

53-1002805-03
24 January, 2014



Brocade MLX Series and NetIron Family

Documentation Updates

Supporting Multi-Service IronWare R05.6.xx

BROCADE

Copyright © 2014 Brocade Communications Systems, Inc. All Rights Reserved.

Brocade, Brocade Assurance, the B-wing symbol, BigIron, DCX, Fabric OS, FastIron, MLX, NetIron, SAN Health, ServerIron, TurboIron, 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

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

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 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>	53-1002805-03	NetIron 05.6.00b Release updates.	24 January, 2014

Contents

About This Document

In this chapter	v
How this document is organized	v
Brocade resources	vi
Getting technical help	vi
Document feedback	vi

Chapter 1

Documentation Updates for the Multi-Service IronWare Configuration Guides

In this chapter	1
Configuring a “null” route	1
ACL deny logging	2
Unsupported features for Brocade NetIron CES and Brocade NetIron CER devices	2
Deployment Scenarios and CLI Configuration	2
Telemetry Solutions	3
Scaling limitations	3
Base-line configuration of telemetry solutions	4
Global level configuration	5

Chapter 2

Documentation updates for Multi-Service IronWare Diagnostic Guide

Chapter 3

Documentation updates for Unified IP MIB Reference

In this chapter	11
Route map configuration table	11
MAC filters	11
RFC 4444: Management Information Base for Intermediate System to Intermediate System (IS-IS)	11
Scalar isisSys objects	11

Chapter 4

Documentation Updates for the MLXe / MLX Series and NetIron XMR Series Hardware Installation Guide

In this chapter	13
Switch fabric modules	13
Brocade MLXe Series	13
Brocade MLX Series and Brocade NetIron XMR	14

About This Document

In this chapter

- [“How this document is organized”](#) on page v
- [“Brocade resources”](#) on page vi
- [“Getting technical help”](#) on page vi
- [“Document feedback”](#) on page vi

How this document is organized

This document contains updates to the Multi-Service IronWare R05.6.00a product manuals. These updates include document fixes and changes covering new features. [Table 1](#) below list the most recently released Multi-Service IronWare R05.6.00a product manuals.

TABLE 1 Documentation supporting Multi-Service IronWare R05.6.00a

Publication Title	Fabric OS Release	Publication Date
<i>Multi-Service IronWare Administration Configuration Guide</i>	R05.6.00a and later	December 2013
<i>Multi-Service IronWare Multiprotocol Label Switch (MPLS) Configuration Guide</i>	R05.6.00a and later	December 2013
<i>Multi-Service IronWare IP Multicast Configuration Guide</i>	R05.6.00a and later	December 2013
<i>Multi-Service IronWare Routing Configuration Guide</i>	R05.6.00a and later	December 2013
<i>Multi-Service IronWare Software Defined Networking (SDN) Configuration Guide</i>	R05.6.00a and later	December 2013
<i>Multi-Service IronWare Security Configuration Guide</i>	R05.6.00a and later	December 2013
<i>Multi-Service IronWare Switching Configuration Guide</i>	R05.6.00a and later	December 2013
<i>Multi-Service IronWare QoS and Traffic Management Configuration Guide</i>	R05.6.00a and later	December 2013
<i>Brocade MLXe Series Hardware Installation Guide</i>	R05.6.00a and later	December 2013
<i>Brocade MLX Series and NetIron XMR Hardware Installation Guide</i>	R05.6.00a and later	December 2013
<i>Brocade NetIron CES Series and NetIron CER Series Hardware Installation Guide</i>	R05.6.00a and later	December 2013
<i>Multi-Service IronWare Software Upgrade Guide</i>	R05.6.00a and later	December 2013
<i>Brocade MLX Series and NetIron XMR Diagnostics Guide</i>	R05.6.00a and later	December 2013
<i>Unified IP MIB Reference</i>	R05.6.00a and later	December 2013
<i>Brocade MLX Series and NetIron XMR YANG Guide</i>	R05.6.00a and later	December 2013

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 Multi-Service IronWare Configuration Guides

In this chapter

The updates in this chapter are for the following *Multi-Service IronWare R05.6.00 Configuration Guides*.

- Multi-Service Ironware Switching Configuration Guide - publication number 53-1003036-03
- Multi-Service Ironware Security Configuration Guide - publication number 53-1003035-03

The following features were added or modified as part of the 5.6.00b release.

- [“Configuring a “null” route”](#) on page 1
- [“ACL deny logging”](#) on page 2
- [“Deployment Scenarios and CLI Configuration”](#) on page 2
- [“Telemetry Solutions”](#) on page 3

Configuring a “null” route

The following section is an update to the Configuring IP Chapter in the *Multi-Service Ironware Switching Configuration Guide*.

The feature support table is updated for the “Dropping Traffic Sent to the Null0 Interface in Hardware” feature.

TABLE 1 Feature support table

Features supported	Brocade NetIron XMR	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
Dropping Traffic Sent to the Null0 Interface in Hardware	Yes	Yes	Yes	Yes	Yes	Yes	Yes

The following note is added in the “Dropping traffic sent to the null0 interface In hardware” section.

NOTE

The `ip hw-drop-on-def-route` command is not supported on the Brocade NetIron CES and Brocade NetIron CER devices. You can drop traffic sent to the default IP route address in hardware without the `ip hw-drop-on-def-route` command.

ACL deny logging

The following section is an update to the Configuring an IPv6 Access Control List Chapter in the *Multi-Service Ironware Security Configuration Guide*.

ACL deny logging is supported on the Brocade NetIron CES and Brocade NetIron CER devices but not in conjunction with acl accounting, hence updating this section by removing the bullet point “ACL deny logging is not supported”.

Unsupported features for Brocade NetIron CES and Brocade NetIron CER devices

The following features are not supported on the Brocade NetIron CES and Brocade NetIron CER devices:

- The **acl-outbound exclude-switched-traffic** command to exclude switched traffic from outbound ACL filtering is not supported.
- The **acl-frag-conservative** command to change the operation of ACLs on fragmented packets is not supported.
- The **suppress-rpf-drop** command to suppress RPF packet drops for a specific set of packets using inbound ACLs is not supported.
- For all NetIron devices, if a port has an IPv4 or IPv6 ACL applied, you must remove the ACL bindings before adding that port to a VLAN that has a VE interface.

Deployment Scenarios and CLI Configuration

The following section is an update to the Provider Backbone Bridging (PBB) Networks for the Brocade NetIron XMR and the Brocade MLX series Chapter in the *Multi-Service Ironware Switching Configuration Guide*.

In the Configuration for CE Devices section, under Configuration for PE Devices, the S-VLAN tag-type is 0x9100 and not 0x900.

Telemetry Solutions

The following section is an update to the Telemetry Solutions Chapter in the *Multi-Service Ironware Administration Guide*.

The update provides information about recommended baseline configuration and scaling limitations for telemetry solutions.

Scaling limitations

- 400 (IPv4 and IPv6 combined) route-map instances per interface.
 - Valid instance is a route-map instance with the permit option and with a valid ACL (ACL is present in configuration).
 - Exceeding this limit results in first come first applied behavior on the port.
 - User should redesign their route-map if this limit is exceeded for proper functioning.
- 200 IPv6 ACLs
 - 20480 IPv6 clauses that can be present in the configuration.
- IPv4 ACL limitations have not changed.
- At maximum scale, this configuration may take up to 30 to 45 minutes to bind ACLs used in the route-maps to the ingress interfaces. Traffic is flooded to all VLAN 1 ports during that time.
- User should execute the show cam-partition usage command under the Rule item, to check if it will accommodate the application of the route-map on the desired number of ports on each tower.
- Usage of transparent-hw-flooding (TVF) and transparent-hw-flooding lag-load-balancing(TVF LAG LDB) is best effort, and may result in data loss for bursty streams.
- Usage of per-packet load balancing on LAGs used for TVF LAG LDB is not supported.
- Dynamic and keep-alive LAGs are not supported with TVF LAG LDB.
- If the SFMs are operating in “normal mode”, the number of TVF LAG LDB instances must not exceed the following values. Run the **show vlan tvf-lag-lb** command:

TABLE 2 Configured System Max Values

tvf-lag-lb-fid-group	tvf-lag-lb-fid-pool	Do not exceed value
2	512	170
2	1024	341
2	2048	682
4	512	102
4	1024	204
4	2048	409
8	512	56
8	1024	113
8	2048	227
16	512	30
16	1024	60
16	2048	120

Configuration examples

Base-line configuration of telemetry solutions

```
no spanning-tree
no dual-mode-default-vlan
```

NOTES: Default VLAN must have TVF enabled as shown.

```
vlan 1 name DEFAULT-VLAN
no untagged ethe 13/1 to 13/3
transparent-hw-flooding
```

NOTES: Egress VLANs must have the following as shown.

- 1 A port present
- 2 TVF or TVF LAG LDB enabled
- 3 Port must be enabled
- 4 Port must be in the up state

```
vlan 1000 name Outer_1000
tagged ethe 13/1
transparent-hw-flooding
```

```
vlan 1001 name Outer_1001
tagged ethe 13/2
transparent-hw-flooding
```

```
vlan 1002 name Outer_1002
tagged ethe 13/3
transparent-hw-flooding
```

Global level configuration

Configuring System max and cam-partition

```
system-max vlan 4095
system-max virtual-interface 4095
system-max ip-filter-sys 40960
system-max receive-cam 512
system-max ipv4-mcast-cam 512
system-max ipv6-mcast-cam 512
cam-partition profile ipv4-ipv6
```

Disabling LFS at global level

```
no link-fault-signaling
link-fault-signaling ignore-rx
link-fault-signaling ignore-rx device-1
```

NOTE

These commands prevent link-fault-signaling (LFS) from taking the tap ports offline due to LFS on the monitored links

Configuring Ingress tap port

```
interface ethernet 1/1
enable
ip policy route-map Outer_Mall
ipv6 policy route-map Outer_Mall
allow-all-vlan pbr
gig-default neg-off
mac access-group Deny_Any out
```

NOTE

gig-default neg-off is required to be configured only for 1G fiber ports.

Configuring Egress port

```
interface ethernet 13/1
enable
link-fault-signaling
link-fault-signaling ignore-rx

interface ethernet 13/2
enable
link-fault-signaling
link-fault-signaling ignore-rx

interface ethernet 13/3
enable
link-fault-signaling
link-fault-signaling ignore-rx
```

NOTE

LFS must be enabled on Egress 10G ports.

Configuring ACL

```

ipv6 access-list v6_Mall_Outer_1001
permit ipv6 host 667:a6db:39c5:f217:4374:435e:ba5e:d402 any

ipv6 access-list v6_Mall_Outer_1002
permit ipv6 host 849e:958:ed:bcd8:577d:5468:edef:8dfc any

ipv6 access-list v6_Mall_Outer_1000
permit ipv6 host 2f12:4a71:704c:8a1a:7de3:7ef9:43a9:550a any

ipv6 access-list v6_Permit_Any
permit ipv6 any any

ip access-list extended v4_Mall_Outer_1001
permit ip host 95.64.50.180 any

ip access-list extended v4_Mall_Outer_1002
permit ip host 126.126.14.76 any

ip access-list extended v4_Mall_Outer_1000
permit ip host 117.218.157.45 any

ip access-list extended v4_Permit_Any
permit ip any any

mac access-list Deny_Any
deny any any any

```

NOTE

For this application always set the ACL rule as “permit”.

NOTE

The only exception to this rule is, the last route-map instance must be set as CATCH-ALL, to avoid all unmatched traffic going to the CPU for forwarding. The only exception is if you have another routing protocol which picks up the unmatched traffic, and allows the usage of deny statement in the ACLs and no need to set CATCH-ALL. All “denied” and unmatched packets will be passed to the routing protocol for forwarding. Traffic to be dropped is handled at the end of the route-map.

Configuring Route-map

```

route-map Outer_Mall permit 1000
rule-name 1000
match ip address v4_Mall_Outer_1000
match ipv6 address v6_Mall_Outer_1000
set next-hop-flood-vlan 1000
set interface null0
route-map Outer_Mall permit 1001
rule-name 1001
match ip address v4_Mall_Outer_1001
match ipv6 address v6_Mall_Outer_1001
set next-hop-flood-vlan 1001
set interface null0
route-map Outer_Mall permit 1002
rule-name 1002
match ip address v4_Mall_Outer_1002
match ipv6 address v6_Mall_Outer_1002
set next-hop-flood-vlan 1002
set interface null0

```

```
route-map Outer_Mall permit 10000
rule-name Catch_All
match ip address v4_Permit_Any
match ipv6 address v6_Permit_Any
set interface null0

end
```

Configuration consideration for Route-map

Route-map instances (The complete **route-map blah permit | deny xxx** configuration section) and route-map configuration must meet the following conditions:

1. The last set of commands must be interface null0, this can be preceded by multiple set of other commands. This prevents the matched traffic from going to the CPU for forwarding, when the egress VLAN is not a valid next hop.
2. Rule names can only be used once per route-map.
3. The last route-map instance must be set as CATCH-ALL, to avoid all unmatched traffic going to the CPU for forwarding. The only exception is if you have another routing protocol which picks up the unmatched traffic, and allows the usage of deny statement in the ACLs and no need to set CATCH-ALL. All denied and unmatched packets will be passed to the routing protocol for forwarding.

Documentation updates for Multi-Service IronWare Diagnostic Guide

NOTE

There are no updates for NetIron 5.6.00b.

Documentation updates for Unified IP MIB Reference

In this chapter

The updates in this chapter are for the *Unified IP MIB Reference*, published December 2013.

Route map configuration table

Name, OID, and syntax	Access	Description
brcdRouteMapRuleName brcdIp.1.1.3.39.1.1.1.1.4 Syntax: DisplayString	Read-create	Identifies the path name for the route map. A maximum of 127 characters is allowed.
NOTE: This object is not supported on the Brocade NetIron CES and Brocade NetIron CER series devices.		

MAC filters

NOTE

MAC filter MIB objects are not supported on the Brocade NetIron XMR, Brocade MLX series, Brocade NetIron CES, and Brocade NetIron CER series devices.

RFC 4444: Management Information Base for Intermediate System to Intermediate System (IS-IS)

Scalar isisSys objects

Object group name	Object identifier	Supported?	Notes
isisSysMaxPathSplits	1.3.6.1.2.1.138.1.1.1.4	Yes	Default value is 4 on the Brocade NetIron devices.

Documentation Updates for the MLXe / MLX Series and NetIron XMR Series Hardware Installation Guide

In this chapter

The updates in this chapter are for the following publications:

- Brocade MLXe Series Hardware installation Guide - publication number 53-1003030-02
- Brocade MLX Series and Brocade NetIron XMR Hardware Installation Guide - publication number 53-1003040-02

Switch fabric modules

Brocade MLXe Series

The following table note is added to the “blinking” state of the switch fabric module LED in the Product Overview chapter of the Brocade MLXe Series Hardware Installation Guide.

TABLE 1 Switch fabric module LEDs

LED	Position	State	Meaning
Pwr	Above Active LED	On	The module is receiving power.
		Off	The module is not receiving power.
Active	Below Pwr LED	On (4-, 8-, and 16-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.

Brocade MLX Series and Brocade NetIron XMR

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 Series Hardware Installation Guide.

TABLE 2 Switch fabric module LEDs

LED	Position	State	Meaning
Pwr	Above Active LED	On	The module is receiving power.
		Off	The module is not receiving power.
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.