter 10 VLANs. The output is displayed as shown in the example below.

Displaying VLAN information for specific ports

To determine which VLANs a port is a member of, enter the following command.

```
Brocade# show vlan e 4/1
```

```
VLANS 1
```

```
VLANS 100
```

show vlan ethernet *slot-number/port-number* **[[begin** *expression* **exclude** *expression* **include** *expression*]

The **ethernet** *slot-number/port-number* parameter specifies a port. The command lists all the VLAN memberships for the port.

The output shows the following information.

Output of show vlan Ethernet Configuration Guide.fm

This field...	Displays...
VLANS	The IDs of the VLANs that the port is a member of.

Sflow sampling on Brocade NetIron CES and Brocade NetIron CER devices

This is an update to the section titled *Configuring and enabling sFlow* in the sFlow chapter.

NOTE

Sflow samples outbound traffic if the sflow enabled port is monitored by a mirror port.

On Brocade NetIron CES and Brocade NetIron CER devices, if mirrored Sflow packets are received in the LP CPU there is no option to distinguish them from regular Sflow packets.

LACP Enhancement

LACP flap counters

The **show lacp flap <port>** command shows the LACP flap counters and the corresponding timestamps and the index. Use the index to get detailed information about this flap instance.

Sample Output

```
MP#show lacp flap <9/3>
Port          : 9/3
Lag ID       : 1
Number of Flaps: 2

Index Timestamps
1  2012.03.20-09:26:28.365
2  2012.03.20-09:36:28.365
```

To obtain the histogram data at the time of the flap, use the existing show command **show sysmon events detail <index>**

```
Brocade# show sysmon events detail 1
-----
Sysmon Event Details
-----
Sysmon Event ID   - 1
Sysmon Event Type - LACP Flap Event
Sysmon Event Time - May  7 01:02:45
Sysmon Event Actions - Histogram(CPU, Buffer)
-----
HISTOGRAM TASK SEQUENCE INFO
-----
THRESHOLD   : 10 ms
DURATION    : 30 s
-----
Seq No  Task Name   Context  HoldTime   Start Time   End Time   Date
      Max (ms)
-----
   1 console      TASK      51 01:02:17.650 01:02:17.901 2013.05.07
   2 console      TASK      51 01:02:14.150 01:02:14.351 2013.05.07
   3 l4           TASK      55 01:02:12.964 01:02:13.020 2013.05.07
   4 snms         TASK      40 01:02:12.915 01:02:12.955 2013.05.07
   5 snms         TASK      17 01:02:12.896 01:02:12.913 2013.05.07
   6 snms         TASK      17 01:02:12.878 01:02:12.895 2013.05.07
   7 snms         TASK      18 01:02:12.859 01:02:12.877 2013.05.07
   8 snms         TASK      12 01:02:12.845 01:02:12.858 2013.05.07
   9 l4           TASK      56 01:02:12.782 01:02:12.838 2013.05.07
  10 snms         TASK      19 01:02:12.754 01:02:12.774 2013.05.07
  11 snms         TASK      17 01:02:12.736 01:02:12.754 2013.05.07
  12 snms         TASK      18 01:02:12.717 01:02:12.735 2013.05.07
  13 snms         TASK      18 01:02:12.698 01:02:12.716 2013.05.07
  14 scp          TASK      10 01:02:12.679 01:02:12.690 2013.05.07
  15 console      TASK      51 01:02:09.400 01:02:09.653 2013.05.07
  16 console      TASK      51 01:01:58.650 01:01:59.601 2013.05.07
  17 scp          TASK      27 01:01:56.804 01:01:56.831 2013.05.07
  18 scp          TASK      42 01:01:56.401 01:01:56.444 2013.05.07
```

19	scp	TASK	49	01:01:56.337	01:01:56.386	2013.05.07
20	scp	TASK	40	01:01:56.294	01:01:56.335	2013.05.07
21	console	TASK	16	01:01:56.170	01:01:56.187	2013.05.07
22	console	TASK	51	01:01:49.150	01:01:49.552	2013.05.07
23	l4	TASK	55	01:01:46.269	01:01:46.325	2013.05.07
24	snms	TASK	22	01:01:46.237	01:01:46.259	2013.05.07
25	snms	TASK	17	01:01:46.219	01:01:46.236	2013.05.07
26	snms	TASK	17	01:01:46.201	01:01:46.219	2013.05.07
27	snms	TASK	17	01:01:46.182	01:01:46.200	2013.05.07
28	snms	TASK	12	01:01:46.169	01:01:46.181	2013.05.07
29	scp	TASK	10	01:01:46.159	01:01:46.169	2013.05.07
30	scp	TASK	12	01:01:17.359	01:01:17.372	2013.05.07
31	scp	TASK	12	01:01:17.345	01:01:17.357	2013.05.07
32	scp	TASK	12	01:01:17.330	01:01:17.343	2013.05.07
33	scp	TASK	12	01:01:17.316	01:01:17.329	2013.05.07
34	scp	TASK	12	01:01:17.301	01:01:17.314	2013.05.07
35	scp	TASK	12	01:01:17.287	01:01:17.300	2013.05.07
36	snms	TASK	17	01:01:09.040	01:01:09.057	2013.05.07
37	snms	TASK	17	01:01:09.021	01:01:09.039	2013.05.07
38	snms	TASK	17	01:01:09.003	01:01:09.021	2013.05.07
39	snms	TASK	13	01:01:08.989	01:01:09.003	2013.05.07
40	snms	TASK	17	01:01:07.636	01:01:07.653	2013.05.07
41	snms	TASK	17	01:01:07.618	01:01:07.636	2013.05.07
42	snms	TASK	17	01:01:07.597	01:01:07.616	2013.05.07
43	snms	TASK	14	01:01:07.583	01:01:07.597	2013.05.07
44	scp	TASK	25	01:01:00.240	01:01:00.266	2013.05.07
45	scp	TASK	25	01:01:00.213	01:01:00.238	2013.05.07
46	scp	TASK	25	01:01:00.186	01:01:00.211	2013.05.07
47	scp	TASK	25	01:01:00.159	01:01:00.184	2013.05.07
48	scp	TASK	25	01:01:00.132	01:01:00.157	2013.05.07
49	scp	TASK	25	01:01:00.105	01:01:00.130	2013.05.07
50	scp	TASK	12	01:01:00.088	01:01:00.101	2013.05.07
51	scp	TASK	12	01:01:00.074	01:01:00.087	2013.05.07
52	scp	TASK	12	01:01:00.059	01:01:00.072	2013.05.07
53	scp	TASK	12	01:01:00.045	01:01:00.057	2013.05.07
54	scp	TASK	12	01:01:00.030	01:01:00.043	2013.05.07
55	scp	TASK	12	01:01:00.016	01:01:00.029	2013.05.07
56	scp	TASK	25	01:00:59.989	01:01:00.014	2013.05.07
57	scp	TASK	25	01:00:59.962	01:00:59.987	2013.05.07
58	scp	TASK	25	01:00:59.934	01:00:59.960	2013.05.07
59	scp	TASK	25	01:00:59.907	01:00:59.933	2013.05.07
60	scp	TASK	25	01:00:59.880	01:00:59.906	2013.05.07
61	scp	TASK	25	01:00:59.854	01:00:59.879	2013.05.07
62	scp	TASK	50	01:00:59.800	01:00:59.851	2013.05.07
63	scp	TASK	50	01:00:59.748	01:00:59.798	2013.05.07
64	scp	TASK	50	01:00:59.695	01:00:59.746	2013.05.07
65	scp	TASK	50	01:00:59.643	01:00:59.693	2013.05.07
66	scp	TASK	50	01:00:59.588	01:00:59.641	2013.05.07
67	scp	TASK	50	01:00:59.534	01:00:59.585	2013.05.07
68	main	TASK	10	01:00:59.512	01:00:59.522	2013.05.07

HISTOGRAM BUFFER SEQUENCE INFO

DURATION : 60 s

CSPF limitation

NOTE

The following is a limitation of the MPLS CSPF fate-sharing group.

CSPF calculates the least cost paths first and then applies the hop limit on the paths.

Fabric Auto Tuning SNMP and syslog enhancement

The following section describes an enhancement to the existing Slow Rate CRC Link Monitoring feature.

TABLE 19 Feature support table

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series BASE package	Brocade Netron CER 2000 Series Advanced Services package
Fabric Auto Tuning enhancement	Yes	Yes	No	No	No	No	No

NOTE

This feature is only applicable for 8x10G series, 2x100G, and FE600 cards when used in an MLXe chassis.

If the total CRC errors in a monitoring period pass the link down threshold, it will start auto tuning for that particular link instead of shutting down the link. After the completion of tuning, optimized Rx parameters obtained from the DFE tuning algorithm are applied to that link and bring that link up and ready to start transmit/receive traffic.

If the tuning algorithm returns error, depending sysmon tm link configuration, a syslog and trap sent and the link may be powered down. If a fabric module/line card is replaced or power cycled, this software starts freshly by allowing all links eligible for tuning again.

TM CLI command changes

The `sysmon tm link` command checks link status for errors such as slow rate CRCs on the TM side of the link and starts tuning the link if the number of CRC errors in a monitoring period exceed pre-defined threshold.

```
Brocade (config)# sysmon tm link
```

Syntax: `[no] sysmon tm link`

The default setting is enabled. Use the `no` command to disable the command.

```
Brocade (config)# sysmon tm link threshold 5 10
```

Syntax: [no] **sysmon tm link threshold** *error-threshold poll-window-size*

Set the *error-threshold* parameter for the error threshold value for a 60 second monitoring period. The minimum value is 1 the maximum value is 300. The default is 5.

Set the *poll-window-size* parameter to the number of 60 second monitoring periods in the polling window. The minimum value is 1 the maximum value is 300. The default is 10.

```
Brocade (config)# sysmon tm link action none
```

Syntax: [no] **sysmon tm link action** none | shutdown-link | syslog

Set the action for the link when the error threshold has been exceeded.

None - takes no action.

Shutdown-link - powers down the link when the error threshold has been exceeded.

Syslog - sends a syslog message when the error threshold has been exceeded.

Shutdown-link and **syslog** can be used together to send a syslog message and power down the link.

```
Brocade (config)# sysmon tm link log-backoff 60
```

Syntax: [no] **sysmon tm link log-backoff** *num*

Specify the log back-off period in which only one log message is sent. This parameter is only applicable when syslog is set as an action.

The *num* parameter is the number of seconds in which only one log message is sent. The acceptable range from 1 - 14400 seconds. The default value is 60 seconds.

Message examples

SYSLOG (If no action taken, just logging message):

```
Apr 30 15:32:16: I: System: Health Monitoring: TM link CRC errors: SNM5/FE1/Link16
? LP15/TM1/Link4
```

SYSLOG (If link is shutdown):

```
Apr 30 15:32:16: I: System: Health Monitoring: TM link shutdown due to CRC errors:
SNM5/FE1/Link16 ? LP15/TM1/Link4
```

TM Log Message (show tm log command output when CRC is detected):

```
Mar 4 20:33:57: TM Link CRC errors: SNM5/FE1/Link16 ? LP15/TM1/Link4
```

TM Log Message (show tm log command output when auto tuning started):

```
Mar 4 20:33:57: TM Link auto tuning started: SNM5/FE1/Link16 ? LP15/TM1/Link4
```

TM Log Message (show tm log command output when auto tuning finished):

```
Mar 4 20:33:57: TM Link auto tuning completed: SNM5/FE1/Link16 ? LP15/TM1/Link4
```

TM Log Message (show tm log command output when auto tuning failed):

```
Mar 4 20:33:57: TM Link Shutdown due to auto tuning failure: SNM5/FE1/Link16 ?
LP15/TM1/Link4
```

Fabric Auto Tuning SNMP and syslog enhancement

TM Log Message (*show tm log command output based on action taken*):

```
Mar 4 20:33:57: TM Link CRC Errors: SFM1/FE1/Link 15-> LP3/TM1/Link3
```

OR

```
Mar 4 20:33:57: TM Link Shutdown due to CRC Errors: SFM1/FE1/Link 15->
LP3/TM1/Link3
```

FE command changes

The `sysmon FE link` command checks link status for errors such as slow rate CRCs on the FE side of the link and starts tuning the link if the number of CRC errors in a monitoring period exceed pre-defined threshold.

```
Brocade (config)# sysmon fe link
```

Syntax: `[no] sysmon fe link`

The default setting is enabled. Use the `no` command to disable the command.

```
Brocade (config)# sysmon fe link threshold 5 10
```

Syntax: `[no] sysmon fe link threshold error-threshold poll-window-size`

Set the *error-threshold* parameter for the error threshold value for a 60 second monitoring period. The minimum value is 1 the maximum value is 300. The default is 5.

Set the *poll-window-size* parameter to the number of 60 second monitoring periods in the polling window. The minimum value is 1 the maximum value is 300. The default is 10.

```
Brocade (config)# sysmon fe link action none
```

Syntax: `[no] sysmon fe link action none | shutdown-link | syslog`

Set the action for the link when the error threshold has been exceeded.

None - takes no action.

Shutdown-link - powers down the link when the error threshold has been exceeded.

Syslog - sends a syslog message when the error threshold has been exceeded.

Shutdown-link and **syslog** can be used together to send a syslog message and power down the link.

```
Brocade (config)# sysmon fe link log-backoff 60
```

Syntax: `[no] sysmon fe link log-backoff num`

Specify the log back-off period in which only one log message is sent. This parameter is only applicable when `syslog` is set as an action.

The *num* parameter is the number of seconds in which only one log message is sent. The acceptable range from 1 - 14400 seconds. The default value is 60 seconds.

Message examples

SYSLOG (If no action taken, just logging message):

Apr 30 15:32:16: I: System: Health Monitoring: Fabric link CRC errors:
LP15/TM1/Link4 ? SNM5/FE1/Link16

SYSLOG (If link is shutdown):

Apr 30 15:32:16: I: System: Health Monitoring: Fabric link shutdown due to CRC
errors: LP15/TM1/Link4 ? SNM5/FE1/Link16

SFM Log Message (show sfm log command output when CRC is detected):

Mar 4 20:33:57: Fabric Link CRC errors: LP15/TM1/Link4 ? SNM5/FE1/Link16

SFM Log Message (show sfm log command output when auto tuning started):

Mar 4 20:33:57: Fabric Link auto tuning started: LP15/TM1/Link4 ? SNM5/FE1/Link16

SFM Log Message (show sfm log command output when auto tuning finished):

Mar 4 20:33:57: Fabric Link auto tuning completed: LP15/TM1/Link4 ?
SNM5/FE1/Link16

SFM Log Message (show sfm log command output when auto tuning failed):

Mar 4 20:33:57: Fabric Link Shutdown due to auto tuning failure: LP15/TM1/Link4 ?
SNM5/FE1/Link16

SFM Log Message (show sfm log command output based on action taken):

Mar 4 20:33:57: SFM Link CRC Errors: LP3/TM1/Link3 -> SFM1/FE1/Link 15
OR
Mar 4 20:33:57: SFM Link Shutdown due to CRC Errors: LP3/TM1/Link3 ->
SFM1/FE1/Link 15

Default global metric for ISIS

TABLE 20 Feature support table

Features supported	Brocade Netron XMR Series	Brocade MLX Series	Brocade Netron CES 2000 Series BASE package	Brocade Netron CES 2000 Series ME_PREM package	Brocade Netron CES 2000 Series L3_PREM package	Brocade Netron CER 2000 Series BASE package	Brocade Netron CER 2000 Series Advanced Services package
Enhanced support for Default metric for ISIS	Yes	Yes	Yes	Yes	Yes	Yes	Yes

ISIS has a default metric of 10 on the ISIS active interfaces. You can change the metric value for a specific interface by using the **isis metric** command or **isis ipv6 metric** command. This feature allows you to change the metric value globally for all the active ISIS interfaces using one command.

You can still configure the interface level metric. If ISIS metric is configured on the interface, it will take the precedence over the global configuration.

Configuration steps

1. Configure router ISIS using the `router isis` command.
2. Go to the appropriate address-family using `address-family [ipv4/ipv6] unicast` command.
3. Configure default metric using `default-link-metric <value>` command.

Configuration example

The following global configuration example ISIS default metric is for the IPv4 address-family. It can be similarly configured for IPv6 address-family.

```
Brocade(config)#router isis
Brocade(config-isis-router)#address-family ipv4 unicast
Brocade(config-isis-router-ipv4u)# default-link-metric 40
```

Syntax: `[no] default-link-metric value [level-1 | level-2]`

The *value* parameter is the default-link-metric value to be set for the given address-family. This is a required parameter for this command. There is no default value for this parameter. For metric-style narrow: 1 to 63. For metric-style wide: 1 to 16777215.

The *level* parameter is an optional parameter used to set the default-metric for only one of the levels. If this parameter is not given, the default-link-metric will be applied to both level-1 and level-2.

The **[no]** version of command will revert the metric value to default, which is 10.

IPv6 metric behavior with multi-topology configuration

The default-link-metric for IPv6 will depend upon the multi-topology configuration.

No multi-topology: The IPv6 default-link-metric will be same as that configured for IPv4 address-family.

Multi-topology: The IPv6 default-link-metric will be equal to the value configured for IPv6 address-family.

Multi-topology transition: The IPv6 default-link-metric will be equal to the value configured for IPv6 address-family.

Metric behavior with change in metric-style

There are two types of metric styles in ISIS, narrow metric and wide metric. The range of the metric value is different in both of these styles. If there is a change in the metric-style configuration, the default-link-metric will also change with it. The new value of the default-link-metric will be equal to the minimum of a) configured value and b) the maximum value supported for the new metric-style.

If the metric style changes from narrow metric to wide metric, there will be no change in the value of default-link-metric.

If the metric style changes from wide metric to narrow metric, and if the value of default-link-metric is greater than 63, the default-link-metric will now take the value 63, as it is the maximum supported in the narrow metric.

ISIS Show command

The `show isis` command and `show ipv6 isis` command output has been modified to reflect the default-link-metric configured.

```
Brocade#sh isis
....
Default redistribution metric: 0
Default link metric for level-1: 33
Default link metric for level-2: 5
Protocol Routes redistributed into IS-IS:
....
Brocade#
```

```
Brocade#sh ipv6 isis
....
Default redistribution metric: 0
Default link metric for level-1: 15
Default link metric for level-2: 9
Protocol Routes redistributed into IS-IS:
....
Brocade#
```

Configuring Secure Shell and Secure Copy

The following section replaces the same titled section in the NetIron 5.4.00 Configuration Guide.

Configuring DSA or RSA public key authentication

With DSA or RSA public key authentication, a collection of clients' public keys are stored on the Brocade device. Clients are authenticated using these stored public keys. Only clients that have a private key that corresponds to one of the stored public keys can gain access to the device using SSH server.

Setting up DSA or RSA private key authentication consists of the following steps.

1. Import authorized public keys into the Brocade device.
2. Enable DSA or RSA public key authentication.

Data Integrity Protection for Metro

The following section is an update to the Data Integrity Protection section of Chapter 3 in the NetIron 5.4.00 Configuration Guide.

TABLE 21 Feature support table

Features supported	Brocade NetIron XMR Series	Brocade MLX Series	Brocade NetIron CES 2000 Series BASE package	Brocade NetIron CES 2000 Series ME_PREM package	Brocade NetIron CES 2000 Series L3_PREM package	Brocade NetIron CER 2000 Series BASE package	Brocade NetIron CER 2000 Series Advanced Services package
Data Integrity Protection for Metro - Phase 2	No	No	Yes	Yes	Yes	Yes	Yes

Data Integrity Protection for Metro for Phase 2 introduces the ability to monitor low level memory corruption events occurring at the external Control Static Random Access Memory (CSRAM) in both Brocade NetIron CER and Brocade NetIron CES. Additionally, monitoring of the Longest Prefix Match (LPM) Memories are included as part of this feature.

There are a total of four LPM memories in total (LPM-0 to LPM-3). Brocade NetIron CER devices uses LPM-0 to LPM-2, which are on external memory chip. Brocade NetIron CES devices use LPM-3 only, which is internal. Brocade NetIron CES devices do not have external LPM memories.

Rolling windows are maintained for each of the monitor points. When any of the monitor points cross their configured thresholds SYSLOGs and traps are generated.

Configuring Data Integrity Protection for Metro

1. Configure the Global Rolling Window Time Frame.
2. Configure the threshold parameters for CSRAM and/or LPM memories.

New configuration commands

The following configuration commands are introduced to configure various parameters.

The **system np control-ram-threshold** command configures the CSRAM error reporting threshold.

```
Brocade(config)# system np control-ram-threshold 20
```

Syntax: [no] **system np control-ram-threshold** *threshold*

The threshold range is 0 - 120 events. The default is 10. A value of 0 disables the monitoring.

The [no] option resets the threshold to default.

The **system np lpm-ram-threshold** command configures the LPM error reporting threshold.

```
Brocade(config)# system np control-ram-threshold 20
```

Syntax: [no] **system np lpm-ram-threshold** *threshold*

The threshold range is 0 - 120 events. The default is 10. A value of 0 disables the monitoring.

The [no] option resets the threshold to default.

New show commands

The following show commands have been added to the feature.

show np control-ram-errors

The **show np control-ram-errors** command displays the Control RAM error event counter.

```
CSRAM
Ports          Current Cumulative
1/1 - 1/24      0         3
2/1 - 2/2      0         0
```

show np lpm-ram-errors

The **show np lpm-ram-errors** command displays the LPM RAM error event counter.

```

          LPM 0          LPM 1          LPM 2
Ports    Current Cumulative Current Cumulative Current Cumulative
1/1 - 1/24  0         3         0         3         0         3
2/1 - 2/2  0         3         0         3         0         3
```

Management module redundancy overview

Syslog messages

The following are examples of Syslog messages that may be displayed.

```
NP CSRAM has 4 error events, exceeding configured threshold for interfaces 1/1 to 1/24.
```

```
NP LPM 1 has 4 error events, exceeding configured threshold for interfaces 1/1 to 1/24.
```

Management module redundancy overview

The following section is an update to the Management module redundancy overview section of Chapter 6 in the NetIron 5.4.00 Configuration Guide.

The following paragraph replaces the first paragraph in the referenced section. The defined active-management module only takes over after a cold reboot (removing power to the chassis), not a reload.

Corrected text:

When you apply power to a Brocade device with two management modules installed, by default, the management module in slot M1 becomes the active module and the module in slot M2 becomes the standby module. (You can change the default active slot from M1 to M2 using the active-management command. Refer to “Changing the default active chassis slot” on page 199.)

Globally changing the IP MTU

The following note has been added to the section titled “Globally changing the IP MTU” in Chapter 28, Configuring IP.

NOTE

The global IP MTU change does not get applied to IP tunnel interfaces such as GRE interface. The MTU for these interfaces has to be changed on interface level.

Tuning multicast parameters

NOTE

The following commands are specific to the Brocade NetIron CER and Brocade NetIron CES devices:

- Multicast to Unicast descriptor ratio
- Multicast weight
- Multicast descriptor limit
- Multicast Traffic class mapping to Fabric Traffic class

Multicast to Unicast descriptor ratio

The `qos-multicast mcast-unicast-desc-ratio` command defines the number of Multicast Descriptors Duplications per Unicast Descriptor. The default value of this field is 1 which means there is one multicast descriptor replication per unicast. In multicast intensive environment user may want this parameter to be increased to 2 or 4 to give a better replication ratio to multicast traffic.

```
Brocade(config)# qos-multicast mcast-ucast-desc 1
```

Syntax: `[no] qos-multicast mcast-ucast-desc-ratio ratio-value [ppcr <ppcr_id>]`

The `ratio_value` parameter has a default value of 1. Valid values include 1, 2, and 4.

The `ppcr_id` parameter is an optional keyword. Valid values: 0 - for ports e1/1 - e1/24; 1 - for ports e1/25 - e1/28; 2 - for ports e2/1 - e2/2.

Without the `ppcr id` the command will set the descriptor ratio on all packet processors. If the `ppcr id` is specified then the command will set the descriptor ratio on only that packet processor.

Multicast weight

The Unicast/Multicast Arbiter in Hardware is WRR (weighted round robin). The defaults weight of Multicast and Unicast is assigned as 1. In a multicast intensive environment you may want to assign a higher multicast weight.

```
Brocade(config)# qos-multicast mcast-weight 15
```

Syntax: `[no] qos-multicast mcast-weight weight-value [ppcr <ppcr_id>]`

The `weight_value` parameter has a default value of 1. Valid values include 0 - 15.

The `ppcr_id` parameter is an optional keyword. Valid values: 0 - for ports e1/1 - e1/24; 1 - for ports e1/25 - e1/28; 2 - for ports e2/1 - e2/2.

Without the `ppcr id` the command will set the `weight value` on all packet processors. If the `ppcr id` is specified then the command will set the `weight value` on only that packet processor.

Multicast descriptor limit

Brocade NetIron CES and Brocade NetIron CER require a descriptor to place the packet in queue. If due to traffic conditions the descriptor limit is reached, the subsequent packets are dropped until the descriptors are available again. This parameter specifies the maximum number of descriptors that can be allocated for multicast packets. In a multicast intensive environment you may want to assign a higher value for the multicast descriptor limit. The default value of this is 2048.

```
Brocade(config)# qos-multicast mcast-desc-limit 8000
```

Syntax: [no] qos-multicast mcast-desc-limit *limit-value* [ppcr <ppcr_id>]

The *limit_value* parameter has a default value of 2048. Valid values include 0 - 8192.

The *ppcr_id* parameter is an optional keyword. Valid values: 0 - for ports e1/1 - e1/24; 1 - for ports e1/25 - e1/28; 2 - for ports e2/1 - e2/2.

Without the *ppcr id* the command will set the *limit value* on all packet processors. If the *ppcr id* is specified then the command will set the *limit value* on only that packet processor.

Multicast Traffic class mapping to Fabric Traffic class

The multicast traffic class is mapped to Fabric traffic class. Multicast TC 0, 1 is mapped to Fabric TC 0, Multicast TC 2,3 is mapped to Fabric TC 1, Multicast TC 4, 5 is mapped to Fabric TC 2, Multicast TC 6,7 is mapped to fabric TC 3. In the Egress Pipeline if Fabric TC is 2 and it is a multicast packet, the buffering of the descriptors occurs in a deep queue. This is a queue generally used in IPTV applications. You may want to change this mapping to assign all the multicast Traffic class to Fabric Traffic class 2 to utilize the deep buffer queue. This will improve the performance of IPTV applications.

```
Brocade(config)# qos-multicast mcast-tc-to-iptv all
```

Syntax: [no] qos-multicast mcast-tc-to-iptv *tc-value* | *all*

The *tc_value* parameter is the Multicast traffic class that needs to be mapped to Fabric traffic class 2. Valid values include 0 - 7. All will map all multicast traffic classes to Fabric class 2.

Displaying a traffic engineering path to a destination

The following is an update to the Displaying a traffic engineering path to a destination section of the Configuring MPLS Traffic Engineering chapter.

The option **path-name** *<path-name>* has been added to the `show mpls ted path` command.

path-name indicates to TED path to follow the path as defined by the MPLS path specified by the "path-name" option.

NOTE

The path specified by the path-name option will be constrained by the setting of "cspf-interface-constraint" in MPLS policies in the same manner as any other use of paths are constrained.

The updated section is updated as follows:

You can display a traffic engineering path to a IPv4 destination address using a specified set of resource parameters. This enhancement allows you to gain insight into a traffic engineering path in a network, before setting it up using RSVP. This will help you to avoid RSVP path setup failure due to unavailable requested resources along the path to the destination host.

To display the traffic engineering path to a IPv4 destination address, enter the following command.

```
Brocade# show mpls ted path 4.4.4.4
```

Syntax: `show mpls ted path <destIPAddress> path-name <path-name> [bandwidth <bw_in_kbps>] [hop-limit <max_hops>] [priority <setup_priority>] [exclude-any <STRING_List_of_Admin_Resource_groups_name_or_and_num>] [include-any <STRING_List_of_Admin_Resource_groups_name_or_and_num>] [include-all <STRING_List_of_Admin_Resource_groups_name_or_and_num>] [tie-breaking {random | least-fill | most-fill}]`

NOTE

When configuring the CLI command, any combination of resource constraint parameters can be used.

The path specified by the path-name option will be constrained by the setting of "cspf-interface-constraint" in MPLS policies in the same manner as any other use of paths are constrained.

The following table describes the parameters of the **show mpls ted path** command

TABLE 22 Parameters from the show mpls ted path command

CLI parameter	Description
<i><destIPAdress></i>	The IPv4 address of the destination host.
<i><path-name></i>	Indicates to TED path to follow the path as defined by the MPLS path specified by the "path-name" option.
bandwidth <i><bw_in_kbps></i>	The minimum bandwidth of the path to its destination. The bandwidth value is entered in decimal in kilo bits per second unit. The valid range is between 0- 2147483647. If the value entered is larger than 2147483647, then the value will be truncated to the max limit of 2147483647 and accepted as the bandwidth input.

Displaying a traffic engineering path to a destination

TABLE 22 Parameters from the show mpls ted path command (Continued)

CLI parameter	Description
hop-limit <max_hops>	The maximum hops for the path to reach to its destination. The valid range is between 0 - 255. If an invalid range is entered, then an error message will display. If a path to the destination is available, but the hop count for the path is greater than <max_hops> value, then MPLS will indicate that path is not available.
priority <setup_priority>	The setup priority of the path. The valid range is between 0 - 7. The default is 7, the lowest setup priority value. If an invalid range is entered, then an error message will display. The priority parameter should be entered along with the bandwidth parameter because while setting up an LSP, the setup priority value decides the ability to reserve a bandwidth amount.
exclude-any <STRING>	The <STRING> variable specifies the admin group name or number. The string must be enclosed in double quotes. The <STRING> variable is a list of any combination of admin group name and number. The valid range for the admin group number is between 0 and 31. The admin group name must start with an alphabet character. If an invalid range is entered for admin group number and admin group name, then the CLI will prompt a warning message. The CLI will be accepted, but the out of range value will be ignored.
include-any <STRING>	The <STRING> variable specifies the admin group name or number. The string must be enclosed in double quotes. The <STRING> variable is a list of any combination of admin group name and number. The valid range for the admin group number is between 0 and 31. The admin group name must start with an alphabet character. If an invalid range is entered for admin group number and admin group name, then the CLI will prompt a warning message. The CLI will be accepted, but the out of range value will be ignored.
include-all <STRING>	The <STRING> variable specifies the admin group name or number. The string must be enclosed in double quotes. The <STRING> variable is a list of any combination of admin group name and number. The valid range for the admin group number is between 0 and 31. The admin group name must start with an alphabet character. If an invalid range is entered for admin group number and admin group name, then the CLI will prompt a warning message. The CLI will be accepted, but the out of range value will be ignored.
tie-breaking {random least-fill most-fill}	The tie-breaking method is used when multiple equal cost paths to the destination exist. The tie-breaking rule will select only one path to be displayed from among multiple equal cost paths. The default is random.

Egress port and priority based rate shaping

The following section provides an update for multicast and IPTV for Brocade NetIron CES and Brocade NetIron CER devices.

Multicast queue size, flow control, rate shaping and egress buffer threshold

There are four internal priorities for multicast or broadcast traffic. These four priorities are mapped from the device's eight internal priorities as described in [Table 23](#)

TABLE 23 Mapping between multicast or broadcast and internal forwarding priorities

Internal Forwarding Priority	0,1	2,3	4,5	6,7
Multicast Internal Priority	0	1	2	3

The internal forwarding priority of a multicast or broadcast packet is determined from the packet's IEEE 802.1p priority, incoming port priority or IP ToS or DSCP as described in the "[Default QoS mappings](#)" on page 73. Four multicast queue types (0 to 3) are used for multicast internal priorities 0 to 3 respectively.

NOTE

Instead of ACL priority, use VLAN or port priority to prioritize multicast traffic.

Configuring multicast queue size

The following example configures a 2 MByte queue size for queue 0.

```
Brocade(config)# qos multicast-queue-type 0 max-queue-size 2048
```

Syntax: [no] qos multicast-queue-type *queue-number* max-queue-size *queue-size*

The *queue-number* variable specifies the queue that you want to configure a maximum size for. Possible values are 0 - 3.

The *queue-size* variable specifies size in KBytes that you want to set as the maximum value for the specified multicast queue. Possible values are 1 - 32768 KBytes. The default queue size is 1 Mbyte. This command is applied per device and takes effect on all Traffic Managers within the configured device.

Configuring multicast flow control

Flow controls are available from egress to Ingress, and from fabric to Ingress. At the egress of each Traffic Manager, there are pre-determined thresholds for consumed resources and available resources and separate thresholds for guaranteed multicast or broadcast traffic and best-effort multicast or broadcast traffic. When a threshold is crossed, flow control can be triggered and multicast or broadcast traffic of the corresponding class is stopped at Ingress until resources are below the threshold again. Flow control is disabled by default and can be enabled on an interface using the command shown in the following.

```
Brocade(config)# interface ethernet 2/2
Brocade(config-if-e10000-2/2)# qos multicast flow-control
```

Syntax: [no] qos multicast flow-control

This command changes the flow control setting on the Traffic Manager where the interface resides.

Egress port and priority based rate shaping

Configuring multicast rate shaping

You can specify either guaranteed or best effort multicast rate shaping for a port in Kilobits per second. Multicast rate shaping is configured per-port to the Ingress port.

The following example changes the best-effort multicast traffic rate to 10 Mbps.

```
Brocade(config)# interface ethernet 2/2
Brocade(config-if-e10000-2/2)# qos multicast shaper best-effort rate 10000
```

Syntax: [no] qos multicast shaper [guaranteed | best-effort] rate *bandwidth*

The **guaranteed** option specifies that the multicast or broadcast shaper applies only to internal multicast priority 3 (the highest multicast priority) traffic.

The **best-effort** option specifies that the multicast or broadcast shaper applies to internal multicast priority 0, 1 and 2 traffic only.

The *bandwidth* variable specifies the maximum bandwidth in Kbps for best effort or guaranteed multicast traffic scheduled by the Traffic Manager across the switch fabric.

Configuration considerations for multicast rate shaping

When applied to a port, the **qos multicast shaper** configuration is applied to all ports on the Interface module that use the same Traffic Manager as the configured port. This is unlike the behavior of rate shaping applied for unicast traffic. The relationship between ports and Traffic Managers is defined in the following tables: [Table 35](#) on page 135 and [Table 36](#) on page 135.

Example

In the following example, multicast rate shaping is applied to port 1 of a Brocade NetIron CER or Brocade NetIron CES.

```
Brocade(config)# interface ethernet 1/1
Brocade(config-if-e10000-1/1)# qos multicast shaper best-effort rate 10000
```

In this example, the configuration will apply to Ingress traffic that arrives on either port 1 or port 2 of the device.

NOTE

When a **qos multicast shaper** command is configured for a port, the configuration command is placed in the running config for all ports that belong to the same Traffic Manager. In the example, that would mean that the **qos multicast shaper best-effort rate 10000** command would appear in the interface configuration section for ports 1 and 2 on the Interface Module.

Configuring multicast egress buffer threshold

With the current configuration egress buffer threshold per port are set to 50% of total egress buffer size, with the new command you can set multicast egress buffer threshold up to 95% of total egress buffer size which helps in reducing egress packet drops when there is a high multicast traffic.

NOTE

It is recommended to use this command when a sudden burst is seen in multicast traffic and not for general use.

The following example explains how to set the thresholds for individual ports so that the port can use the total egress buffer. This is useful when the multicast traffic is very high.

```
Brocade#config terminal
Brocade(config)#interface ethernet 1/1
Brocade(config-if-e10000-1/1)# qos multicast egress-max-buffer port 95% 90%
```

Syntax: `qos multicast egress-max-buffer port { [guaranteed_max_buffer] [best-effort_max_buffer] }`

The **guaranteed_max_buffer** specifies the maximum buffer size allowed per port for guaranteed traffic flow (multicast port priority 3). Specified as percentage of total buffer size.

The **best-effort_max_buffer** specifies the maximum buffer size allowed per port for guaranteed traffic flow (multicast port priority 2-0). Specified as percentage of total buffer size.

Additionally you can also have the threshold configured for individual ports so that each port has its dedicated buffer space.

The example uses a 8x10 line card and has four ports per TM. The buffer size is divided into 4 times the total buffer size so that each port has its dedicated buffer space.

```
Brocade#config terminal
Brocade(config)#interface ethernet 1/1
Brocade(config-if-e10000-1/1)# qos multicast egress-max-buffer port 24% 23%
```

Syntax: `qos multicast egress-max-buffer port { [guaranteed_max_buffer] [best-effort_max_buffer] }`

The **guaranteed_max_buffer** specifies the maximum buffer size allowed for guaranteed traffic flow (multicast port priority 3). Specified as percentage of total buffer size.

The **best-effort_max_buffer** specifies the maximum buffer size allowed for guaranteed traffic flow (multicast port priority 2-0). Specified as percentage of total buffer size.

Ingress traffic shaping per multicast stream

Internet Protocol Television (IPTV) multicast streams on an individual inbound physical port are rate shaped to a specified rate and are prioritized over the broadcast or unknown-unicast traffic. Each IPTV multicast stream is queued separately and is scheduled independently to the outbound ports. The IPTV rate shaping reduces burstiness in the source stream.

NOTE

The number of active IPTV multicast streams for which per stream ingress shaping can be applied is limited to 512.

Implementation considerations

The considerations for implementing the multicast rate shaping are as follows:

- At least the rate or the priority value must be specified.
- A maximum of 32 profiles and 64 profile-ACL bindings are allowed for multicast traffic.
- A single profile can be bound to multiple ACLs.
- An ACL can be associated only to one profile at a time.

Egress port and priority based rate shaping

- Either standard or extended ACLs can be used for multicast traffic shaping.
 - When a standard ACL is used, the address specified is treated as a group address and not as a source address.
 - When an extended ACL is used, the source and the destination addresses are treated as the source and group address of the multicast stream, respectively.

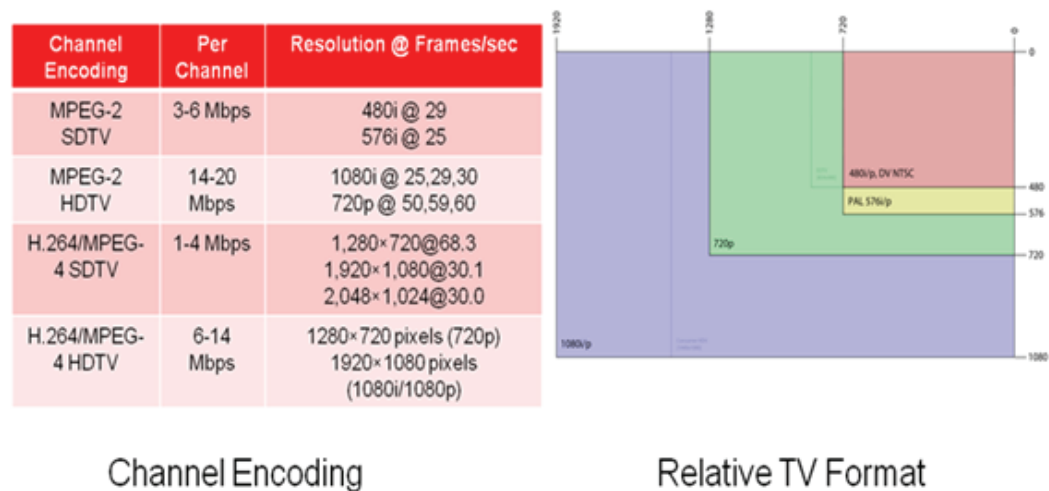
Figure 4 shows the type of IPTV channels and the bandwidth requirements for each type of channel. The following conventions are used in the figure:

- SDTV denotes Standard Definition Television
- HDTV denotes High Definition Television

NOTE

This feature is supported only on the 4x10 and 24x1 interface modules.

FIGURE 4 IPTV Bandwidth Requirements



Configuring multicast traffic policy maps

You can define profiles to match the IPTV multicast traffic of the individual ingress streams. To configure a policy map for the multicast streams, enter the following command.

```
Brocade(config)# policy-map multicast sd_prof rate 2000 burst-size 2500 priority 2
queue-type 3
```

Syntax: [no] policy-map multicast *profile_name* [rate *r*] [burst-size *b*] [priority 0-7] [queue-type 0-3]

The *profile_name* is used to provide the parameters for traffic policing the multicast traffic.

The **rate** *r* variable specifies the shaping rate in bits per second. The value ranges from 100 kilobits per second (Kbps) through 20 gigabits per second (Gbps).

The **priority** 0-7 variable specifies the multicast traffic priority. The default value is four.

The **burst-size** *b* variable specifies the maximum number of bytes the multicast traffic is allowed to burst. The value ranges from 3 kilobytes (KB) through 128 KB. The default value is four KB.

The **queue-type** 0-3 variable specifies the queue type for which you want to set the priority. The default value is two. Optionally, you can specify the **burst-size** *b* and the **queue-type** 0-3 value to define a profile.

To delete a defined profile, enter the following command with the profile name.

```
Brocade(config)# policy-map multicast sd_prof
```

The **no** form of the command resets the parameters to their default values.

NOTE

The rate and the priority value cannot be reset for a defined profile.

Binding multicast traffic policy maps

NOTE

A profile must exist in the configuration before it can be used for binding.

A standard or an extended ACL is used to define the IPTV streams that can be bound to a defined profile. The profile binding associates the properties of the profile to all the IPTV streams identified by the ACL. Binding of multicast streams can be done for Layer 3 multicast routing and Layer 2 multicast snooping.

Profile binding for Layer 3 multicast routing

You can bind a defined profile to a defined ACL for the default VRF or a specific VRF.

To bind the profile to the ACL in the default VRF, enter the following commands.

```
Brocade(config)# ip multicast-routing policy-map r1 sd-1
Brocade(config)# ipv6 multicast-routing policy-map r1q0 sdv61
```

To bind the profile to the ACL for a specified VRF, enter the following commands within the VRF “red” configuration context.

```
Brocade(config)# ip vrf red
Brocade(config-vrf-red)# address-family ipv4
Brocade(config-vrf-red-ipv4)# ip multicast-routing policy-map r1 sd-1
Brocade(config-vrf-red-ipv4)# exit-address-family
Brocade(config-vrf-red)# exit-vrf
```

```
Brocade(config)# ip vrf red
Brocade(config-vrf-red)# address-family ipv6
Brocade(config-vrf-red-ipv6)# ipv6 multicast-routing policy-map r1q0 sdv61
Brocade(config-vrf-red-ipv6)# exit-address-family
Brocade(config-vrf-red)# exit-vrf
```

Syntax: [no] ip multicast-routing policy-map *profile_name* *acl_id* | *acl_name*

The **ip multicast-routing policy-map** specifies ACL binding for IPv4 multicast routing.

The *profile_name* variable specifies the profile name of the multicast stream.

The *acl_id* | *acl_name* variable specifies the number and name of the standard ACL or an extended ACL. Enter a number from 1 through 99 for a standard ACL, and a number from 100 through 199 for an extended ACL.

Syntax: [no] ipv6 multicast-routing policy-map *profile_name* *acl_id* | *acl_name*

Egress port and priority based rate shaping

The **ipv6 multicast-routing policy-map** specifies ACL binding for IPv6 multicast routing.

The **no** form of the command removes the profile binding with the ACL in default VRF.

Profile binding for Layer 2 multicast snooping

You can bind a defined profile to a defined ACL per VLAN or per VPLS instance. To bind the profile to the ACL on VLAN 10, enter the following commands.

```
Brocade(config)# vlan 10
Brocade(config-vlan-10)# ip multicast policy-map r2q1 2
Brocade(config-vlan-10)# ipv6 multicast policy-map r4q3 sdv64
```

Syntax: [no] ip multicast policy-map *profile_name acl_id | acl_name*

The **ip multicast policy-map** specifies the ACL binding for IPv4 multicast snooping.

Syntax: [no] ipv6 multicast policy-map *profile_name acl_id | acl_name*

The **ipv6 multicast policy-map** specifies the ACL binding for IPv6 multicast snooping.

The **no** form of the command removes the profile binding with the ACL on the VLAN or VPLS.

In the following example, binding for Layer 2 multicast snooping is applied to VPLS instance V1.

```
Brocade(config)# router mpls
Brocade(config-mpls)# vpls v1 10
Brocade(config--mpls-vpls-v1)# multicast policy-map r2q1 2
Brocade(config--mpls-vpls-v1)# multicast policy-map r4q3 sdv64
```

Syntax: [no] multicast policy-map *profile_name acl_id | acl_name*

NOTE

A profile that is bound cannot be deleted.

Configuration example for rate shaping IPTV multicast stream

The following example shows how to rate shape a multicast stream. In this example, to rate shape a multicast stream, profiles (*sd_prof* and *hd_prof*) are defined with rate and priority, ACLs (*hd_streams* and *sd_streams*) are configured which permit packets from four host IP addresses and denies all packets that are not explicitly permitted by the first four ACL entries, and then the profiles are bound to the defined ACLs.

```
Brocade(config)# ip access-list standard hd_streams
Brocade(config-std-nacl)# permit host 239.1.1.200
Brocade(config-std-nacl)# permit host 239.1.1.201
Brocade(config-std-nacl)# permit host 239.1.1.202
Brocade(config-std-nacl)# permit host 239.1.1.203
Brocade(config-std-nacl)# deny any
Brocade(config-std-nacl)# exit

Brocade(config)# ip access-list extended sd_streams
Brocade(config-ext-nacl)# permit ip any 239.1.1.1/32
Brocade(config-ext-nacl)# permit ip any 239.1.1.2/32
Brocade(config-ext-nacl)# permit ip any 239.1.1.3/32
Brocade(config-ext-nacl)# permit ip any 239.1.1.5/32
Brocade(config-ext-nacl)# deny ip any any
Brocade(config-ext-nacl)# exit
```

```

Brocade(config)# policy-map multicast profile sd_prof rate 2000
Brocade(config)# policy-map multicast profile hd_prof rate 14000 queue-type 2

Brocade(config)# ip multicast-routing policy-map sd_prof sd_streams
Brocade(config)# ip multicast-routing policy-map hd_prof hd_streams

```

Naming convention used in example

- sd_prof = Standard Definition profile
- hd_prof = High Definition profile
- sd_streams = Standard Definition TV stream
- hd_stream = High Definition TV stream

Tuning multicast parameters

NOTE

The following commands are specific to the Brocade NetIron CER and Brocade NetIron CES devices:

- Multicast to Unicast descriptor ratio
- Multicast weight
- Multicast descriptor limit
- Multicast Traffic class mapping to Fabric Traffic class

Multicast to Unicast descriptor ratio

The `qos-multicast mcast-unicast-desc-ratio` command defines the number of Multicast Descriptors Duplications per Unicast Descriptor. The default value of this field is 1 which means there is one multicast descriptor replication per unicast. In multicast intensive environment user may want this parameter to be increased to 2 or 4 to give a better replication ratio to multicast traffic.

```
Brocade(config)# qos-multicast mcast-ucast-desc 1
```

Syntax: `[no] qos-multicast mcast-ucast-desc-ratio ratio-value [ppcr <ppcr_id>]`

The `ratio_value` parameter has a default value of 1. Valid values include 1, 2, and 4.

The `ppcr_id` parameter is an optional keyword. Valid values: 0 - for ports e1/1 - e1/24; 1 - for ports e1/25 - e1/28; 2 - for ports e2/1 - e2/2.

Without the `ppcr id` the command will set the descriptor ratio on all packet processors. If the `ppcr id` is specified then the command will set the descriptor ratio on only that packet processor.

Multicast weight

The Unicast/Multicast Arbiter in Hardware is WRR (weighted round robin). The defaults weight of Multicast and Unicast is assigned as 1. In a multicast intensive environment you may want to assign a higher multicast weight.

```
Brocade(config)# qos-multicast mcast-weight 15
```

Syntax: `[no] qos-multicast mcast-weight weight-value [ppcr <ppcr_id>]`

The `weight_value` parameter has a default value of 1. Valid values include 0 - 15.

The `ppcr_id` parameter is an optional keyword. Valid values: 0 - for ports e1/1 - e1/24; 1 - for ports e1/25 - e1/28; 2 - for ports e2/1 - e2/2.

Egress port and priority based rate shaping

Without the *ppcr id* the command will set the *weight value* on all packet processors. If the *ppcr id* is specified then the command will set the *weight value* on only that packet processor.

Multicast descriptor limit

Brocade NetIron CES and Brocade NetIron CER require a descriptor to place the packet in queue. If due to traffic conditions the descriptor limit is reached, the subsequent packets are dropped until the descriptors are available again. This parameter specifies the maximum number of descriptors that can be allocated for multicast packets. In a multicast intensive environment you may want to assign a higher value for the multicast descriptor limit. The default value of this is 2048.

```
Brocade(config)# qos-multicast mcast-desc-limit 8000
```

Syntax: [no] qos-multicast mcast-desc-limit *limit-value* [ppcr <ppcr_id>]

The *limit_value* parameter has a default value of 2048. Valid values include 0 - 8192.

The *ppcr_id* parameter is an optional keyword. Valid values: 0 - for ports e1/1 - e1/24; 1 - for ports e1/25 - e1/28; 2 - for ports e2/1 - e2/2.

Without the *ppcr id* the command will set the *limit value* on all packet processors. If the *ppcr id* is specified then the command will set the *limit value* on only that packet processor.

Multicast Traffic class mapping to Fabric Traffic class

The multicast traffic class is mapped to Fabric traffic class. Multicast TC 0, 1 is mapped to Fabric TC 0, Multicast TC 2,3 is mapped to Fabric TC 1, Multicast TC 4, 5 is mapped to Fabric TC 2, Multicast TC 6,7 is mapped to fabric TC 3. In the Egress Pipeline if Fabric TC is 2 and it is a multicast packet, the buffering of the descriptors occurs in a deep queue. This is a queue generally used in IPTV applications. You may want to change this mapping to assign all the multicast Traffic class to Fabric Traffic class 2 to utilize the deep buffer queue. This will improve the performance of IPTV applications.

```
Brocade(config)# qos-multicast mcast-tc-to-iptv all
```

Syntax: [no] qos-multicast mcast-tc-to-iptv *tc-value* | *all*

The *tc_value* parameter is the Multicast traffic class that needs to mapped to Fabric traffic class 2. Valid values include 0 - 7. All will map all multicast traffic classes to Fabric class 2.

Applying traffic policing parameters directly to a port

A correction to the Configuring Traffic Policing for the Brocade NetIron CES and Brocade NetIron CER chapter of the Configurations guide. The “maximum burst” values are expressed in bits and not bytes. A new table is also added to this section.

Maximum burst

Maximum burst provides a higher than average rate to traffic that meet the rate limiting criteria. Traffic will be allowed to pass through the port for a short period of time. The unused bandwidth can be accumulated up to a maximum of “maximum burst” value expressed in bits.

Maximum burst size is adjusted according the configured average line rate. If the user configured maximum burst size value exceeds the maximum burst size allowed, the maximum burst size will be automatically adjusted to the values indicated in [Table 24](#) on page 79.

TABLE 24 Maximum Burst Size

Average rate (bps)	Maximum burst size (Bits)
1 Mbps	66,535
1 - 10 Mbps	524,280
10 - 100 Mbps	4,194,240
100 Mbps - 1 Gbps	33,553,920
1 Gbps - 10 Gbps	268,431,230

Deprecated MPLS Commands

The following MPLS commands are deprecated starting from R5.4 release:

```
show mpls rsvp traffic
show mpls ldp traffic
```

OSPF with MCT

OSPF is supported as passive using MCT in R5.4 release.

Site-Local IPv6 Addresses

Site-local IPv6 addresses are not supported starting from R5.4 release. Unique local IPv6 unicast addresses (ULAs) are supported, but treated same as global unicast addresses. You need to configure filters to ensure that ULAs are not advertised outside the sites. For more information on ULAs, refer to RFC 4193. For the complete list of supported IPv6 addresses, refer to RFC 4291.

IPv6 ND Router Advertisement Control

IPv6 ND Router Advertisement Control allows for disabling sending out router advertisements at the interface level. The **no ipv6 nd suppress-ra** command at the interface level allows the user to disable and enable the sending of the ND Router Advertisement on an interface. By default, the sending of ND Router Advertisement (RA) is enabled on all interfaces, except for the tunnel and loopback interfaces, providing that the IPv6 Unicast Routing is enabled and the interfaces are active for IPv6.

The IPv6 ND Router Advertisement Control gives the ability to quickly turn off the sending of IPv6 ND Router Advertisement message on an IPv6 enabled interfaces.

By default,

- The ND Router Advertisement is enabled.
- Interface is enabled to send ND Router Advertisements.
- The **ipv6 nd suppress-ra** and **ipv6 nd send-ra** interface commands, when configured, override the system and VRF global **ipv6 nd global-suppress-ra** command.

Client-role-revertible-delay timer

Users sometimes require the ability to quickly turn off the sending of IPv6 ND Router Advertisement message on an IPv6 enabled interfaces. This is achieved by providing the following additional configuration command at interface level:

```
Brocade(config-if-e10000-1/1)#no ipv6 nd suppress-ra
```

The **ipv6 nd send-ra** command is a new interface level command added as part of this enhancement. This allows the user to configure the sending of RA messages on some selected interfaces when the **ipv6 nd global-suppress-ra** command is set to disable the sending of RA messages on all other interfaces.

Syntax: [no]ipv6 nd suppress-ra

Client-role-revertible-delay timer

The **client-role-revertible-delay timer** command is used for configuring client role revertible delay mode in case of port flaps, and to revert the client role after client role switch.

```
Brocade(config-cluster-c-client-c1)#client-role-revertible-delay timer 5
```

Syntax: [no] client-role-revertible-delay timer <value>

Enter a value of 1-240 minutes. Default value is 1 minute.

Port-based Rate Limiting

NOTE

Port-based rate-limiting is not supported with PBR on the same interface.

MSTP support for PBB

The following configuration consideration has been added.

MSTP should not be configured for:

- topology groups having layer 2 (L2) member VLANs
- member VLANs configured in a topology group. A

If a topology group is configured with a master VLAN running MSTP, layer 2 (L2) VLANs should not be configured as members until MSTP is disabled on the master VLAN of this topology group. Such configurations via CLI are blocked.

RPF

Configuration considerations for RPF

The following configuration consideration has changed.

The item stating Brocade MLX series and Brocade NetIron XMR devices do not support uRPF for VE interfaces is incorrect. The following configuration consideration replaces the incorrect information.

- RPF can only be configured at the physical port level. It should not be configured on virtual interfaces on the Brocade MLX series and Brocade NetIron XMR.
- Brocade MLX series and Brocade NetIron XMR devices support uRPF for VE interfaces, but they must be configured at the physical port level.

Keep-alive VLAN

The following configuration consideration has been added.

NOTE

A port in keep-alive-vlan cannot be assigned to another VLAN.

Configuring FDP

The following section describes how to enable FDP and how to change the FDP update and hold timers.

Enabling FDP globally

To enable a Brocade device to globally send FDP packets, enter the following command at the global CONFIG level of the CLI.

```
Brocade(config)# fdp run
```

The feature is disabled by default.

NOTE

If FDP is globally enabled on a Brocade device, all the interfaces by default, will have FDP enabled on it. In this case, the **show run** command will not display any running information about the FDP configuration in its output.

Enabling FDP at the interface level

You can enable FDP at the interface level by entering the following commands.

```
Brocade(config)# int e2/1
Brocade(config-if-e10000-2/1)# fdp enable
```

Enabling interception of CDP packets globally

Syntax: [no] fdp enable

By default, the feature is enabled on an interface once FDP is enabled on the device. It is not enabled globally.

NOTE

To remove an interface from the global configuration, run the **no fdp enable** command in the interface mode explicitly. In this case, the **show run** displays the running configuration information for the specific interface at that instance.

NOTE

By removing FDP from the configuration, the **no fdp enable** stays in the configuration of the VPLS endpoint, which cannot be removed.

NOTE

FDP is not supported on VPLS/VLL endpoints.

Enabling interception of CDP packets globally

To enable the device to intercept and display CDP packets, enter the following command at the global CONFIG level of the CLI.

```
Brocade(config)# cdp run
```

Syntax: [no] cdp run

The feature is disabled by default.

NOTE

If CDP is globally on a Brocade device, all of the interfaces, by default, will have CDP enabled on it. In this case, the **show run** will not display any running information about CDP configuration in its output.

Enabling interception of CDP packets on an interface

You can disable and enable CDP at the interface level.

You can enter the following commands.

```
Brocade(config)# int e2/1
Brocade(config-if-e10000-2/1)# cdp enable
```

Syntax: [no] cdp enable

By default, the feature is enabled on an interface on CDP is enabled on the device.

NOTE

To remove an interface from the global CDP configuration, run the **no cdp enable** command in the interface mode explicitly. In this case, the **show run** displays the running configuration information for the specific interface at that instance.

NOTE

By removing FDP from the configuration, the **no cdp enable** stays in the configuration of the VPLS endpoint, which cannot be removed.

NOTE

CDP is not supported on VPLS/VLL endpoints.

Configuring VPLS endpoint over FDP/CDP interface

Configuring VPLS endpoint over a FDP/CDP enabled interface will implicitly disable the FDP/CDP configuration on that specific interface for that instance, considering FDP/CDP is enabled globally. In this case, the **show run** command will display the running configuration information as shown below.

The following examples explain the **show run** output for different instances:

- The **show run** output when the VPLS endpoint is configured over a globally enabled FDP/CDP interface:

```
Brocade(config-mpls-vpls-svlan-vlan-100)# tag eth 4/3 eth 4/5 eth 4/7
FDP/CDP is disabled on port 4/3
FDP/CDP is disabled on port 4/5
FDP/CDP is disabled on port 4/7
```

- The **show run** output when the VPLS endpoint is configured over a globally enabled FDP/CDP interface:

```
Brocade(config-mpls-vpls-svlan-vlan-100)# tag eth 4/3 eth 4/5 eth 4/7
FDP/CDP is disabled on port 4/3
FDP/CDP is disabled on port 4/5
FDP/CDP is disabled on port 4/7
```

- The **show run** output when the VPLS output is removed over a globally enabled FDP/CDP interface:

```
FDP/CDP is enabled on port 4/3
FDP/CDP is enabled on port 4/5
FDP/CDP is enabled on port 4/7
```

- The **show run** output when the VPLS endpoint is removed over a globally enabled FDP/CDP interface:

```
FDP/CDP is enabled on port 4/3
FDP/CDP is enabled on port 4/5
FDP/CDP is enabled on port 4/7
```

NOTE

If a VPLS endpoint is configured over a globally enabled FDP/CDP interface, the **show run** will not display FDP/CDP information for that specific interface until the VPLS endpoint is deleted. On deleting the VPLS endpoints, the previous FDP/CDP configuration is retained over that specific interface and the **show run** displays the FDP/CDP information again for that interface.

NOTE

By removing the FDP/CDP from the configuration, the **no cdp enable** or **no fdp enable** stays in the configuration of the VPLS endpoint, both of which cannot be removed.

Configuring VLL endpoint over FDP/CDP enabled interface

Configuring VLL endpoint over an FDP/CDP enabled interface will implicitly disable the FDP/CDP configuration and also will be enable back implicitly when the VLL endpoint is deleted on that specific interface, considering the FDP/CDP is enabled globally.

Information messages will be displayed to notify the user as below in these cases:

For example, when VLL endpoint is created, the information messages are as below.

1. When only FDP is enabled globally

```
Brocade(config-mpls-vll-vll11-vlan-100)# tag eth 4/3 eth 4/5 eth 4/7
info- FDP is disabled on port 4/3
info- FDP is disabled on port 4/5
info- FDP is disabled on port 4/7
```

2. When only CDP is enabled globally

```
Brocade(config-mpls-vll-vll11-vlan-100)# tag eth 4/3 eth 4/5 eth 4/7
info- FDP is disabled on port 4/3
info- FDP is disabled on port 4/5
info- FDP is disabled on port 4/7
```

3. When both FDP/CDP are enabled globally

```
Brocade(config-mpls-vll-vll11-vlan-100)# tag eth 4/3 eth 4/5 eth 4/7
info- FDP is disabled on port 4/3
info- FDP is disabled on port 4/5
info- FDP is disabled on port 4/7
```

For example, when the VLL endpoint is deleted the information messages are displayed as below.

1. When only FDP is enabled globally

```
Brocade(config-mpls-vll-vll11-vlan -100)# no tag eth 4/3 eth 4/5 eth 4/7
info - FDP is enabled on port 4/3
info - FDP is enabled on port 4/5
info - FDP is enabled on port 4/7
```

2. When only CDP is enabled globally

```
Brocade(config-mpls-vll-vll11-vlan-100)# no tag eth 4/3 eth 4/5 eth 4/7
info - FDP is enabled on port 4/3
info - FDP is enabled on port 4/5
info - FDP is enabled on port 4/7
```

3. When both FDP/CDP are enabled globally

```
Brocade(config-mpls-vll-vll11-vlan-100)# no tag eth 4/3 eth 4/5 eth 4/7
info - FDP/CDP is enabled on port 4/3
```

```
info - FDP/CDP is enabled on port 4/5  
info - FDP/CDP is enabled on port 4/7
```

NOTE

If the VLL endpoint is configured over a globally enabled FDP/CDP interface, the show run command does not display the FDP/CDP information for that specific interface.

NOTE

By removing FDP/CDP from the configuration, the **no fdp enable** and **no cdp enable** stays in the configuration of the VPLS endpoints, which cannot be removed.

PIM over MCT

The MCT feature interaction matrix, Chapter 27, has been updated to reflect that PIM over MCT is supported in NetIron 5.4.00e and newer.

MCT feature interaction

Use the following feature matrix when configuring MCT:

MCT feature interaction matrix

Supported	Not Supported
LACP on both ICL and CCEP.	MSTP, VSRP, RIP, OSPF, IS-IS, IPv6, VRRPv3, IPv6 VRRP-E, and BGP
PIM is supported over ICL and CCEP links	ESI VLANs on CCEP or ICL ports.
VRRP on the CCEP.	GRE is not supported on the ICL ve interfaces
MRP and MRP II supported with the restriction that ICL port cannot be the secondary port of the MRP ring.	DAI on the CCEP ports.
802.1ad on CCEP or ICL ports.	802.1ah on CCEP or ICL ports.
Flooding features (VLAN CPU protection, multicast flooding etc.) on cluster VLANs.	VPLS on CCEP or ICL ports.
Multi-VRF	VLL on CCEP or ICL ports.
UDLD as independent boxes.	MPLS on CCE or ICL ports.
Link OAM as independent boxes.	
802.1ag as independent boxes.	
ARP as independent boxes.	
STP and RSTP	MSTP
L3 Routing - The IP address assignment is OK on CCEP ports for VRRP purpose. However, routing protocols would not be enabled on CCEP ports.	Hitless Fail over is NOT supported on the Brocade MLX series, and NetIron XMR, however it is compatible. If the operation is performed with cluster configuration the TCP session is reestablished. The MACs from the cluster peers will be revalidated and programmed accordingly. Brocade recommends shutting down all the CCEP ports on the cluster node so that there is graceful failover and then hitless operation can be performed.
Port MAC Security on the node where it is programmed.	Hitless Upgrade is NOT supported, on the Brocade MLX series, and NetIron XMR, however it is compatible. If the operation is performed with cluster configuration the TCP session is reestablished. The MACs from the cluster peers will be revalidated and programmed accordingly. Brocade recommends shutting down all the CCEP ports on the cluster node so that there is graceful failover and then hitless operation can be performed.

<p>802.1x on the node where it is programmed.</p>	<p>Multi-port MAC are not supported on ICL or CCEP ports. Configuration will be rejected when trying to configure multi-port MAC addresses with a port mask which contains either a CCEP port or ICL port and vice-versa on the Brocade MLXe, NetIron MLX and NetIron XMR.</p>
<p>Static MAC configuration - Static MACs are programmed on both local and remote peers as static entries.</p>	<p>Multi-port ARP will not be allowed on ICL or CCEP ports on the Brocade MLX series and NetIron XMR.</p>

Show lag_ecmp_port command

This new CLI command takes the incoming port on which the traffic is entering and the flow type and their relevant parameters as input. It identifies the actual outgoing port of the particular flow, based on the hashing mechanism used in the XPP.

The lag port cannot be identified if per packet load balancing is enabled on that trunk.

Handling switched traffic

For switched traffic flows, the Layer 2 parameters and a valid *lag_id* are to be supplied as input. The lag port is identified based on the trunk hash calculated for input parameters and the *lag_id* supplied by the user.

The supplied destination MAC address must not match the input ports MAC address for L2 switched and MPLS switched traffic flows or an error message will be displayed.

Handling routed traffic

For routed traffic, the destination MAC of the traffic flow must match the MAC address of the input port or VRRP MAC address or an error message will be displayed.

For routed traffic, the *lag_id* should be specified as zero, if route to be identified. When the *lag_id* is a non zero value, it is assumed that the ECMP path is known and only the trunk output port is calculated and displayed. The IP address specified will not be validated for route entries.

For routed IP flows, when the IP destination address is provided as an input, the next hop details are fetched based on the destination IP address and the ECMP hash value is calculated from the input parameters. The next hop details will have the outgoing trunk or port information. When the next hop outgoing interface is a LAG, then the output port is identified using the trunk hash calculated from the input parameters.

For MPLS flows, the ECMP path is selected based on the outer MPLS label.

All the flow parameters must be specified through the CLI configuration. The fields can be specified as zero if those fields are masked in the load balancing configuration.

Show lag_ecmp_port command

Hash-diversify and hash rotate options

The CLI command also takes the hash-diversify and hash rotate values as part of the CLI command. When these values are specified in the command, then these values are used for calculating the ECMP or Trunk hash instead of using the real values used in the XPP. This helps in identifying the impact of changes to these values in XPP, instead of changing them.

Syntax `show lag_ecmp_port [inport incoming port | lag_id lag_id | flowtype [I2_switched_non_ip | I2_mpls_switched | ipv4 | ipv4_ipv4 | ipv4_gre_ipv4 | ipv6_ipv4_ipv6 | ipv4_gre_ipv6 | gre_mpls | gre | macinmac] [src_mac | dst_mac | eth_type | vlan_id | inn_vlan_id | isid | label | inn_src_mac | inn_dst_mac | inn_vlan_id | src_ip | dst_ip | inn_src_ip | inn_dst_ip | proto | sport | dport] hash-diversify val hash-rotate val`

Parameters inport

Input port number of the traffic.

lag-id

Egress LAG ID of the traffic.

flowtype

Traffic type for the incoming traffic.

I2_switched_non-IP

This traffic type is for switched traffic. The payload must not be IP packets. The destination MAC address must not match the ports MAC address.

1. *src_mac* (Source MAC value)
2. *dst_mac* (Destination MAC value)
3. *eth_type* (Ethernet type following first Tag)
4. *vlan_id* (VLAN ID)
5. *inn_vlan_id* (Inner VLAN ID)

I2_mpls_switched

This traffic type is for switched traffic. The MPLS labels are optional. It should be provided if not masked in the load balancing configuration. The payload must not be IP packets. The destination MAC address must not match the ports MAC address.

1. *src_mac* (Source MAC value)
2. *dst_mac* (Destination MAC value)
3. *eth_type* (Ethernet type following first Tag)
4. *vlan_id* (VLAN ID)
5. *inn_vlan_id* (Inner VLAN ID)
6. **label** *lbl0* (MPLS Label 0, outermost label)
 - *lbl1* (MPLS Label 1)
 - *lbl2* (MPLS Label 2)
 - *lbl3* (MPLS Label 3)

ipv4

This traffic type is for routed and switched IP traffic. When the destination MAC address matches the ports MAC or VRRP MAC, then the traffic is considered as routed traffic. Otherwise, the traffic is considered as switching traffic.

1. *src_mac* (Source MAC value)
2. *dst_mac* (Destination MAC value)
3. *vlan_id* (VLAN ID)
4. *src_ip* (IPv4 source address)
5. *dst_ip* (IPv4 destination address)
6. **proto** (Protocol)
7. **sport** (TCP/UDP source port)
8. **dport** (TCP/UDP destination port)

ipv4_ipv4 and *ipv4_gre_ipv4*

This traffic type is for tunneled IPv4 traffic flows. The traffic type *ipv4_gre_ipv4* is used for IPv4 over GRE and *ipv4_ipv4* is used for IPv4 over IPv4 tunnel. When the destination MAC address matches the ports MAC or VRRP MAC, then the traffic is considered as routed traffic. Otherwise, the traffic is considered as switching traffic. The outer destination IP address specified must have a valid routing entry for the routing flows.

1. *src_mac* (Source MAC value)
2. *dst_mac* (Destination MAC value)
3. *vlan_id* (VLAN ID)
4. *src_ip* (IPv4 source address)
5. *dst_ip* (IPv4 destination address)
6. *inn_src_ip* (Inner IPv4 source address)
7. *inn_dst_ip* (Inner IPv4 destination address)
8. *inn_proto* (Protocol)
9. **sport** (TCP/UDP source port)
10. **dport** (TCP/UDP destination port)

ipv6

This traffic type is for routed IPv6 traffic. When the destination MAC address matches the ports MAC or VRRP MAC, then the traffic is considered as routed traffic. Otherwise, the traffic is considered as switching traffic.

The destination IPv6 address specified must have a valid routing entry for routing flows.

1. *src_mac* (Source MAC value)
2. *dst_mac* (Destination MAC value)
3. *vlan_id* (VLAN ID)
4. *src_ip* (IPv6 source address)
5. *dst_ip* (IPv6 destination address)

Show lag_ecmp_port command

6. **proto** (Next Header)
7. **sport** (TCP/UDP source port)
8. **dport** (TCP/UDP destination port)

ipv4_ipv6 and *ipv4_gre_ipv6*

This traffic type is for tunneled IPv6 traffic flows. The traffic type *ipv4_gre_ipv6* is used for IPv6 over GRE and *ipv4_ipv6* is used for IPv6 over IPv4 tunnel. When the destination MAC address matches the ports MAC or VRRP MAC, then the traffic is considered as routed traffic. Otherwise, the traffic is considered as switching traffic. The outer destination IP address specified must have a valid routing entry for routing flows.

1. *src_mac* (Source MAC value)
2. *dst_mac* (Destination MAC value)
3. *vlan_id* (VLAN ID)
4. *src_ip* (IPv4 source address)
5. *dst_ip* (IPv4 destination address)
6. *inn_src_ip* (IPv6 source address)
7. *inn_dst_ip* (IPv6 destination address)
8. *inn_proto* (Next Header)
9. **sport** (TCP/UDP source port)
10. **dport** (TCP/UDP destination port)

mpls

This traffic type is used for router MPLS traffic flow. The outer MPLS label must have valid entry in MPLS forwarding table. The destination MAC address of the packet must match the ports MAC address or the VRRP MAC address.

For inner payload details, the input parameters are based on the *mpls_payload_type* parameters. For IP type only, IP address and port details are required.

1. *src_mac* (Source MAC value)
2. *dst_mac* (Destination MAC value)
3. *vlan_id* (VLAN ID)
4. **label lbl0** (MPLS Label 0, outermost label)
 - *lbl1* (MPLS Label 1)
 - *lbl2* (MPLS Label 2)
 - *lbl3* (MPLS Label 3)
5. *mpls_payload_type* (IPv4/IPv6//ETH/ETH_ITAG)
6. *inn_src_mac* (Inner SA MAC for L2VPN)
7. *inn_dst_mac* (Inner DA MAC for L2 VPN)
8. *inn_vlanid* (Inner VLAN ID)
9. *isid* (ISID)
10. **ipv4/ipv6**

- *src_ip* (IPv4/IPv6 source address)
- *dst_ip* (IPv4/IPv6 destination address)

11. **proto** (Protocol/next Header)
12. **sport** (TCP/UDP Source port)
13. **dport** (TCP/UDP Destination port)

gre_mpls

The traffic type is used for routed GRE egress and transit traffic flows. The outer destination IP address specified must have a valid routing entry. The destination MAC address specified must match the ports MAC address or VRRP MAC address configured.

1. *src_mac* (Source MAC value)
2. *dst_mac* (Destination MAC value)
3. *vlan_id* (VLAN ID)
4. *dst_ip* (GRE IPv4 destination address)
5. **label** *lbl0* (MPLS Label 0, outermost label)
 - *lbl1* (MPLS Label 1)
 - *lbl2* (MPLS Label 2)
 - *lbl3* (MPLS Label 3)

mac-in-mac

1. *src_mac* (Source MAC address value)
2. *dst_mac* (Destination MAC address value)
3. *vlan_id* (VLAN ID)
4. *isid* (ISID)
5. *inn_src_mac* (Inner source MAC)
6. *inn_dst_mac* (Inner destination MAC)
7. **ipv4/ipv6**
 - *src_ip* (IPv4/IPv6 source address)
 - *dst_ip* (IPv4/IPv6 destination address)
8. **proto** (Protocol/Next Header)

Modes This command operates in all modes.

Usage guidelines This CLI command can be used identify the output trunk port and ECMP interface of the given traffic flow. The flow type and other relevant packet fields are taken as CLI command input.

[Table 25](#) provides the mappings between the traffic types and the associated flow_types to be used for the **lag_ecmp_port** commands.

Show lag_ecmp_port command

TABLE 25 Traffic type and flow type for the **show lag-ecmp-port** command

Number	Traffic type	Flow type for CLI command
1	L2 switched traffic with non IP payload	l2_switched_non_ip
2	L2 switched MPLS packet	l2_switched_mpls
3	L2 switched traffic with IPv4 payload (header following Ethernet header)	ipv4
4	L2 switched with IPv6 payload (header following Ethernet header)	ipv6
5	IPv4 routed traffic	ipv4
6	Ipv6 routed traffic	ipv6
7	IPv4 over IPv4 tunnel	ipv4_ipv4
8	IPv4 payload over GRE tunnel (IPv4)	ipv4_gre_ipv4
9	IPv6 router traffic	ipv6
10	Ipv6 over IPv4 tunnel	ipv4_ipv6
11	Ipv6 over GRE tunnel (IPv4)	ipv4_gre_ipv6
12	Routed MPLS traffic with inner payload as IP (L2VPN)	mpls
13	Routed MPLS traffic with inner payload is Ethernet (L2VPN)	mpls
14	Routed MPLS traffic with inner payload is Ethernet ITAG (L2VPN)	mpls
15	Routed MPLS traffic over GRE tunnel (IPv4)	gre_mpls
16	MAC in MAC traffic flow	macinmac

Error Conditions

- When both the destination IP address and the *lag_id* are not provided as input, the CLI command throws an error.
- When the destination IP address provided did not have a valid next hop entry, or when the *lag_id* specified is zero.
- When the destination MAC matches the ports MAC address for switched traffic flows.
- When the destination MAC does not match the ports MAC address for routed traffic.
- When the lag ID supplied is not a deployed LAG.
- When the MPLS label supplied for the routed traffic does not have a valid cross connect entry.
- When per packet load balancing is enabled on the line card.

History

Release	Command history
Multi-Service NetIron Release 05.4.00	This command was introduced.

Related commands None.

Match command

The following match command information is an update to the Policy-based Routing (PBR), chapter 32.

Match command options

The Match command utilizes the following command options..

Match command options

TABLE 26

Command option	Function
as-path	Match IP AS Path Access Lists
community	Match IP Community Access Lists
extcommunity	Match BGP/VPN extended community list
interface	Match interface list (maximum 5 interfaces per match clause)
ip	Match ip specific information
ipv6	Match ipv6 specific information
metric	Match metric of route
protocol	Match route on protocol type and sub-type
route-type	Match route-type of route
tag	Match tag of route

NOTE

Unless noted, the access-control lists used are IPv4 access-control lists.

Command examples

```
Brocade (config)# match as-path
```

Syntax: [no] match

[as-path | community | extcommunity | interface | ip | ipv6 | metric | protocol | route-type | tag]

```
Brocade (config)#
```

Syntax: [no] match as-path <ACL1> <ACL2> <ACL3> <ACL4> <ACL5>

```
Brocade (config)# match community exact-match
```

Syntax: [no] match community <ACL1> <ACL2> <ACL3> <ACL4> <ACL5> [exact-match]

```
Brocade (config)# match extcommunity
```

Syntax: [no] match extcommunity <ACL1> <ACL2> <ACL3> <ACL4> <ACL5>

Match command

Brocade (config)# match interface

Syntax: [no] match interface [ethernet|loopback|null0|tunnel|ve] <interface1>
 [ethernet|loopback|null0|tunnel|ve] <interface2> [ethernet|loopback|null0|tunnel|ve]
 <interface3> [ethernet|loopback|null0|tunnel|ve] <interface4>
 [ethernet|loopback|null0|tunnel|ve] <interface5>

Brocade(config)# match ip next-hop

Syntax: [no] match ip [address|next-hop|route-source] [prefix-list] <ACL1> <ACL2> <ACL3>
 <ACL4> <ACL5>

Brocade (config)# match ipv6 next-hop

Syntax: [no] match ipv6 [address|next-hop|route-source] [prefix-list] <ACL1> <ACL2> <ACL3>
 <ACL4> <ACL5>

ACLs for the match IPv6 statement should be IPv6 ACLs.

Brocade (config)# match metric

Syntax: [no] match metric <Route metric>

Brocade (config)# match protocol bgp external

Syntax: [no] match protocol bgp [external|internal|static-network]

Brocade (config)# match protocol isis level-1

Syntax: [no] match protocol isis [level-1|level-2]

Brocade (config)# match protocol rip

Syntax: [no] match protocol rip

Brocade (config)# match protocol static

Syntax: [no] match protocol static

Brocade (config)# match route-type

Syntax: [no] match route-type [external-type1|external-type2|internal|level-1|level-2]

Brocade (config)# match tag

Syntax: [no] match tag <tag1> <tag2> <tag3> <tag4> <tag5> <tag6> <tag7> <tag8> <tag9>
 <tag10> <tag11> <tag12> <tag13> <tag14> <tag15> <tag16>

Documentation updates for Multi-Service IronWare Diagnostic Guide

Management module diagnostics

The management modules control Brocade NetIron XMR and Brocade MLX series hardware components, run networking protocols, and provide the Real Time Operating System (RTOS).

Each chassis requires one management module, and can accept a second module for redundancy that works in conjunction with the active management module. If the active management module becomes unavailable, the redundant management module automatically takes over the system operation, minimizing system downtime.

Running management module diagnostics

You can run diagnostics on the management modules to check if the devices needed for proper operation are accessible and in working order. The diagnostics for the Line Processor (LP) modules begin after the completion of diagnostics for the management processor (MP) modules, if the LP modules are present in the chassis.

MP module is considered to have passed the diagnostics if the result of all the checks is "Passed". If an MP or an LP does not pass the **diag burn-in** command, contact Brocade Technical Support for further assistance.

NOTE

Remove the standby management module from the chassis before running the diagnostics. If the standby management module is present, running the diagnostics on the interface module fails.

To run diagnostics on management modules, perform the following steps.

1. Reload the system and immediately press the **B** key repeatedly until the system boots into monitor mode.
2. Type **boot os flash primary** to enter the OS.
The prompt will change from MP Monitor> to MP OS>.
3. From the MP OS> prompt, enter **diag burn-in**, as shown in the following example.

```
MP-1 OS>diag burn-in
PCI access                - Passed
88E1145 PHY               - Passed
Storage Card              - Passed
M41T11 RTC                - Passed
  FE (slot 0; FE 0; 0x11fe6000) access passed;
  FE (slot 0; FE 1; 0x11fe6000) access passed;
  FE (slot 0; FE 2; 0x11fe6000) access passed;

  FE (slot 1; FE 0; 0x11fe6000) access passed;
  FE (slot 1; FE 1; 0x11fe6000) access passed;
  FE (slot 1; FE 2; 0x11fe6000) access passed;
```

Management module diagnostics

```
SAND access - Passed
Valere power Supply 0 Passed
Valere power Supply 1 Passed
Power Supply access - Passed
  Port 0 passed
  Port 1 passed
  Port 2 passed
  Port 3 passed
  Port 4 passed
  Port 5 passed
  Port 6 passed
  Port 7 passed
  Port 8 passed
Port 9 passed
  Port 10 passed
  Port 11 passed
  Port 12 passed
  Port 13 passed
  Port 14 passed
  Port 15 passed
  Port 16 passed
  Port 17 passed
  Port 18 passed
  Port 19 passed
  Port 23 passed
Dx246 Switch Port Loopback - Passed

###- PASS -###
MP-1 OS>
LP (6) [MLX-X 1Gx24 Copper] burn-in started
LP (6) PING test passed
LP (7) [MLX-X 1Gx24 Copper] burn-in started
LP (7) PING test passed

LP (6) (MLX-X 1Gx24 Copper) diagnostic Passed
LP (7) (MLX-X 1Gx24 Copper) diagnostic Passed
```

```
###- PASS -###
```

NOTE

After the completion of diagnostics for the MP modules, the system displays the MP-1 OS> prompt and then starts the diagnostics for the LP modules.

NOTE

Brocade requires that you remove physical connections to all ports on the module, and all optics to all ports on the module, so the module does not receive traffic while the diagnostics are running.

4. Enter the **reset** command to return the system to normal operation (system reboot).

```
MP-1 OS>reset

REBOOT S1: NI-XMR-1Gx20-SFP 20-port 1GbE/100FX Module CARD_STATE_REBOOT 20 0000.003d.8500
BOOT S1: NI-XMR-1Gx20-SFP 20-port 1GbE/100FX Module CARD_STATE_BOOT 20 0000.003d.8500
CARD_STATE_UP S1: NI-XMR-1Gx20-SFP 20-port 1GbE/100FX Module CARD_STATE_SW_LOADED 20
0000.003d.8500
UP S1: NI-XMR-1Gx20-SFP 20-port 1GbE/100FX Module CARD_STATE_UP 20 0000.003d.8500
```

After the system reboots, you can display the status of the module using the **show module** command, as shown in the following example.

```
BigIron# show module
      Module                               Status                               Ports  Starting MAC
M1 (upper): NI-XMR-MR Management Module  Active
M2 (lower):
F1: NI-X-SF Switch Fabric Module        Active
F2: NI-X-SF Switch Fabric Module        Active
F3: NI-X-SF Switch Fabric Module        Active
F4: NI-X-SF Switch Fabric Module        Active
S1: NI-XMR-1Gx20-SFP 20-port 1GbE/100FX Module  CARD_STATE_SW_LOADED    20    0000.003d.8500
S2:
S3: NI-XMR-1Gx20-SFP 20-port 1GbE/100FX Module  CARD_STATE_UP           20    0000.003d.8550
```

Management module diagnostics

Documentation updates for Unified IP MIB Reference

RFC 4293: Management Information Base for the Internet Protocol (IP)

RFC 4293, Management Information Base for the Internet Protocol (IP) obsoletes the following:

- RFC 2011: SNMPv2 Management Information Base for the Internet Protocol using SMIv2
- RFC 2465: Management Information Base for IP Version 6: Textual Conventions and General Group
- RFC 2466: Management Information Base for IP Version 6: ICMPv6 Group

This RFC is supported on the Brocade NetIron XMR, Brocade NetIron MLX, Brocade MLXe, Brocade NetIron CES, and Brocade NetIron CER series devices.

The following table summarizes the tables from the RFC that are supported.

Object group name	Object identifier	Supported IP version	Access
IP scalar variables	1.3.6.1.2.1.4	IPv4 and IPv6	Only the following objects have read-write access: <ul style="list-style-type: none"> • ipDefaultTTL • ipv6IpDefaultHopLimit • ipv6IpForwarding All other scalar variables are read-only.
ipv4InterfaceTable	1.3.6.1.2.1.4.28	IPv4	All objects are read-only.
ipv6InterfaceTable	1.3.6.1.2.1.4.30	IPv6	All objects are read-only.
ipSystemStatsTable			
ipSystemStatsInOctets	1.3.6.1.2.1.4.31.1.1.5	None	Always returns 0.
ipSystemStatsHCInOctets	1.3.6.1.2.1.4.31.1.1.6	None	Always returns 0.
ipSystemStatsInAddrErrors	1.3.6.1.2.1.4.31.1.1.9	IPv4	IPv6 returns 0.
ipSystemStatsInUnknownProtos	1.3.6.1.2.1.4.31.1.1.10	IPv4	IPv6 returns 0.
ipSystemStatsInTruncatedPkts	1.3.6.1.2.1.4.31.1.1.11	IPv6	IPv4 returns 0.
ipSystemStatsInDiscards	1.3.6.1.2.1.4.31.1.1.17	IPv4	IPv6 returns 0.
ipSystemStatsOutNoRoutes	1.3.6.1.2.1.4.31.1.1.22	IPv4	IPv6 returns 0.
ipSystemStatsOutFragReqds	1.3.6.1.2.1.4.31.1.1.26	IPv4	IPv6 returns 0.
ipSystemStatsOutFragFails	1.3.6.1.2.1.4.31.1.1.28	IPv4	IPv6 returns 0.
ipSystemStatsOutTransmits	1.3.6.1.2.1.4.31.1.1.30	IPv4	IPv6 returns 0.

Object group name	Object identifier	Supported IP version	Access
ipSystemStatsHCOutTransmits	1.3.6.1.2.1.4.31.1.1.31	IPv4	IPv6 returns 0.
ipSystemStatsOutOctets	1.3.6.1.2.1.4.31.1.1.32	None	Always returns 0.
ipSystemStatsHCOutOctets	1.3.6.1.2.1.4.31.1.1.33	None	Always returns 0.
ipSystemStatsInMcastPkts	1.3.6.1.2.1.4.31.1.1.34	None	Always returns 0.
ipSystemStatsHCInMcastPkts	1.3.6.1.2.1.4.31.1.1.35	None	Always returns 0.
ipSystemStatsInMcastOctets	1.3.6.1.2.1.4.31.1.1.36	None	Always returns 0.
ipSystemStatsHCInMcastOctets	1.3.6.1.2.1.4.31.1.1.37	None	Always returns 0.
ipSystemStatsOutMcastPkts	1.3.6.1.2.1.4.31.1.1.38	None	Always returns 0.
ipSystemStatsHCOutMcastPkts	1.3.6.1.2.1.4.31.1.1.39	None	Always returns 0.
ipSystemStatsOutMcastOctets	1.3.6.1.2.1.4.31.1.1.40	None	Always returns 0.
ipSystemStatsHCOutMcastOctets	1.3.6.1.2.1.4.31.1.1.41	None	Always returns 0.
ipSystemStatsInBcastPkts	1.3.6.1.2.1.4.31.1.1.42	None	Always returns 0.
ipSystemStatsHCInBcastPkts	1.3.6.1.2.1.4.31.1.1.43	None	Always returns 0.
ipSystemStatsOutBcastPkts	1.3.6.1.2.1.4.31.1.1.44	None	Always returns 0.
ipSystemStatsHCOutBcastPkts	1.3.6.1.2.1.4.31.1.1.45	None	Always returns 0.
ipSystemStatsDiscontinuityTime	1.3.6.1.2.1.4.31.1.1.46	None	Always returns 0.
ipIfStatsTableLastChange	1.3.6.1.2.1.4.31.2	IPv4 and IPv6	All objects are read-only.
ipIfStatsTable			
ipIfStatsInOctets	1.3.6.1.2.1.4.31.3.1.5	None	Always returns 0.
ipIfStatsHCInOctets	1.3.6.1.2.1.4.31.3.1.6	None	Always returns 0.
ipIfStatsInHdrErrors	1.3.6.1.2.1.4.31.3.1.7	IPv6	IPv4 returns 0.
ipIfStatsInNoRoutes	1.3.6.1.2.1.4.31.3.1.8	IPv6	IPv4 returns 0.
ipIfStatsInAddrErrors	1.3.6.1.2.1.4.31.3.1.9	IPv6	IPv4 returns 0.
ipIfStatsInUnknownProtos	1.3.6.1.2.1.4.31.3.1.10	IPv6	IPv4 returns 0.
ipIfStatsInTruncatedPkts	1.3.6.1.2.1.4.31.3.1.11	IPv6	IPv4 returns 0.
ipIfStatsInForwDatagrams	1.3.6.1.2.1.4.31.3.1.12	IPv4	IPv6 returns 0.
ipIfStatsReasmReqds	1.3.6.1.2.1.4.31.3.1.14	IPv6	IPv4 returns 0.
ipIfStatsReasmOKs	1.3.6.1.2.1.4.31.3.1.15	IPv6	IPv4 returns 0.
ipIfStatsReasmFails	1.3.6.1.2.1.4.31.3.1.16	IPv6	IPv4 returns 0.
ipIfStatsInDiscards	1.3.6.1.2.1.4.31.3.1.17	IPv6	IPv4 returns 0.
ipIfStatsInDelivers	1.3.6.1.2.1.4.31.3.1.18	IPv6	IPv4 returns 0.
ipIfStatsHCInDelivers	1.3.6.1.2.1.4.31.3.1.19	IPv6	IPv4 returns 0.
ipIfStatsOutRequests	1.3.6.1.2.1.4.31.3.1.20	IPv6	IPv4 returns 0.
ipIfStatsHCOutRequests	1.3.6.1.2.1.4.31.3.1.21	IPv6	IPv4 returns 0.
ipIfStatsOutForwDatagrams	1.3.6.1.2.1.4.31.3.1.23	IPv6	IPv4 returns 0.

RFC 4293: Management Information Base for the Internet Protocol (IP)

Object group name	Object identifier	Supported IP version	Access
ipIfStatsHCOuForwDatagrams	1.3.6.1.2.1.4.31.3.1.24	IPv6	IPv4 returns 0.
ipIfStatsOutDiscards	1.3.6.1.2.1.4.31.3.1.25	IPv6	IPv4 returns 0.
ipIfStatsOutFragReqds	1.3.6.1.2.1.4.31.3.1.26	None	Always returns 0.
ipIfStatsOutFragOKs	1.3.6.1.2.1.4.31.3.1.27	IPv6	IPv4 returns 0.
ipIfStatsOutFragFails	1.3.6.1.2.1.4.31.3.1.28	IPv6	IPv4 returns 0.
ipIfStatsOutFragCreates	1.3.6.1.2.1.4.31.3.1.29	IPv6	IPv4 returns 0.
ipIfStatsOutTransmits	1.3.6.1.2.1.4.31.3.1.30	IPv4	IPv6 returns 0.
ipIfStatsHCOuTransmits	1.3.6.1.2.1.4.31.3.1.31	IPv4	IPv6 returns 0.
ipIfStatsOutOctets	1.3.6.1.2.1.4.31.3.1.32	None	Always returns 0.
ipIfStatsHCOuOctets	1.3.6.1.2.1.4.31.3.1.33	None	Always returns 0.
ipIfStatsInMcastPkts	1.3.6.1.2.1.4.31.3.1.34	IPv6	IPv4 returns 0.
ipIfStatsHCInMcastPkts	1.3.6.1.2.1.4.31.3.1.35	IPv6	IPv4 returns 0.
ipIfStatsInMcastOctets	1.3.6.1.2.1.4.31.3.1.36	None	Always returns 0.
ipIfStatsHCInMcastOctets	1.3.6.1.2.1.4.31.3.1.37	None	Always returns 0.
ipIfStatsOutMcastPkts	1.3.6.1.2.1.4.31.3.1.38	IPv6	IPv4 returns 0.
ipIfStatsHCOuMcastPkts	1.3.6.1.2.1.4.31.3.1.39	IPv6	IPv4 returns 0.
ipIfStatsOutMcastOctets	1.3.6.1.2.1.4.31.3.1.40	None	Always returns 0.
ipIfStatsHCOuMcastOctets	1.3.6.1.2.1.4.31.3.1.41	None	Always returns 0.
ipIfStatsInBcastPkts	1.3.6.1.2.1.4.31.3.1.42	None	Always returns 0.
ipIfStatsHCInBcastPkts	1.3.6.1.2.1.4.31.3.1.43	None	Always returns 0.
ipIfStatsOutBcastPkts	1.3.6.1.2.1.4.31.3.1.44	None	Always returns 0.
ipIfStatsHCOuBcastPkts	1.3.6.1.2.1.4.31.3.1.45	None	Always returns 0.
ipIfStatsDiscontinuityTime	1.3.6.1.2.1.4.31.3.1.46	None	Always returns 0.
ipAddressPrefixTable	1.3.6.1.2.1.4.32	IPv4 and IPv6	All objects are read-only.
ipAddressTable	1.3.6.1.2.1.4.34	IPv4 and IPv6	All objects are read-only.
ipNetToPhysicalTable NOTE: Only ARP entries that are currently being used are included in the ARP table.	1.3.6.1.2.1.4.35	IPv4 and IPv6	Only the following objects have read-create access: <ul style="list-style-type: none"> ipNetToPhysicalPhysAddress ipNetToPhysicalType ipNetToPhysicalRowStatus All other objects are read-only.
ipV6ScopeZoneIndexTable	1.3.6.1.2.1.4.36	IPv6	All objects are read-only.
ipDefaultRouterTable	1.3.6.1.2.1.4.37	IPv4 and IPv6	All objects are read-only.

Object group name	Object identifier	Supported IP version	Access
Ipv6RouterAdvertTable	1.3.6.1.2.1.4.39	IPv6	Only the following objects have read-write access; all others are read-only: <ul style="list-style-type: none"> • ipv6RouterAdvertSendAdverts • ipv6RouterAdvertManagedFlag • ipv6RouterAdvertOtherConfigFlag • ipv6RouterAdvertReachableTime • ipv6RouterAdvertRetransmitTime • ipv6RouterAdvertCurHopLimit • ipv6RouterAdvertDefaultLifetime
icmpStatsTable	1.3.6.1.2.1.5.29	IPv4 and IPv6	All objects are read-only.
icmpMsgStatsTable	1.3.6.1.2.1.5.30	IPv4 and IPv6	All objects are read-only.

Fabric drop count

The Brocade NetIron MLX, Brocade MLXe, and Brocade NetIron XMR devices are provided with Simple Network Management Protocol (SNMP) Management Information Base (MIB) support for the fabric drop count. The fabric drop counters are maintained by the system and are updated automatically whenever there is a packet drop at switch fabric level. The `brcdFabricStatsTable` contains information of Switch Fabric Module (SFM) related information specific to the Brocade NetIron MLX, Brocade MLXe, and Brocade NetIron XMR devices.

NOTE

The following `brcdFabricStatsTable` is supported only on the High-speed SFM (HSFM) cards. The table support GET and GET-NEXT requests.

Name, OID, and syntax	Access	Description
<code>brcdFabricStatsTable</code> <code>brcdIp.1.1.13.1.1.1</code>	None	The <code>brcdFabricStatsTable</code> contains information of various SFM counters supported by the system.
<code>brcdFabricSfmId</code> <code>brcdIp.1.1.13.1.1.1.1</code> Syntax: Unsigned32	None	The SFM ID.
<code>brcdFabricSfmFeld</code> <code>brcdIp.1.1.13.1.1.1.1.2</code> Syntax: Unsigned32	None	The Fabric Element (FE) ID.
<code>brcdFabricDropMAC0Count</code> <code>brcdIp.1.1.13.1.1.1.1.3</code> Syntax: Counter32	Read-only	The number of packets dropped for MAC0 (links 0 through 23) link group.
<code>brcdFabricDropMAC1Count</code> <code>brcdIp.1.1.13.1.1.1.1.4</code> Syntax: Counter32	Read-only	The number of packets dropped for MAC1 (links 24 through 47) link group.
<code>brcdFabricDropMAC2Count</code> <code>brcdIp.1.1.13.1.1.1.1.5</code> Syntax: Counter32	Read-only	The number of packets dropped for MAC2 (links 48 through 71) link group.
<code>brcdFabricDropMAC3Count</code> <code>brcdIp.1.1.13.1.1.1.1.6</code> Syntax: Counter32	Read-only	The number of packets dropped for MAC3 (links 72 through 95) link group.

brcdNPCSRAMErrorTable (to query for NP CSRAM errors)

The brcdNPCSRAMErrorTable displays information of Network Processor (NP) Control Static Random Access Memory (CSRAM) MIB objects.

NOTE

The following MIB objects are supported on the Brocade NetIron CES and Brocade NetIron CER series devices.

Name, OID, and syntax	Access	Description
brcdNPCSRAMErrorTable brcdIp.1.14.2.1.1.4	None	The table contains information of various Network Processor (NP) CSRAM error event counters supported by the system. The objects in this table are refreshed every second, based on request. This table is only supported on CES/CER.
brcdNPCSRAMErrorSlotId brcdIp.1.14.2.1.1.4.1.1 Syntax: Unsigned32	None	Slot-ID of the module that uniquely identifies it in the system. The module must be a UP and physically present. This is an 1-based index.
brcdNPCSRAMErrorDeviceId brcdIp.1.14.2.1.1.4.1.2 Syntax: Unsigned32	None	The Network Processor device-ID. A number that uniquely identifies the NP within a module in the system. This is an 1-based index.
brcdNPCSRAMErrorDescription brcdIp.1.14.2.1.1.4.1.3 Syntax: DisplayString	Read-only	Details the range of ports serviced by the NP identified by brcdNPCSRAMErrorSlotId and brcdNPCSRAMErrorDeviceId objects.
brcdNPCSRAMErrorCurrentEvents brcdIp.1.14.2.1.1.4.1.4 Syntax: Counter32	Read-only	Counter for NP CSRAM errors recorded within a configured window.
brcdNPCSRAMErrorCumulativeEvents brcdIp.1.14.2.1.1.4.1.5 Syntax: Counter32	Read-only	Counter for total NP CSRAM errors recorded.

brcdNPLPMRAMErrorTable (to query for NP LPM-RAM errors)

NOTE

The following MIB objects are supported only on the Brocade NetIron CES and Brocade NetIron CER series devices.

Name, OID, and syntax	Access	Description
brcdNPLPMRAMErrorTable brcdIp.1.14.2.1.1.5	None	A list of brcdNPLPMRAMError entries. The table contains information of various LPM RAM error event counters supported by the Network processor in the system. The objects in the table are refreshed every second, based on the request.
brcdNPLPMRAMErrorIndex brcdIp.1.14.2.1.1.5.1.1 Syntax: Unsigned32	None	This object uniquely identifies a LPM within a Network Processor. Brocade NetIron CER series devices use LPM-0, LPM-1 and LPM-2 memories, whereas Brocade NetIron CES devices use LPM-3 memory. This is an 1-based index. Index value of 1 maps to LPM0, 2 maps to LPM1 and so on.
brcdNPLPMRAMErrorSlotId brcdIp.1.14.2.1.1.5.1.1 Syntax: Unsigned32	None	Slot-ID of the module that is uniquely identifies it in the system. The module must be a UP and physically present. This is an 1-based index.
brcdNPLPMRAMErrorDeviceld brcdIp.1.14.2.1.1.5.1.1 Syntax: Unsigned32	None	The Network Processor device-ID. A number that uniquely identifies the NP within a module in the system. This is an 1-based index.
brcdNPLPMRAMErrorName brcdIp.1.14.2.1.1.5.1.1 Syntax: DisplayString	Read-only	Details a string representing the LPM identified by brcdNPLPMRAMErrorIndex.
brcdNPLPMRAMErrorDescription brcdIp.1.14.2.1.1.5.1.1 Syntax: DisplayString	Read-only	Details the range of ports serviced by the NP identified by brcdNPLPMRAMErrorSlotId and brcdNPLPMRAMErrorDeviceld objects.
brcdNPLPMRAMErrorCurrentEvents brcdIp.1.14.2.1.1.5.1.1 Syntax: Counter32	Read-only	Counter for the error events recorded within a configured window in the LPM identified by brcdNPLPMRAMErrorIndex, brcdNPLPMRAMErrorSlotId, and brcdNPLPMRAMErrorDeviceld objects.
brcdNPLPMRAMErrorCumulativeEvents brcdIp.1.14.2.1.1.5.1.1 Syntax: Counter32	Read-only	Counter for the error events recorded within a configured window in the LPM identified by brcdNPLPMRAMErrorIndex, brcdNPLPMRAMErrorSlotId, and brcdNPLPMRAMErrorDeviceld objects.

Traps

The following new traps are added to report the CSRAM and LPMRAM errors on the Brocade NetIron CES and Brocade NetIron CER series devices.

Trap name and number	Object ID	Severity	Description
brcdNPCSRAMErrorThresholdExceeded brcdIp.1.14.2.0.3	brcdNPCSRAMErrorDescription, brcdNPCSRAMErrorCurrentEvents	Alerts	The SNMP trap that is generated when the Network Processor CSRAM error event count within a window exceeds the configured threshold. Sample syslog message: NP CSRAM has 4 error events, exceeding configured threshold for interfaces 1/1 to 1/24.
brcdNPLPMRAMErrorThresholdExceeded brcdIp.1.14.2.0.4	brcdNPLPMRAMErrorName, brcdNPLPMRAMErrorDescription, brcdNPLPMRAMErrorCurrentEvents	Alerts	The SNMP trap that is generated when the Network Processor LPMRAM error event count within a window exceeds the configured threshold. Sample syslog message: NP LPM 1 has 4 error events, exceeding configured threshold for interfaces 1/1 to 1/24.

Agent board table

Updated snAgentBrdIndex (OID brodIp.1.1.2.2.1.1.1) object index range from 1 through 42 in the snAgentBrdTable.

Agent board table

Documentation Updates for the MLX Series and NetIron XMR Series Hardware Installation Guide

In this chapter

The updates in this chapter are for the *Brocade MLX Series and Brocade NetIron XMR Hardware Installation Guide*, publication number 53-1002424-03, published May 2012.

- “4-slot router”

For determining the number of power supplies required for redundancy, refer to Chapter 7, “Hardware Specifications”.

- “8-slot router”

Because power is supplied over a common power bus, any power supply installed in addition to the minimum required provides backup for any supply that fails. For power redundancy, you must purchase additional power supplies depending on how you populate your router. For determining the number of power supplies required for redundancy, refer to Chapter 7, “Hardware Specifications”.

- “16-slot router”

Because power is supplied over a common power bus, any power supply installed in addition to the minimum required provides backup for any power supply that fails. For power redundancy, you must purchase additional power supplies depending on how you populate your router. For determining the number of power supplies required for redundancy, refer to Chapter 7, “Hardware Specifications”.

- “32-slot router”

Because power is supplied over a common power bus, any power supply installed in addition to the minimum required provides backup for any power supply that fails. For power redundancy, you must purchase additional power supplies depending on how you populate your router. For determining the number of power supplies required for redundancy, refer to Chapter 7, “Hardware Specifications”.

- “[Brocade MLXe router power consumption values](#)”

TABLE 1 Brocade MLXe router power consumption values

Model	@100 VAC			@200 VAC			@-48VDC			Minimum number of 1200W power supplies needed	Minimum number of 1800W power supplies needed	Minimum number of 2400W power supplies needed	Minimum number of 3000W power supplies needed
	Amps	Watts	BTU/hr	Amps	Watts	BTU/hr	Amps	Watts	BTU/hr				
MAXIMUM PER MLXe (using 8x10G-D, 8x10G-M, 4x10G, 2x10G, 1G modules only)													
MLXe-4	17	1730	5905	9	1730	5905	36	1730	5905	2	1		
MLXe-8	34	3356	11453	17	3356	11453	70	3356	11453	3	2		
MLXe-16	57	5698	19446	28	5698	19446	119	5698	19446	4	3		

4 100xGbE 2-port interface module

TABLE 1 Brocade MLXe router power consumption values (Continued)

Model	@100 VAC			@200 VAC			@-48VDC			Minimum number of 1200W power supplies needed	Minimum number of 1800W power supplies needed	Minimum number of 2400W power supplies needed	Minimum number of 3000W power supplies needed
	Amps	Watts	BTU/hr	Amps	Watts	BTU/hr	Amps	Watts	BTU/hr				
MLXe-32	N/A	N/A	N/A	57	11414	38958	238	11414	38958			4	4
MAXIMUM PER MLXe (any module)													
MLXe-4	21	2083	7108	10	2083	7108	43	2083	7108	2	1		
MLXe-8	41	4060	13858	20	4060	13858	85	4060	13858	3	2		
MLXe-16	71	7107	24255	36	7107	24255	148	7107	24255	5	4		
MLXe-32	N/A	N/A	N/A	71	14232	48575	297	14232	48575			5	4

100xGbE 2-port interface module

NOTE

The following section corrects typographical error from 1.5 MB to 1.5 GB, as stated below.

The 100xGbE 2-port interface module supports 1.5 GB buffering per port.

Cooling system and fans

NOTE

The following section corrects a typographical error. Fan speeds should not be changed.

The following statement from the manual is void. "If desired, you can change the settings of the temperature thresholds associated with fan speed devices."

High-speed switch fabric modules

The following note text was removed:

NOTE

Hot-swapping or installing high-speed switch fabric modules one-at-a-time may affect the performance of 2x100GbE interface modules. If you have 2x100GbE modules in your device, it is recommended that you bring your system down and install all required high-speed switch fabric modules before you bring the system back up.

Switch fabric modules

The following table note is added to the “blinking” state of the switch fabric module LED in the Product Overview chapter of the Brocade MLX Series and Brocade NetIron XMR Hardware Installation Guide.

TABLE 2 Switch fabric module LEDs

LED	Position	State	Meaning
Active	Below Pwr LED	On (4-, 8-, and 16-slot routers only)	The switch fabric is on (active) and ready to switch user packets.
		On (32-slot routers only)	The switch fabric is on (active) and ready to switch user packets.
		Blinking (32-slot routers only)	The switch fabric is on (active) and being accessed by the Management Module CPU. This indicates normal operation. NOTE: On devices supporting software version R05.3.00 and earlier, when you insert an SFM or during powering on the device, the Active LED was off for a short duration, up to 15 seconds because the monitoring of the Fabric module is stopped for this duration. After this delay, the LED indicated the monitoring status. In version R05.4.00 and later, the Active LED reads the switch fabric continuously even during module insertion or powering on the device, and thus the Active LED blinks.
		Off for extended period	The switch fabric is not active and cannot switch user packets.

Managing Routers and Modules

NOTE

Wait at least 10 seconds before issuing the **power-off** command and the **power-on** command.

Maintenance and Field Replacement

NOTE

Wait at least 10 seconds before issuing the **power-off** command and the **power-on** command.

Cable specifications



CAUTION

Before plugging a cable to any port, be sure to discharge the voltage stored on the cable by touching the electrical contacts to ground surface.

Enabling and disabling management module CPU usage calculations

Removed the following note:

NOTE

When you are finished gathering statistics for debugging purposes, it is recommended that you disable the usage averaging calculations, which are CPU-intensive and can affect the performance of the management module.

Installing NIBI-16-FAN-EXH-A fan assemblies

NOTE

This updated section applies to the Brocade MLX Series and Brocade NetIron XMR Hardware Installation Guide. The figure titled “High-speed fan assemblies for Brocade MLX 16-slot routers” was removed because it pictured an MLXe fan assembly.

NOTE

You can remove and replace fan assemblies and switch fabric modules while the Brocade MLX 16-slot router is powered on and running.



CAUTION

To avoid overheating of the router, remove one fan assembly at a time, and replace it promptly. Do not remove all fans from the device at once.

To install the high speed fans, you need the following:

- Two high-speed fan assemblies.
- A small flat-blade screwdriver.
- An ESD wrist strap with a plug for connection to the ESD connector on the router chassis.



DANGER

For safety reasons, the ESD wrist strap should contain a 1 megohm series resistor.


Perform these steps to replace standard rear fan assemblies with high-speed rear fan assemblies in a 16-slot router.

1. Put on the ESD wrist strap and ground yourself by inserting the plug into the ESD connector on the front of the device.
2. Using the flat-blade screwdriver, loosen the four captive screws that secure each fan assembly to the device.
3. Remove the standard fan assemblies by inserting your fingers underneath each assembly and pulling it toward you as shown in [Figure 1](#).

**DANGER**

Be careful not to insert your fingers into the fan while removing it. The fan may still be spinning at a high speed.


FIGURE 1 Removing a standard fan assembly from the Brocade MLX-16-slot router



- 1 Power connector
 - 2 Fan assembly
 - 3 Captive screws
4. Insert the high-speed fan assembly into the fan slot and push the assembly in until the connection plate is flush with the device. Pushing the assembly in seats the fan connector with the device connector. Refer to [Figure 2](#).

4 Installing NIBI-16-FAN-EXH-A fan assemblies

FIGURE 2 Installing high-speed fan assemblies



5. Secure each high-speed fan assembly to the device by tightening the four captive screws.
6. Access the CLI and enter the **show chassis** command to verify that all fans are operating normally.

```
Brocade# show chassis
*** Brocade 16-slot Chassis ***

---POWERS ---
Power 1 (32015000 - AC 1200W): Installed (OK)
Power 2 (32015000 - AC 1200W): Installed (OK)
Power 3 (H1250CFN - AC 1200W): Installed (OK)
Power 4 (32015000 - AC 1200W): Installed (OK)
Power 5: not present
Power 6: not present
Power 7 (32015000 - AC 1200W): Installed (OK)
Power 8 (32015000 - AC 1200W): Installed (OK)
Total power budget for chassis = 7200 W
Total power used by system core = 762 W
Total power used by LPs = 3788 W
Total power available = 2650 W
Slot Power-On Priority and Power Usage
Slot1 pri=1 module type=NI-MLX-10Gx8-M 8-port 10GbE (M) Module power usage=246W
Slot2 pri=1 module type=NI-XMR-10Gx4 4-port 10GbE Module power usage=245W
Slot3 pri=1 module type=NI-MLX-10Gx4 4-port 10GbE Module power usage=225W
Slot4 pri=1 module type=NI-MLX-10Gx8-M 8-port 10GbE (M) Module power usage=246W
Slot5 pri=1 module type=NI-MLX-10Gx8-M 8-port 10GbE (M) Module power usage=246W
Slot6 pri=1 module type=NI-MLX-10Gx4 4-port 10GbE Module power usage=225W
Slot7 pri=1 module type=NI-MLX-10Gx4 4-port 10GbE Module power usage=225W
Slot8 pri=1 module type=NI-MLX-10Gx8-M 8-port 10GbE (M) Module power usage=246W
Slot9 pri=1 module type=NI-MLX-10Gx8-M 8-port 10GbE (M) Module power usage=246W
Slot10 pri=1 module type=NI-MLX-10Gx4 4-port 10GbE Module power usage=225W
```

```
Slot11 pri=1 module type=NI-MLX-10Gx4 4-port 10GbE Module power usage=225W
Slot12 pri=1 module type=NI-MLX-10Gx8-M 8-port 10GbE (M) Module power usage=246W
Slot13 pri=1 module type=NI-MLX-10Gx8-M 8-port 10GbE (M) Module power usage=246W
Slot14 pri=1 module type=NI-MLX-10Gx4 4-port 10GbE Module power usage=225W
Slot15 pri=1 module type=NI-MLX-10Gx4 4-port 10GbE Module power usage=225W
Slot16 pri=1 module type=NI-MLX-10Gx8-M 8-port 10GbE (M) Module power usage=246W
```

```
--- FANS ---
```

```
Bottom fan tray (fan 1): Status = OK, Speed = HI (100%)
Bottom fan tray (fan 2): Status = OK, Speed = HI (100%)
Bottom fan tray (fan 3): Status = OK, Speed = HI (100%)
Bottom fan tray (fan 4): Status = OK, Speed = HI (100%)
Bottom fan tray (fan 5): Status = OK, Speed = HI (100%)
Bottom fan tray (fan 6): Status = OK, Speed = HI (100%)
Rev A Back Fan A (revision 0x09): Status = OK, Speed = HI (100%)
Rev A Back Fan B (revision 0x09): Status = OK, Speed = HI (100%)
```

```
--- TEMPERATURE READINGS ---
```

```
Active Mgmt Module: 40.250C 60.875C
Standby Mgmt Module: 34.500C 821384312.545015484C
SNM1: FE1:50.0C FE2:50.0C FE3:50.625C
SNM2: FE1:46.375C FE2:43.875C FE3:41.750C
SNM3: FE1:48.875C FE2:47.375C FE3:48.250C
SNM4: FE1:47.375C FE2:42.375C FE3:39.875C
LP1 Sensor1: 47.625C
LP1 Sensor2: 51.625C
LP1 Sensor3: UNUSED
LP1 Sensor4: 51.375C
LP1 Sensor5: 48.500C
LP1 Sensor6: UNUSED
LP1 Sensor7: 183.0C
```

```
.
.
.
```

```
Fans are in auto mode (current speed is HI (100%)). Temperature monitoring poll
period is 60 seconds.
```

```
--- MISC INFO ---
```

```
Backplane EEPROM MAC Address: 000c.dbdf.6400
```

This output shows firmware Revision A (Rev A) for NIBI-16-FAN-EXH-A fan modules. Rev A indicates that the Brocade MLX-16 contains the required rear fan modules to support NI-MLX-10Gx8 and NI-MLX-1Gx48-T-A modules. The RPM value thresholds (LOW/MED/MED-HI/HI) are also displayed for rear fan modules.

NOTE

If the Brocade MLX-16 does not contain NIBI-16-FAN-EXH-A modules, the **show chassis** command output will not list Rev A for rear fan modules.

Module identification correction

BR-MLX-1GFx24-X 24 XMR/MLXE 24-port 1-GbE (X) (SFP) Module with IPv4/IPv6/MPLS hardware support. Supports 1M IPv4 routes in FIB.

The mentioned module is a Fiber SFP module and not a RJ45 module.

Brocade MLX Series and NetIron XMR supplemental upgrade procedures

Upgrading MBRIDGE or MBRIDGE32 images on management modules

NOTE

Always use TELNET on the MLX-32 chassis (instead of SSH). PROM write operations consume substantial CPU cycles, starving other tasks such as SSH. The end result includes timeouts within affected tasks. TELNET does not have similar issues (i.e. hello exchanges) and hence is not impacted.
