BROCADE

CONFIGURATION GUIDE

Brocade Monitoring and Alerting Policy Suite Configuration Guide

Supporting Fabric OS 8.0.1

53-1004121-01
22 April 2016

Contents

Preface **7**
- Document conventions 7
  - Text formatting conventions 7
  - Command syntax conventions 7
  - Notes, cautions, and warnings 8
- Brocade resources 8
- Contacting Brocade Technical Support 8
  - Brocade customers 8
  - Brocade OEM customers 9
- Document feedback 9

About This Document **11**
- Supported hardware and software 11
  - Brocade Gen 5 (16-Gbps) fixed-port switches 11
  - Brocade Gen 5 (16-Gbps) DCX 8510 Directors 11
  - Brocade Gen 6 fixed-port switches 11
  - Brocade Gen 6 Directors 12
- What's new in this document 12
  - Changes made for this release 12
- MAPS commands altered in this release 13
- MAPS rules and groups altered in this release 13
- Glossary 16

Monitoring and Alerting Policy Suite Overview **19**
- MAPS overview 19
- MAPS license requirements 19
- MAPS activation 20
- MAPS configuration 20
  - Deleting a user-created MAPS configuration 20
- MAPS interaction with other Fabric OS features 20
- Restrictions on MAPS monitoring 20
- Firmware upgrade and downgrade considerations for MAPS 21
  - Firmware upgrade considerations for MAPS 21
  - Firmware downgrade considerations for MAPS 22
- Features that do not require a Fabric Vision license 24
  - Monitors that do not require a Fabric Vision license 24
  - MAPS commands that do not require a Fabric Vision license 24

MAPS Setup and Operation **25**
- Initial MAPS setup 25
  - Activating MAPS without a Fabric Vision license 25
  - Activating MAPS with a Fabric Vision license 26
  - Quickly monitoring a switch with predefined policies 28
- Monitoring across different time windows 28
- Setting the active MAPS policy to a default policy 29
- Pausing MAPS monitoring 30
- Resuming MAPS monitoring 30

MAPS Elements and Categories 31
MAPS structural elements 31
MAPS monitoring categories 31
Port Health 32
Back-end Health 33
FRU Health 34
Security Violations 34
Fabric State Changes 34
Switch Resource 35
Traffic Performance 36
FCIP Health 36
Fabric Performance Impact 37
Switch Status Policy 38
Security certificate monitoring 38
MAPS Groups, Conditions, Rules, and Policies 41
MAPS groups overview 41
Viewing group information 41
Predefined groups 42
User-defined groups 44
Cloning a group 48
Deleting groups 48
MAPS conditions 49
Threshold values 49
Timebase 49
MAPS rules overview 52
MAPS rule actions 52
MAPS policies overview 61
Viewing policy values 62
Predefined policies 62
User-defined policies 63
Fabric Watch legacy policies 63
Working with MAPS policies 64
Automatic creation of MAPS rules and policies 67
Working with MAPS rules and actions 68
MAPS Dashboard 79
MAPS dashboard overview 79
MAPS dashboard sections 79
MAPS dashboard display options 82
Viewing the MAPS dashboard 83
Viewing a summary switch status report 85
Viewing a detailed switch status report 87
Viewing historical data 89
Viewing data for a specific time window 91
Clearing MAPS dashboard data 95
Port Monitoring Using MAPS 97
Monitoring groups of ports using the same conditions 97
Port monitoring using port names 97
Port monitoring using device WWNs 98
Adding a port to an existing static group 98

Adding missing ports to a dynamic group 99
Removing ports from a group 99
D_Port monitoring 100
Back-end port monitoring 102
Dashboard output of back-end port rule violations 103
Port monitoring and pausing 103
Gigabit Ethernet port monitoring 104
GE port monitoring CRC rule creation 105

Monitoring Flow Vision Flows with MAPS 107
Monitoring Flow Vision Flow Monitor data with MAPS 107
Importing flows 108
Adding monitoring flows after importing 108
Monitoring traffic performance 109
Monitoring end-to-end performance 109
Monitoring frames for a specified set of criteria 110
Monitoring learned flows 110
Excessive throughput notification 110
I/O latency monitoring 110
Monitoring I/O latency 111

Fabric performance impact monitoring using MAPS 113
MAPS latency monitoring 113
Frame timeout latency monitoring 115
Transient queue latency counter monitoring 115
Buffer credit zero counter monitoring 117
Latency state clearing 117
Zoned device ratio monitoring 117
MAPS and Bottleneck Detection 119
Port toggling support 120
Slow Drain Device quarantining 121
Slow Drain Device Quarantine licensing 123
Notes on Slow Drain Device Quarantining 123
Enabling Slow Drain Device Quarantining 124
Disabling Slow Drain Device Quarantining 124
Confirming the slow-draining status of a device 125
Displaying quarantined ports 126
Clearing quarantined ports 127
Slow Drain Device quarantining and FICON 128

Other MAPS monitoring capabilities 131
Scalability limit monitoring 131
Layer 2 fabric device connection monitoring 132
LSAN device connection monitoring in a metaSAN 132
Backbone fabric Fibre Channel router count monitoring 132
Zone configuration size monitoring 133
Monitoring NPIV logins to F_Ports 133
Scalability limit monitoring assumptions and dependencies 134
Default rules for scalability limit monitoring 135
Examples of scalability limit rules 135
MAPS Service Availability Module 136
MAPS monitoring for Extension platforms 138

Brocade FCIP monitoring parameters and groups....138
Quality of Service monitoring example....140
IPEXT monitoring....140
IPEXT rule creation....141
E-mail delivery monitoring....142
Fan air-flow direction monitoring....143
Updating monitoring policies for devices with four PSUs....145

MAPS Threshold Values....147
Viewing monitoring thresholds....147
Back-end port monitoring thresholds....148
Fabric state change monitoring thresholds....148
Extension monitoring thresholds....149
FRU state monitoring thresholds....149
Port Health monitoring thresholds....150
D_Port default Port Health monitoring thresholds....150
E_Port default Port Health monitoring thresholds....151
F_Port default Port Health monitoring thresholds....151
Non-F_Port default Port Health monitoring thresholds....152
Resource monitoring thresholds....153
Security monitoring thresholds....153
SFP monitoring thresholds....154
10 Gbps, 16 Gbps, and 32 Gbps SFP monitoring threshold defaults....154
Quad SFPs and all other SFP monitoring threshold defaults....155
Fabric Performance Impact thresholds....155
Switch status policy monitoring thresholds....156
Traffic Performance thresholds....158

Preface

- Document conventions 7
- Brocade resources 8
- Contacting Brocade Technical Support 8
- Document feedback 9

Document conventions

The document conventions describe text formatting conventions, command syntax conventions, and important notice formats used in Brocade technical documentation.

Text formatting conventions

Text formatting conventions such as boldface, italic, or Courier font may be used in the flow of the text to highlight specific words or phrases.

Format	Description
bold text	Identifies command names
	Identifies keywords and operands
	Identifies the names of user-manipulated GUI elements
	Identifies text to enter at the GUI
italic text	Identifies emphasis
	Identifies variables
	Identifies document titles
`Courier font`	Identifies CLI output
	Identifies command syntax examples

Command syntax conventions

Bold and italic text identify command syntax components. Delimiters and operators define groupings of parameters and their logical relationships.

Convention	Description
bold text	Identifies command names, keywords, and command options.
italic text	Identifies a variable.
value	In Fibre Channel products, a fixed value provided as input to a command option is printed in plain text, for example, **--show** WWN.
[ ]	Syntax components displayed within square brackets are optional.
	Default responses to system prompts are enclosed in square brackets.
{ **x** | **y** | **z** }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
	In Fibre Channel products, square brackets may be used instead for this purpose.
x | **y**	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, for example, passwords, are enclosed in angle brackets.

Convention	Description
...	Repeat the previous element, for example, *member*[*member*...].
\	Indicates a "soft" line break in command examples. If a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Notes, cautions, and warnings

Notes, cautions, and warning statements may be used in this document. They are listed in the order of increasing severity of potential hazards.

NOTE
A Note provides a tip, guidance, or advice, emphasizes important information, or provides a reference to related information.

ATTENTION
An Attention statement indicates a stronger note, for example, to alert you when traffic might be interrupted or the device might reboot.

CAUTION
A Caution statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.

DANGER
A Danger statement indicates conditions or situations that can be potentially lethal or extremely hazardous to you. Safety labels are also attached directly to products to warn of these conditions or situations.

Brocade resources

Visit the Brocade website to locate related documentation for your product and additional Brocade resources.

You can download additional publications supporting your product at www.brocade.com. Select the Brocade Products tab to locate your product, then click the Brocade product name or image to open the individual product page. The user manuals are available in the resources module at the bottom of the page under the Documentation category.

To get up-to-the-minute information on Brocade products and resources, go to MyBrocade. You can register at no cost to obtain a user ID and password.

Release notes are available on MyBrocade under Product Downloads.

White papers, online demonstrations, and data sheets are available through the Brocade website.

Contacting Brocade Technical Support

As a Brocade customer, you can contact Brocade Technical Support 24x7 online, by telephone, or by e-mail. Brocade OEM customers contact their OEM/Solutions provider.

Brocade customers

For product support information and the latest information on contacting the Technical Assistance Center, go to http://www.brocade.com/services-support/index.html.

If you have purchased Brocade product support directly from Brocade, use one of the following methods to contact the Brocade Technical Assistance Center 24x7.

Online	Telephone	E-mail
Preferred method of contact for non-urgent issues: • My Cases through MyBrocade • Software downloads and licensing tools • Knowledge Base	Required for Sev 1-Critical and Sev 2-High issues: • Continental US: 1-800-752-8061 • Europe, Middle East, Africa, and Asia Pacific: +800-AT FIBREE (+800 28 34 27 33) • For areas unable to access toll free number: +1-408-333-6061 • Toll-free numbers are available in many countries.	support@brocade.com Please include: • Problem summary • Serial number • Installation details • Environment description

Brocade OEM customers

If you have purchased Brocade product support from a Brocade OEM/Solution Provider, contact your OEM/Solution Provider for all of your product support needs.

- OEM/Solution Providers are trained and certified by Brocade to support Brocade® products.
- Brocade provides backline support for issues that cannot be resolved by the OEM/Solution Provider.
- Brocade Supplemental Support augments your existing OEM support contract, providing direct access to Brocade expertise. For more information, contact Brocade or your OEM.
- For questions regarding service levels and response times, contact your OEM/Solution Provider.

Document feedback

To send feedback and report errors in the documentation you can use the feedback form posted with the document or you can e-mail the documentation team.

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. You can provide feedback in two ways:

- Through the online feedback form in the HTML documents posted on www.brocade.com.
- By sending your feedback to documentation@brocade.com.

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

About This Document

- Supported hardware and software....11
- What's new in this document....12
- MAPS commands altered in this release....13
- MAPS rules and groups altered in this release....13
- Glossary....16

Supported hardware and software

The following hardware platforms are supported by Fabric OS 8.0.1.

NOTE

Although many different software and hardware configurations are tested and supported by Brocade Communication Systems, Inc for Fabric OS 8.0.1, documenting all possible configurations and scenarios is beyond the scope of this document.

Brocade Gen 5 (16-Gbps) fixed-port switches

- Brocade 6505 switch
- Brocade 6510 switch
- Brocade 6520 switch
- Brocade M6505 blade server SAN I/O module
- Brocade 6543 blade server SAN I/O module
- Brocade 6545 blade server SAN I/O module
- Brocade 6546 blade server SAN I/O module
- Brocade 6547 blade server SAN I/O module
- Brocade 6548 blade server SAN I/O module
- Brocade 6558 blade server SAN I/O module
- Brocade 7840 Extension Switch

Brocade Gen 5 (16-Gbps) DCX 8510 Directors

NOTE

For ease of reference, Brocade chassis-based storage systems are standardizing on the term "Director". The legacy term "Backbone" can be used interchangeably with the term "Director".

- Brocade DCX 8510-4 Director
- Brocade DCX 8510-8 Director

Brocade Gen 6 fixed-port switches

- Brocade G620 switch

Brocade Gen 6 Directors

- Brocade X6-4 Director
- Brocade X6-8 Director

Fabric OS support for the Brocade Analytics Monitoring Platform (AMP) device depends on the specific version of the software running on that platform. Refer to the AMP Release Notes and documentation for more information.

What's new in this document

This document includes new and modified information for the Fabric OS 8.0.1 release of MAPS.

Changes made for this release

The following content is new or has been significantly revised for this release of this document.

- Supported hardware and software on page 11
- MAPS commands altered in this release on page 13
- MAPS rules and groups altered in this release on page 13
- Firmware upgrade and downgrade considerations on page 21
- Security Violations on page 34
- Security certificate monitoring on page 38
- Viewing group information on page 41
- Predefined groups on page 42
- MAPS SNMP traps on page 56
- Viewing MAPS rules on page 68
- Creating a rule on page 69
- Creating an exact clone on page 72
- Cloning a rule and changing its values on page 72
- Cloning a rule and changing its timebase on page 72
- Quieting a rule on page 72
- D_Port monitoring on page 100
- Dashboard high-level information section on page 79
- Summary Report section on page 80
- Viewing a summary switch status report on page 85
- Viewing historical data on page 89
- Viewing data for a specific time window on page 91
- Clearing MAPS dashboard data on page 95
- Frame timeout latency monitoring on page 115

- Buffer credit zero counter monitoring on page 117
- Latency state clearing on page 117
- Port toggling support on page 120
- Scalability limit monitoring on page 131
- MAPS monitoring for Extension platforms on page 138
- I/O latency monitoring on page 110
- Zoned device ratio monitoring on page 117
- Updating monitoring policies for devices with four PSUs on page 145
- F_Port default Port Health monitoring thresholds on page 151
- SFP monitoring thresholds on page 154
- Fabric Performance Impact thresholds on page 155
- Switch status policy monitoring thresholds on page 156

MAPS commands altered in this release

The following existing commands have been altered for this release of Fabric OS. Refer to the *Fabric OS Command Reference* for more complete information on modified and deleted commands.

TABLE 1 Altered MAPS commands

Affected Command	Alteration
All Administration Domain (AD) configuration commands	Support for Administration Domains has been deprecated. A warning message will be displayed and a RASLog entry will be generated for any AD configuration commands or if an AD is activated through a command or zone merge.
bottleneckmon	This command has been deprecated.
mapsconfig --action	The keyword **--action** lets you configure actions. However, SW_MARGINAL and SW_CRITICAL actions are no longer valid options, because these actions are already enabled and the user cannot disable them. Also, the keyword no longer requires a license.
mapsconfig --disableFPImon	The keyword **--disableFPImon** is no longer supported in the **mapsconfig** command.
mapsconfig --enableFPImon	The keyword **--enableFPImon** is no longer supported in the **mapsconfig** command.
mapsconfig --show	The keyword **--show** no longer displays FPI status. Also, it no longer requires a license.
mapsdb --show	The keyword **--show** no longer displays latency and congestion errors.

MAPS rules and groups altered in this release

For this release of Fabric OS, several changes have been made to the way MAPS handles rules and groups. You should ensure your policies are using the correct rules and groups.

Rules that are obsolete on various platforms

For this release of Fabric OS, MAPS automatically generates a configuration file for each platform. Therefore, on different platforms, some rules have become obsolete, because they cannot be used on those platforms. The following tables list the rules that have been made obsolete for various platforms.

NOTE

Obsolete rules are automatically removed from user-defined policies, and they will be removed from default policies in future releases.

The following table lists the rules that have been made obsolete for non-FCIP platforms. Obsolete rules are automatically removed from user-defined policies, and they will be removed from default policies in future releases.

TABLE 2 FCIP rules that cannot be used on non-FCIP platforms

defALL_CIRCUITSCIR_PKTLOSS_PER_05	defALL_EXT_GE_PORTSCRC_10
defALL_CIRCUITSCIR_PKTLOSS_PER_1	defALL_EXT_GE_PORTSCRC_20
defALL_CIRCUITSCIR_PKTLOSS_PER_5	defALL_EXT_GE_PORTSCRC_5
defALL_CIRCUITSCIR_STATE_0	defALL_EXT_GE_PORTSINV_LEN_10
defALL_CIRCUITSCIR_STATE_3	defALL_EXT_GE_PORTSINV_LEN_20
defALL_CIRCUITSCIR_STATE_5	defALL_EXT_GE_PORTSINV_LEN_5
defALL_CIRCUITSCIR_UTIL_60	defALL_EXT_GE_PORTSLOS_10
defALL_CIRCUITSCIR_UTIL_75	defALL_EXT_GE_PORTSLOS_20
defALL_CIRCUITSCIR_UTIL_90	defALL_EXT_GE_PORTSLOS_5
defALL_CIRCUITS_JITTER_PER_05	defALL_TUNNELSSTATE_CHG_0
defALL_CIRCUITS_JITTER_PER_15	defALL_TUNNELSSTATE_CHG_1
defALL_CIRCUITS_JITTER_PER_20	defALL_TUNNELSSTATE_CHG_3
defALL_CIRCUITS_RTT_250	defALL_TUNNELSUTIL_PER_50
defALL_CIRCUIT_F_QOS_PKTLOSS_PER_05	defALL_TUNNELSUTIL_PER_75
defALL_CIRCUIT_F_QOS_PKTLOSS_PER_1	defALL_TUNNELSUTIL_PER_90
defALL_CIRCUIT_F_QOS_PKTLOSS_PER_5	defALL_TUNNEL_F_QOS_PKTLOSS_PER_05
defALL_CIRCUIT_F_QOS_UTIL_PER_50	defALL_TUNNEL_F_QOS_PKTLOSS_PER_1
defALL_CIRCUIT_F_QOS_UTIL_PER_75	defALL_TUNNEL_F_QOS_PKTLOSS_PER_5
defALL_CIRCUIT_F_QOS_UTIL_PER_90	defALL_TUNNEL_F_QOS_UTIL_PER_50
defALL_CIRCUIT_HIGH_QOS_PKTLOSS_PER_05	defALL_TUNNEL_F_QOS_UTIL_PER_75
defALL_CIRCUIT_HIGH_QOS_PKTLOSS_PER_1	defALL_TUNNEL_F_QOS_UTIL_PER_90
defALL_CIRCUIT_HIGH_QOS_PKTLOSS_PER_5	defALL_TUNNEL_HIGH_QOS_PKTLOSS_PER_05
defALL_CIRCUIT_HIGH_QOS_UTIL_PER_50	defALL_TUNNEL_HIGH_QOS_PKTLOSS_PER_1
defALL_CIRCUIT_HIGH_QOS_UTIL_PER_75	defALL_TUNNEL_HIGH_QOS_PKTLOSS_PER_5
defALL_CIRCUIT_HIGH_QOS_UTIL_PER_90	defALL_TUNNEL_HIGH_QOS_UTIL_PER_50
defALL_CIRCUIT_LOW_QOS_PKTLOSS_PER_05	defALL_TUNNEL_HIGH_QOS_UTIL_PER_75
defALL_CIRCUIT_LOW_QOS_PKTLOSS_PER_1	defALL_TUNNEL_HIGH_QOS_UTIL_PER_90
defALL_CIRCUIT_LOW_QOS_PKTLOSS_PER_5	defALL_TUNNEL_LOW_QOS_PKTLOSS_PER_05
defALL_CIRCUIT_LOW_QOS_UTIL_PER_50	defALL_TUNNEL_LOW_QOS_PKTLOSS_PER_1
defALL_CIRCUIT_LOW_QOS_UTIL_PER_75	defALL_TUNNEL_LOW_QOS_PKTLOSS_PER_5
defALL_CIRCUIT_LOW_QOS_UTIL_PER_90	defALL_TUNNEL_LOW_QOS_UTIL_PER_50
defALL_CIRCUIT_MED_QOS_PKTLOSS_PER_05	defALL_TUNNEL_LOW_QOS_UTIL_PER_75
defALL_CIRCUIT_MED_QOS_PKTLOSS_PER_1	defALL_TUNNEL_LOW_QOS_UTIL_PER_90
defALL_CIRCUIT_MED_QOS_PKTLOSS_PER_5	defALL_TUNNEL_MED_QOS_PKTLOSS_PER_05
defALL_CIRCUIT_MED_QOS_UTIL_PER_50	defALL_TUNNEL_MED_QOS_PKTLOSS_PER_1
defALL_CIRCUIT_MED_QOS_UTIL_PER_75	defALL_TUNNEL_MED_QOS_PKTLOSS_PER_5
defALL_CIRCUIT_MED_QOS_UTIL_PER_90	defALL_TUNNEL_MED_QOS_UTIL_PER_50
	defALL_TUNNEL_MED_QOS_UTIL_PER_75
	defALL_TUNNEL_MED_QOS_UTIL_PER_90

The following table lists the rules that have been made obsolete for platforms that do not support BE_Ports. Obsolete rules are automatically removed from user-defined policies, and they will be removed from default policies in future releases.

TABLE 3 BE_Port monitoring rules that cannot be used on platforms without BE_Ports

defALL_BE_PORTSBAD_OS_5M_10	defALL_BE_PORTSITW_5M_10
defALL_BE_PORTSBAD_OS_D_100	defALL_BE_PORTSITW_D_100
defALL_BE_PORTSCRC_5M_10	defALL_BE_PORTSLR_5M_10
defALL_BE_PORTSCRC_D_100	defALL_BE_PORTSLR_D_100
defALL_BE_PORTSFRM_LONG_5M_10	defALL_BE_PORTS_LATENCY_CLEAR
defALL_BE_PORTSFRM_LONG_D_100	defALL_BE_PORTS_LATENCY_IMPACT
defALL_BE_PORTSFRM_TRUNC_5M_10	
defALL_BE_PORTSFRM_TRUNC_D_100	

The following table lists the rules that have been made obsolete for non-extension platforms. Obsolete rules are automatically removed from user-defined policies, and they will be removed from default policies in future releases.

TABLE 4 Extension monitoring rules that cannot be used on non-extension platforms

defALL_CIRCUITS_IP_JITTER_PER_05 defALL_CIRCUITS_IP_JITTER_PER_15 defALL_CIRCUITS_IP_JITTER_PER_20 defALL_CIRCUITS_IP_PKTLOSS_P_05 defALL_CIRCUITS_IP_PKTLOSS_P_1 defALL_CIRCUITS_IP_PKTLOSS_P_5 defALL_CIRCUITS_IP_RTT_250 defALL_CIRCUITS_IP_UTIL_P_60 defALL_CIRCUITS_IP_UTIL_P_75 defALL_CIRCUITS_IP_UTIL_P_90 defALL_CIRCUIT_IP_HIGH_QOS_PKTLOSS_P_05 defALL_CIRCUIT_IP_HIGH_QOS_PKTLOSS_P_1 defALL_CIRCUIT_IP_HIGH_QOS_PKTLOSS_P_5:1 defALL_CIRCUIT_IP_HIGH_QOS_UTIL_P_50:1 defALL_CIRCUIT_IP_HIGH_QOS_UTIL_P_75:1 defALL_CIRCUIT_IP_HIGH_QOS_UTIL_P_90:1 defALL_CIRCUIT_IP_LOW_QOS_PKTLOSS_P_05:1 defALL_CIRCUIT_IP_LOW_QOS_PKTLOSS_P_1:1 defALL_CIRCUIT_IP_LOW_QOS_PKTLOSS_P_5:1 defALL_CIRCUIT_IP_LOW_QOS_UTIL_P_50:1 defALL_CIRCUIT_IP_LOW_QOS_UTIL_P_75:1 defALL_CIRCUIT_IP_LOW_QOS_UTIL_P_90 defALL_CIRCUIT_IP_MED_QOS_PKTLOSS_P_05 defALL_CIRCUIT_IP_MED_QOS_PKTLOSS_P_1 defALL_CIRCUIT_IP_MED_QOS_PKTLOSS_P_5	defALL_CIRCUIT_IP_MED_QOS_UTIL_P_50 defALL_CIRCUIT_IP_MED_QOS_UTIL_P_75 defALL_CIRCUIT_IP_MED_QOS_UTIL_P_90 defALL_TUNNELS_IP_UTIL_P_50 defALL_TUNNELS_IP_UTIL_P_75 defALL_TUNNELS_IP_UTIL_P_90 defALL_TUNNEL_IP_HIGH_QOS_PKTLOSS_P_05 defALL_TUNNEL_IP_HIGH_QOS_PKTLOSS_P_1 defALL_TUNNEL_IP_HIGH_QOS_PKTLOSS_P_5 defALL_TUNNEL_IP_HIGH_QOS_UTIL_P_50 defALL_TUNNEL_IP_HIGH_QOS_UTIL_P_75 defALL_TUNNEL_IP_HIGH_QOS_UTIL_P_90 defALL_TUNNEL_IP_LOW_QOS_PKTLOSS_P_05 defALL_TUNNEL_IP_LOW_QOS_PKTLOSS_P_1 defALL_TUNNEL_IP_LOW_QOS_PKTLOSS_P_5 defALL_TUNNEL_IP_LOW_QOS_UTIL_P_50 defALL_TUNNEL_IP_LOW_QOS_UTIL_P_75 defALL_TUNNEL_IP_LOW_QOS_UTIL_P_90 defALL_TUNNEL_IP_MED_QOS_PKTLOSS_P_05 defALL_TUNNEL_IP_MED_QOS_PKTLOSS_P_1 defALL_TUNNEL_IP_MED_QOS_PKTLOSS_P_5 defALL_TUNNEL_IP_MED_QOS_UTIL_P_50 defALL_TUNNEL_IP_MED_QOS_UTIL_P_75 defALL_TUNNEL_IP_MED_QOS_UTIL_P_90

Rules that are replaced with other rules or are automatically deleted from custom policies

The following rules either are replaced by other rules or are automatically deleted from custom policies for certain systems and devices.

Old rule	New rule
defALL_F_PORTSDEV_NPIV_LOGINS_90	defALL_F_PORTSDEV_NPIV_LOGINS_PER_90
defALL_F_PORTSDEV_NPIV_LOGINS_75	defALL_F_PORTSDEV_NPIV_LOGINS_PER_75
defALL_PORTSSFP_STATE_ON	Not used; automatically removed from custom policies for all systems.
defCHASSISBAD_PWR_CRIT	Automatically removed from custom policies for fixed-port switches.

Groups that are obsolete on various platforms

The following groups are removed when they do not apply to a specific platform. Make sure that you do not use these rules in custom rules. If they already exist in custom rules, delete those rules before upgrading.

Removed group	Removed from these platforms
ALL_CIRCUITS	Removed from non-FCIP platforms

Removed group	Removed from these platforms
ALL_BE_PORTS	Removed from platforms that do not support BE_Ports
ALL_CORE_BLADES ALL_SW_BLADES ALL_SLOTS	Removed from systems that are not chassis-based.

Deprecated rules

The following rules have been deprecated in this release, but they are still present in policies in order to support backward compatibility. However, they will be obsoleted and replaced in future releases.

The following table lists the rules that have been deprecated but are still available for backward compatibility. They will **not** be replaced automatically. You should start using the replacement rules listed in Table 6 .

TABLE 5 Deprecated rules that are still available for backward compatibility

defNON_E_F_PORTSLF_0 defNON_E_F_PORTSLF_3 defNON_E_F_PORTSLF_5 defALL_HOST_PORTSLF_0 defALL_HOST_PORTSLF_3 defALL_HOST_PORTSLF_5 defALL_OTHER_F_PORTSLF_0 defALL_OTHER_F_PORTSLF_3 defALL_OTHER_F_PORTSLF_5 defALL_E_PORTSLF_0 defALL_E_PORTSLF_3 defALL_E_PORTSLF_5 defALL_TARGET_PORTSLF_0 defALL_TARGET_PORTSLF_3 defALL_TARGET_PORTSLF_5 defALL_D_PORTSLF_1 defALL_D_PORTSLF_H30 defALL_D_PORTSLF_D500 defALL_D_PORTSLF_2 defALL_D_PORTSLF_H60 defALL_D_PORTSLF_D1000	defALL_D_PORTSLF_3 defALL_D_PORTSLF_H90 defALL_D_PORTSLF_D1500 defNON_E_F_PORTSLOSS_SIGNAL_0 defNON_E_F_PORTSLOSS_SIGNAL_3 defNON_E_F_PORTSLOSS_SIGNAL_5 defALL_HOST_PORTSLOSS_SIGNAL_0 defALL_HOST_PORTSLOSS_SIGNAL_3 defALL_HOST_PORTSLOSS_SIGNAL_5 defALL_OTHER_F_PORTSLOSS_SIGNAL_0 defALL_OTHER_F_PORTSLOSS_SIGNAL_3 defALL_OTHER_F_PORTSLOSS_SIGNAL_5 defALL_E_PORTSLOSS_SIGNAL_0 defALL_E_PORTSLOSS_SIGNAL_3 efALL_E_PORTSLOSS_SIGNAL_5 defALL_TARGET_PORTSLOSS_SIGNAL_0 defALL_TARGET_PORTSLOSS_SIGNAL_3 defALL_TARGET_PORTSLOSS_SIGNAL_5

The following table lists the rules that will replace the deprecated rules in future releases (listed in Table 5). These rules will **not** be automatically replaced for deprecated ones. You should start including these rules in policies instead of the deprecated rules.

TABLE 6 New rules that will replace deprecated rules in future releases

defALL_PORTSLF_0 defALL_PORTSLF_3 defALL_PORTSLF_5 defALL_PORTLOSS_SIGNAL_0 defALL_PORTLOSS_SIGNAL_3 defALL_PORTLOSS_SIGNAL_5

Glossary

The following terminology, acronyms, and abbreviations are used in this document.

TABLE 7 MAPS-related terminology

Term	Description
Action	An activity, such as RASlog, performed by MAPS if a condition defined in a rule evaluates to true.
ASIC	Application-Specific Integrated Circuit; the chip within a switch or blade that controls its operation, including packet switching and routing.
Back-end port	Port that connects a core switching blade to a port or application blade (and vice versa).
Condition	Combination of monitoring system , timebase, threshold value, and the logical operation (examples: ==, >, <, >=) that needs to be evaluated; for example: CRC/MIN > 10.
Congestion	A circumstance in which the offered traffic load exceeds the capacity to transmit on the link. This could occur because the offered or incoming traffic exceeds the capacity of the link or the receiver of the traffic is receiving slower than the incoming traffic.
CRC	Cyclic redundancy check; an error-detecting code used to detect accidental changes to raw data.
Flow	A set of Fibre Channel frames that share similar traits.
Flow Monitor	A Flow Vision application that can monitor all the traffic passing through E_Ports, EX_Ports, F_Ports, and XISL_Ports. Refer to the *Flow Vision Administrator's Guide* for details.
Front-end port	Port that connects a switch to another switch, host, or target device, such as a storage unit.
FRU	Field-Replaceable Unit; Power supply or other physical element of a device that can be replaced by an end-user.
HBA	Host Bus Adapter; Fibre Channel interface that allows a host system to connect to other network or storage devices.
ISL	Inter-Switch Link; a link joining two Fibre Channel switches using E_Ports.
ITW	Invalid transmission word; if a data word does not transmit successfully, encoding errors may result.
Logical group	A collection of similar objects.
NPIV	N_Port ID Virtualization; a feature permitting multiple Fibre Channel node port (N_Port) IDs to share a single physical N_Port.
Policy	Set of rules which are activated at the same time.
Rule	Setting that associates a condition with actions that need to be triggered when the specified condition is evaluated to true.
QoS	Quality of Service; mechanism for assigning priorities to data flows. See Traffic Management.
RSCN	Registered State Change Notification; a Fibre Channel fabric notification sent to all specified switches and nodes identifying significant fabric changes.
SAN	Storage Area Network; a dedicated network that provides access to consolidated, block level data storage.
Slow-drain(ing) device	Device that does not process frames at the rate generated by the source.
Static flow	Flow created when learning is not used and all the parameters required are contained in the flow definition.
Sub-flow	System auto-created flow based on a root flow. A root flow can have more than one sub-flow.
Threshold	Value used in evaluating a condition.
Timebase	The time period across which the change in a counter is to be monitored.
Traffic management	Fabric OS mechanism that assigns high, medium, or low priority to data flows.
VC	Virtual Circuit; a mechanism in a Brocade switch for creating multiple logical data paths across a single physical link or connection.
Virtual channel	Synonym for Virtual Circuit.
XISL	An eXtended ISL; a type of ISL connection that carries traffic for multiple logical fabrics on the same physical link while maintaining fabric isolation.

Monitoring and Alerting Policy Suite Overview

- MAPS overview 19
- MAPS license requirements 19
- MAPS activation 20
- MAPS configuration 20
- MAPS interaction with other Fabric OS features 20
- Restrictions on MAPS monitoring 20
- Firmware upgrade and downgrade considerations for MAPS 21
- Features that do not require a Fabric Vision license 24

MAPS overview

The Monitoring and Alerting Policy Suite (MAPS) is a storage area network (SAN) health monitor supported on all switches running Fabric OS 7.2.0 or later. Monitoring SAN switches enables early fault detection and isolation as well as permitting performance measurements. This allows each MAPS-enabled switch to constantly monitor itself for potential faults and automatically alert you to problems before they become costly failures.

MAPS is implicitly activated when you install Fabric OS 7.4.0 or a later release; however, doing so without a license provides reduced functionality. Refer to MAPS activation on page 20 and Firmware upgrade considerations for MAPS on page 21 for specific details on activating MAPS. When you upgrade to Fabric OS 7.4.0 or a later version, MAPS starts monitoring the switch as soon as the switch is active. Refer to Features that do not require a Fabric Vision license on page 24 for information on these monitors.

When the Fabric Vision license is installed and activated , MAPS provides you with a set of predefined monitoring policies that allow you to use it immediately. These are described in Predefined policies on page 62. In addition, MAPS provides you with the ability to create customized monitoring rules and policies, allowing you to define specific groups of ports or other elements that will share a common threshold. MAPS lets you define how often to check each switch and fabric measure and specify notification thresholds. This allows you to configure MAPS to provide notifications before problems arise, for example, when network traffic through a port is approaching a bandwidth limit. Whenever fabric measures exceed these thresholds, MAPS can automatically notify you through e-mail messages, SNMP traps, log entries, or combinations. For information on using these features, refer to MAPS groups overview on page 41, MAPS rules overview on page 52, and MAPS policies overview on page 61.

MAPS provides a dashboard that allows you to quickly view what is happening on the switch (for example, the kinds of errors, the error count, and so on). Refer to the MAPS dashboard overview on page 79 for more information.

MAPS license requirements

The Monitoring and Alerting Policy Suite (MAPS) is no longer optional; it replaces Fabric OS Fabric Watch.

In order to provide full functionality, MAPS requires an active and valid Fabric Vision license. Alternatively if you are upgrading and already have a license for Fabric Watch plus a license for Advanced Performance Monitoring, you will automatically get MAPS functionality without having to obtain a separate license. Refer to the *Fabric OS Software Licensing Guide* for more information about licensing and how to obtain the necessary license keys. If you only have one of these licenses, you will need to acquire the other in order to use all the MAPS features.

MAPS activation

The Fabric Vision license must be activated (enabled) in Fabric OS for you to use the full set of MAPS options. In addition, to reuse the Fabric Watch thresholds, you must enable MAPS in Fabric OS 7.3.x before upgrading to Fabric OS 8.0.1.

The following information must be kept in mind when activating MAPS:

- Activating MAPS will activate in all the logical switches.
- On any given chassis there can be multiple logical switches.
- Activating MAPS enables it for all logical switches in the chassis, but each logical switch can have its own MAPS configuration.

MAPS configuration

The MAPS user configuration is persistent across reboot and can be uploaded or downloaded. A configuration upload or download affects only the user-created configuration. You cannot upload or download the default MAPS configuration.

> NOTE
> On a reboot or HA failover, MAPS remains in the same state as it was before the event.

Deleting a user-created MAPS configuration

To remove the entire user-created MAPS configuration, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **mapsconfig --purge**.
   For more information on this command, refer to the *Fabric OS Command Reference*.

MAPS interaction with other Fabric OS features

MAPS interacts in different ways with different Fabric OS features, such as High Availability and Virtual Fabrics.

The following table outlines how MAPS interacts with specific features in Fabric OS.

TABLE 8 Interactions between Fabric OS features and MAPS

Feature	MAPS interaction
Admin Domains	Admin Domains has been deprecated in Fabric OS 8.0.1. NOTE If MAPS is enabled, do not download a configuration that has Admin Domains defined in it, because this might cause unpredictable behavior.
High Availability	MAPS configuration settings are maintained across a High Availability (HA) failover or HA reboot; however, MAPS will restart monitoring after a HA failover or HA reboot and the cached MAPS statistics are not retained.
Virtual Fabrics	When using Virtual Fabrics, different logical switches in a chassis can have different MAPS configurations.

Restrictions on MAPS monitoring

The following restrictions apply globally to MAPS monitoring:

- Small form-factor pluggable (SFP) transceivers on simulated mode (SIM) ports cannot be monitored using MAPS.
- If an event occurs before the dashboard starts monitoring (such as an SCN or an alert), then the event might not be shown in the dashboard.

Refer to Monitoring Flow Vision Flow Monitor data with MAPS on page 107 for additional details about monitoring Flow Vision flows.

Firmware upgrade and downgrade considerations for MAPS

Brocade Monitoring and Alerting Policy Suite (MAPS) has specific considerations when upgrading or downgrading firmware.

Firmware upgrade considerations for MAPS

When upgrading to Fabric OS 8.0.1, there are four main scenarios:

- Upgrading if basic Fabric Watch monitoring is in use without either a Fabric Vision license or Fabric Watch and APM licenses.
- Upgrading if Fabric Watch is installed with a Fabric Watch license.
- Upgrading if basic MAPS monitoring is in use without a Fabric Vision license.
- Upgrading if MAPS is installed with a Fabric Vision license.

Each of these scenarios is discussed in the following sections. In addition, MAPS Fabric Performance Impact monitoring and the legacy bottleneck monitoring feature are mutually exclusive. If the legacy bottleneck monitoring feature was enabled before the upgrade, MAPS will not monitor fabric performance, because by default the Bottleneck Monitor functionality is deprecated in Fabric OS 8.0.1 and MAPS Fabric Performance Impact (FPI) monitoring is enabled.

Upgrading if basic Fabric Watch monitoring is in use without a Fabric Vision license or Fabric Watch and APM licenses.

If basic Fabric Watch monitoring is installed on the switch, but at the time of upgrading you have not activated either the Fabric Vision license or you do not have the Fabric Watch license and the APM license installed, then MAPS is installed and will monitor switches using the policy named “dflt_base_policy” which has the basic monitoring capabilities.

Upgrading if Fabric Watch is installed with a Fabric Watch license

If Fabric Watch is installed on the switch and is enabled, then MAPS is installed with the conservative policy, dft_conservative_policy. If you are upgrading from Fabric OS 7.3.x, you can migrate any existing Fabric Watch thresholds as part of the change to MAPS.

Upgrading if basic MAPS monitoring is in use without a Fabric Vision license

If MAPS gets enabled implicitly (without the Fabric Vision license), then any existing user-defined configurations are not initialized, and MAPS is enabled with the policy “dflt_base_policy”. This policy has rules to monitor only unlicensed features. Any existing user-defined configurations are not initialized. After installing the Fabric Vision license, you can use one of the following setups to enable MAPS with a user-defined configuration:

- Enable MAPS using a different policy in all logical switch contexts.
- Enable MAPS by downloading a user-defined policy containing user-defined configurations.

Upgrading if MAPS is installed with a Fabric Vision license

If MAPS is installed and enabled on the switch, then MAPS will continue to monitor the switch afterwards with no change in operation. MAPS will be enabled with the same active policy that was previously in force on each logical switch and will continue to monitor the fabric based on that policy.

If the Fabric Vision license is installed on the switch but MAPS is not enabled at the time the firmware is upgraded, then after the upgrade, MAPS will be enabled either with the active policy named in the configuration or, if there is no active policy named, with the default conservative policy "dft_conservative_policy". You can choose a different MAPS policy, such as one of the default policies or one of the converted Fabric Watch policies, which are based on the Fabric Watch custom thresholds that were converted to MAPS policies before the upgrade.

If the legacy Latency and Congestion monitoring feature is enabled on a switch using a previous version of Fabric OS, you will receive the following warning when you try to upgrade to Fabric OS 8.0.1:

```
Warning:  Latency and Congestion monitoring features of the bottleneck module are being discontinued in
Fabric OS 8.0.1.  MAPS monitors the same functionality using FPI, and after the upgrade, FPI will start
monitoring the switch.
Do you want to continue [Y/N]?
```

Firmware downgrade considerations for MAPS

The following are the primary downgrade scenarios from Fabric OS 8.0.1 for MAPS.

- When downgrading to any supported version of Fabric OS without an active Fabric Vision license, your switch will continue to use the same MAPS policy.
- When downgrading from Fabric OS 8.0.1 to Fabric OS 7.4.x or 7.3.x with an active Fabric Vision license, Fabric Watch will not be available after the downgrade, but you can still use MAPS at the earlier functionality level. In this case, the Fabric Watch customer thresholds are available, but you would have to do several downgrades to obtain version 7.1.0 of Fabric OS.

When downgrading from Fabric OS 8.0.1 to a previous version, the following MAPS-related behaviors should be expected:

- Downgrading to previous versions of Fabric OS will fail if some features are not supported in the earlier firmware, and their loss could impact MAPS functionality. In this case, MAPS provides instructions on how to disable these features before firmware downgrade. An example of this is if either MAPS actions or rules include Fabric Performance Impact monitoring or port decommissioning.
- Downgrading from Fabric OS to prior versions triggers a warning message if any feature is not supported in the earlier firmware and keeping the feature configuration has no impact. In this case, the downgrade will not be blocked; however, MAPS will display a warning message similar to the following:

  ```
  WARNING: <A>, <B>, <C> feature(s) is/are enabled. These features are not available in FOS <a.b.c>
  release.
  Do you want to continue?
  ```

  Examples of this condition include MAPS having any user-created rules pertaining to monitoring the following: IO_LATENCY_CLEAR, ALL_CIRCUIT_HIGH_QOS, ALL_CIRCUIT_MED_QOS, ALL_CIRCUIT_LOW_QOS, ALL_CIRCUIT_F_QOS, DEV_NPIV_LOGINS, or QT.
- Downgrading to versions of Fabric OS prior to 7.3.0 is not allowed if the MAPS Fabric Performance Impact monitoring feature is enabled. You must disable FPI before starting the firmware downgrade.
- Downgrading is not allowed if automatic VC quarantining has been enabled and there are slow-drain isolations already performed. The isolated slow-drain flows need to be restored to their original virtual circuits before proceeding with the downgrade because the commands to restore them will not be available in the downgraded firmware version.

- Downgrading is not allowed if there is any quarantined port in the logical group ALL_QUARANTINED_PORTS. Before you can downgrade the switch firmware, you must clear the ports from the quarantined state using the **sddquarantine --clear** *slot/port* or **all** command.
- When downgrading from Fabric OS 8.0.1 to a prior version, MAPS issues a warning and asks whether you want to proceed if user-defined rules are present that contain the following monitoring systems:
  - SFPs that were added for Fabric OS 8.0.x
  - Fan air-flow direction
  - Port zoned device ratio
  - IO Insight
  - Extension rule
  - Certificate rules
  - GE_Port rules

  Rules containing these monitoring systems will not be monitored in the versions prior to Fabric OS 8.0.0 and the rules should be deleted. These monitoring systems include the following:

ALL_4_32GSWL_QSFP ALL_32GSWL_SFP ALL_32GLWL_SFP ALL_25Km_16GLWL_SFP ALL_CIRCUIT_IP_HIGH_QOS ALL_CIRCUIT_IP_MED_QOS ALL_CIRCUIT_IP_LOW_QOS ALL_EXT_GE_PORTS ALL_LOCAL_PIDS ALL_TUNNEL_IP_HIGH_QOS ALL_TUNNEL_IP_MED_QOS ALL_TUNNEL_IP_LOW_QOS DAYS_TO_EXPIRE EXPIRED_CERTS FAN_AIRFLOW_MISMATCH GE_CRC GE_INV_LEN GE_LOS_OF_SIG IP_JITTER IP_PKTLOSS IP_RTT IP_UTIL IT_FLOW	RD_1stDATA_TIME_LT_8K RD_1stDATA_TIME_8_64K RD_1stDATA_TIME_64_512K RD_1stDATA_TIME_GE_512K RD_PENDING_IO_LT_8K RD_PENDING_IO_8_64K RD_PENDING_IO_64_512K RD_PENDING_IO_GE_512K RD_STATUS_TIME_LT_8K RD_STATUS_TIME_8_64K RD_STATUS_TIME_64_512K RD_STATUS_TIME_GE_512K WR_1stXFER_RDY_LT_8K WR_1stXFER_RDY_8_64K WR_1stXFER_RDY_64_512K WR_1stXFER_RDY_GE_512K WR_PENDING_IO_LT_8K WR_PENDING_IO_8_64K WR_PENDING_IO_64_512K WR_PENDING_IO_GE_512K WR_STATUS_TIME_LT_8K WR_STATUS_TIME_8_64K WR_STATUS_TIME_64_512K WR_STATUS_TIME_GE_512K

- When downgrading from Fabric OS 8.0.1 to a prior version, the AN modules continue to run in the last known state prior to upgrading to Fabric OS 8.0.1, as long as the switch does *not* have a Fabric Vision license. Otherwise, FPI continues to monitor the switch.

  NOTE
  Other factors might affect downgrading the firmware version; these are only the ones that MAPS affects or is affected by.

Features that do not require a Fabric Vision license

Some features are monitored by MAPS even when the Fabric Vision license is not active.

Monitors that do not require a Fabric Vision license

The following features are monitored by both the unlicensed and licensed versions of MAPS:

- Switch status policies
- Switch resource changes
- FPI monitoring

For more information, refer to MAPS commands that do not require a Fabric Vision license on page 24.

MAPS commands that do not require a Fabric Vision license

The following table lists the MAPS commands that are accessible without an installed Fabric Vision license. This provides similar functionality to that available as an unlicensed Fabric Watch feature in versions of Fabric OS earlier than this version.

TABLE 9 MAPS commands accessible without a Fabric Vision license

Command	Effect	Description
logicalgroup --show	Displays unlicensed feature groups details.	Viewing group information on page 41
mapsconfig --show	Displays the global actions that are currently allowed on a switch.	Enabling or disabling rule actions at a global level on page 53
mapsconfig --action	Specifies which global actions are allowed on a switch.	Enabling or disabling rule actions at a global level on page 53
mapsdb --show	Displays the MAPS dashboard.	MAPS dashboard sections on page 79
mapspolicy --show –summary	Displays a summary of all the policies on the switch.	Viewing policy information on page 64
mapspolicy --show *policy_name*	Displays the rules for the specified MAPS policy on the switch.	Viewing policy information on page 64
mapsrule --show *rule_name*	Displays the details of the specified MAPS rule on the switch.	Viewing MAPS rules on page 68
mapsrule --show -all	Displays all the MAPS rules on the switch.	Viewing MAPS rules on page 68
mapssam --show flash	Displays the flash memory usage as a percentage.	MAPS Service Availability Module on page 136
mapssam --show cpu	Displays the CPU usage as a percentage.	MAPS Service Availability Module on page 136
mapssam --show memory	Displays the general RAM memory usage as a percentage, along with total, used, and free memory values.	MAPS Service Availability Module on page 136
mapshelp	Displays list of MAPS commands.	

NOTE

Even without a Fabric Vision license, help for all the MAPS commands can be displayed.

MAPS Setup and Operation

- Initial MAPS setup 25
- Monitoring across different time windows 28
- Setting the active MAPS policy to a default policy 29
- Pausing MAPS monitoring 30
- Resuming MAPS monitoring 30

Initial MAPS setup

The Brocade Monitoring and Alerting Policy Suite (MAPS) is installed by default, but to enable more than the basic functionality, you must activate the Fabric Vision license.

MAPS is enabled by default starting with FOS 8.0.1. Without a Fabric Vision license the functionality is limited to rules included in the base policy.

Activating MAPS without a Fabric Vision license

If you do not have a Fabric Vision license (or if the license has not been activated), you can still use a limited set of MAPS functions.

MAPS is automatically enabled when you install Fabric OS 8.0.1; however, if you have not installed and activated the Fabric Vision license, you will only be able to use rules that are included in the base policy.

Example of activating MAPS without activating a Fabric Vision license

The following example shows the results when MAPS is automatically enabled without having the Fabric Vision license installed or activated.

```
=============================================================

switch:admin> mapsdb --show

1 Dashboard Information:
========================

DB start time:                  Wed Apr 13 15:29:10 2016
Active policy:                  dflt_base_policy
Configured Notifications:       SW_CRITICAL,SW_MARGINAL
Fenced circuits :               N/A
Quarantined Ports :             None
Top Zoned PIDs <pid(it-flows)>:

2 Switch Health Report:
=======================

Current Switch Policy Status: HEALTHY


3.1 Summary Report:
===================

Category                |Today                    |Last 7 days               |
-------------------------------------------------------------------------------
Port Health             |No Errors                |No Errors                 |
Fru Health              |No Errors                |No Errors                 |
Switch Resource         |No Errors                |No Errors                 |
Traffic Performance     |No Errors                |No Errors                 |
Fabric Performance Impact|No Errors               |No Errors                 |


3.2 Rules Affecting Health:
===========================

Category(Rule Count)|RepeatCount|Rule Name  |Execution Time   |Object   |Triggered Value(Units)|
-------------------------------------------------------------------------------------------------

MAPS is not Licensed. MAPS extended features are available ONLY with License.
switch:admin>
```

Activating MAPS with a Fabric Vision license

When you install and activate a Fabric Vision license you can use all of MAPS functions.

To enable the full functionality of MAPS, you must first install and activate the Fabric Vision license. Perform the following steps to install the license:

1. To make MAPS functionality available for your use, perform the following:
   a) Enter **licenseadd**, followed by the license key provided for your license.
   b) Use **licenseshow** to ensure the license was installed.

   After the license is installed, full MAPS functionality is available. The following example shows these steps:

```
switch:admin> licenseadd H7L73ETXZMFfBQJrKDFNfBWBrABA3N7J7K

switch:admin> licenseshow

H7L73ETXZMFfBQJrKDFNfBWBrABA3N7J7K:
    Fabric Vision license
```

2. To take advantage of MAPS functionality, perform the following:

   a) Enter **mapspolicy --show --summary** to display a list of default policies.

   b) Use **mapspolicy --enable** *default_policy_name* to enable one of the default policies.

      NOTE
      If you have installed a Fabric Vision license, then you should use the conservative, aggressive, or moderate policies. Use the base policy only for basic monitoring, similar to using MAPS without a license. Refer to Features that do not require a Fabric Vision license on page 24 for details.

```
=====================================================================
switch:admin> mapspolicy --show -summary
        Policy Name                    Number of Rules
---------------------------------------------------------------
dflt_aggressive_policy           :               291
dflt_moderate_policy             :               293
dflt_conservative_policy         :               293
dflt_base_policy                 :                46

Active Policy is 'dflt_base_policy'.

switch:admin> mapspolicy --enable dflt_aggressive_p
2016/04/14-15:42:19, [MAPS-1113], 32408, SLOT 1 FID 128, INFO, GEN_6_SWITCH, Policy
dflt_aggressive_policy activated.
```

3. To configure the actions, perform the following:

   a) Use **mapsconfig --actions** *actions* to specify the actions.

      NOTE
      SW_CRITICAL and SW_MARGINAL are enabled by default.

   b) Use **mapsconfig --show** to display and verify the configure notifications and other details.

   The following example shows the use of these commands:

```
switch:admin> mapsconfig --actions raslog,snmp,email,sfp_marginal
2016/04/14-15:44:41, [MAPS-1130], 32411, SLOT 1 FID 128, INFO, GEN_6______ALLIG______GEN_6___,
                     Actions raslog,snmp,email,sfp_marginal configured.

Values for action:              RASLOG, SNMP, EMAIL, FENCE,
----------------------          SDDQ, DECOM, TOGGLE, FMS,
                                SFP_MARGINAL, NONE

switch:admin> mapsconfig --show
Configured Notifications:       RASLOG,SNMP,EMAIL,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL
Mail Recipient:                 admin333@yourcompany.com
Paused members :
===============
PORT :
CIRCUIT :
SFP :
```

   For additional information, refer to the following topics:

   - Refer to Predefined policies on page 62 for details on the default MAPS policies.
   - Refer to Viewing the MAPS dashboard on page 83 for details on the **mapsdb** command output.
   - Refer to MAPS rule actions on page 52 for details on configuring MAPS rule actions.

Quickly monitoring a switch with predefined policies

You can use MAPS to quickly start monitoring your switch with one of the predefined policies delivered with MAPS.

Perform the following steps to quickly monitor a switch with a predefined policy:

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **mapspolicy --enable** followed by the name of the policy you want to enable. You must include an existing policy name in this command. The default policies are:
   - dflt_conservative_policy
   - dflt_aggressive_policy
   - dflt_moderate_policy
   - dflt_base_policy

   NOTE
   If you have installed a Fabric Vision license, then you should use the conservative, aggressive, or moderate policies. Use the base policy only for basic monitoring, similar to using MAPS without a license. Refer to Features that do not require a Fabric Vision license on page 24 for details.
3. Set global actions on the switch to "none" by entering **mapsconfig --actions none**.

   This configuration allows you to test the configured thresholds before enabling their related actions.
4. Monitor the switch by entering **mapsdb --show** or **mapsdb --show all**.
5. Fine-tune the rules used by the policy, as necessary.
6. Set global actions on the switch to the allowed actions by using **mapsconfig --actions** and specifying all of the actions that you want to allow on the switch.

Monitoring across different time windows

You can create rules that monitor across multiple time windows or timebases.

For example, if you want to monitor both for severe conditions and separately for non-critical but persistent conditions, you would construct rules similar to the following.

1. Enter **mapsrule --create** *severe_rule_name* **-monitor** *monitor_name* **-group** *group_name* **-timebase** *time_base* **-op** *operator* **-value** *time* **-action** *action_1, action_2, ...*
2. Enter **mapsrule --create** *persistent_rule_name* **-monitor** *monitor_name* **-group** *group_name* **-timebase** *time_base* **-op** *operator* **-value** *time* **-action** *action_1, action_2, ...*
3. Enter **mapsrule --show** *severe_rule_name* to confirm the rule values.
4. Enter **mapsrule --show** *persistent_rule_name* to confirm the rule values.

Both of the following cases could indicate potential issues in the fabric. Configuring rules to monitor these conditions allows you to correct issues before they become critical.

In the following example, the definition for crc_severe specifies that if the change in the CRC counter in the last minute is greater than 5, it must trigger an e-mail alert and SNMP trap. This rule monitors for the severe condition. It monitors sudden spikes in the CRC error counter over a period of one minute. The definition for crc_persistent specifies that if the change in the CRC counter in the last day is greater than 20, it must trigger a RASLog message and e-mail alert. This rule monitors for slow occurrences of CRC errors that could accumulate to a bigger number over the period of a day.

```
switch1234:admin> mapsrule --create crc_severe -monitor crc -group ALL_PORTS -t min -op g -value 5 -action
email,snmp

switch1234:admin> mapsrule --create crc_persistent -monitor crc -group ALL_PORTS -t day -op g -value 20 -
action raslog,email

switch1234:admin> mapsrule --show crc_severe
Rule Data:
----------
RuleName: crc_severe
Condition: ALL_PORTS(crc/min>5)
Actions: email,snmp
Policies Associated: none

switch1234:admin> mapsrule --show crc_persistent
Rule Data:
----------
RuleName: crc_persistent
Condition: ALL_PORTS(crc/day>20)
Actions: raslog,email
Policies Associated: none
```

Setting the active MAPS policy to a default policy

MAPS allows you to easily set the active MAPS policy to one of the default policies.

To set the active MAPS policy, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **mapspolicy --enable** followed by the name of the policy you want to enable. The default policies are:
   - dflt_conservative_policy
   - dflt_aggressive_policy
   - dflt_moderate_policy
   - dflt_base_policy

     NOTE
     If you have installed a Fabric Vision license, then you should use the conservative, aggressive, or moderate policies. Use the base policy only for basic monitoring, similar to using MAPS without a license. Refer to Features that do not require a Fabric Vision license on page 24 for details.)

   There is no acknowledgment that you have made this change.
3. Enter **mapspolicy --show -summary** to confirm that the policy you specified is active.

The following example sets “dflt_moderate_policy” as the active MAPS policy, and then displays the list of policies and names the active policy.

```
switch:admin> mapspolicy --enable dflt_moderate_policy
switch:admin> mapspolicy --show -summary
        Policy Name                     Number of Rules
-------------------------------------------------------
dflt_aggressive_policy        :                196
dflt_conservative_policy      :                198
dflt_moderate_policy          :                198
dflt_base_policy              :                 20
fw_default_policy             :                109
fw_custom_policy              :                109
fw_active_policy              :                109
Active Policy is 'dflt_moderate_policy'.
```

For more information, refer to Predefined policies on page 62.

Pausing MAPS monitoring

You can stop monitoring ports, FCIP circuits, or SFPs in MAPS. You might do this during maintenance operations, such as device or server upgrades.

To temporarily stop monitoring an element in MAPS, complete the following steps. This suspends MAPS monitoring for the specified element member.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **mapsConfig --config pause** followed by both the element type and the specific members for which you want monitoring paused.

   You must specify both the type and the member information in the command; you specify multiple members by separating them with a comma for individual members, or a hyphen for a range of members.

The following example pauses MAPS monitoring for ports 5 and 7.

```
switch:admin> mapsConfig --config pause -type port -members 5,7
```

Resuming MAPS monitoring

Once you have paused monitoring, you can resume monitoring at any time.

To resume monitoring a paused port or other element in MAPS, complete the following steps. This resumes MAPS monitoring for the specified element member.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **mapsConfig --config continue** followed by both the element type and the specific members for which you want monitoring resumed.

   You must specify both the type and the member information in the command; you specify multiple members by separating them with a comma for individual members, or a hyphen for a range of members.

The following example resumes MAPS monitoring for port 5.

```
switch:admin> mapsConfig --config continue -type port -members 5
```

MAPS Elements and Categories

- MAPS structural elements 31
- MAPS monitoring categories 31

MAPS structural elements

The Monitoring and Alerting Policy Suite (MAPS) has the following structural elements: categories, groups, rules, and policies.

The following table provides a brief description of each structural element in MAPS.

TABLE 10 MAPS structural elements

Element	Description
Action	The activity performed by MAPS if a condition defined in a rule evaluates to true. For more information, refer to Working with MAPS rules and actions on page 68.
Category	A grouping of similar elements that can be monitored (for example, "Security Violations"). For more information, refer to MAPS monitoring categories on page 31.
Condition	An arithmetic or logical expression using either (1) a timebase and a threshold value with a relational operator or (2) a monitor value and a state with a boolean operator. For more information, refer to MAPS conditions on page 49.
Group	A collection of similar objects that you can monitor as a single entity. For example, a collection of ports can be assembled as a group. For more information, refer to MAPS groups overview on page 41.
Monitoring system	A value (measure or statistic) that can be monitored. For more information, refer to MAPS monitoring categories on page 31.
Rule	A direction associating a condition with one or more actions that must occur when the specified condition is evaluated to be true. For more information, refer to MAPS rules overview on page 52.
Policy	A set of rules defining thresholds for triggering actions MAPS is to take when that threshold is triggered. When a policy is enabled, all of the rules in the policy are in effect. For more information, refer to MAPS policies overview on page 61.

MAPS monitoring categories

When you create a rule, you must specify a category to be monitored.

MAPS provides you with the following monitorable categories:

- Port Health on page 32
- Back-end Health on page 33
- FRU Health on page 34
- Security Violations on page 34
- Fabric State Changes on page 34
- Switch Resource on page 35
- Traffic Performance on page 36
- FCIP Health on page 36
- Fabric Performance Impact on page 37
- Switch Status Policy on page 38

In addition to being able to set alerts and other actions based on these categories, the MAPS dashboard displays their status. Refer to MAPS dashboard overview on page 79 for information on using the MAPS dashboard.

Port Health

The Port Health category monitors port statistics and takes action based on the configured thresholds and actions. You can configure thresholds per port type and apply the configuration to all ports of the specified type. Ports whose thresholds can be monitored include physical ports, D_Ports, E_Ports, and F_Ports. The Port Health category also monitors the physical aspects of a small form-factor pluggable (SFP) transceiver, such as voltage, current, receive power (RXP), and transmit power (TXP) in physical ports, D_Ports, E_Ports, and F_Ports.

The following table describes the monitored parameters in this category. In the "Monitored parameter" column, the value in parentheses is the parameter name you specify in **mapsrule -monitor** command.

TABLE 11 Port Health category parameters

Monitored parameter	Description
Cyclic redundancy check (CRC with good EOF (crc g_eof) markers)	The number of times an invalid cyclic redundancy check error occurs on a port or a frame that computes to an invalid CRC. Invalid CRCs can represent noise on the network. Such frames are recoverable by retransmission. Invalid CRCs can indicate a potential hardware problem.
Invalid transmission words (ITW)	The number of times an invalid transmission word error occurs on a port. A word did not transmit successfully, resulting in encoding errors. Invalid word messages usually indicate a hardware problem. **NOTE** The ITW counter includes any physical coding sub-layer (PCS) violations. ITW violations can occur due to an "encoding in" or "encoding out" violation, a PCS violation, or all of these. Encoding violations occur only at slow (8 Gbps or lower) speeds, and PCS violations occur only at high (10 Gbps or higher) speeds.
Sync loss (LOSS_SYNC)	The number of times a synchronization error occurs on the port. Two devices failed to communicate at the same speed. Synchronization errors are always accompanied by a link failure. Loss of synchronization errors frequently occur due to a faulty SFP transceiver or cable. For MAPS 8.0.1, the Loss of sync monitoring happens only for online ports.
Link failure (LF)	The number of times a link failure occurs on offline ports or sends or receives the Not Operational Primitive Sequence (NOPS). Both physical and hardware problems can cause link failures. Link failures also frequently occur due to a loss of synchronization or a loss of signal.
Signal loss (LOSS_SIGNAL)	The number of times that a signal loss occurs in offline ports. Signal loss indicates that no data is moving through the port. A loss of signal usually indicates a hardware problem.
Protocol errors (PE)	The number of times a protocol error occurs on a port. Occasionally, protocol errors occur due to software glitches. Persistent errors generally occur due to hardware problems.
Power-on time (PWR_HRS)	The number of hours that an SFP transceiver has been powered up.
NPIV logins (DEV_NPIV_LOGINS)	The number of NPIV logins to the device. Refer to Monitoring NPIV logins to F_Ports on page 133 for details.
Link reset (LR)	The ports on which the number of link resets exceed the specified threshold value.
Class 3 timeouts (C3TXTO)	The number of Class 3 discard frames because of timeouts.
State changes (STATE_CHG)	The state of the port has changed for one of the following reasons: • The port has gone offline. • The port has come online. • The port is faulty.

TABLE 11 Port Health category parameters (continued)

Monitored parameter	Description
SFP current (CURRENT)	The amperage supplied to the SFP transceiver in milliamps (mA). Current area events indicate hardware failures.
SFP receive power (RXP)	The power of the incoming laser in microwatts (µW). This is used to help determine if the SFP transceiver is in good working condition. If the counter often exceeds the threshold, the SFP transceiver is deteriorating.
SFP transmit power (TXP)	The power of the outgoing laser in microwatts (µW). This is used to help determine if the SFP transceiver is in good working condition. If the counter often exceeds the threshold, the SFP transceiver is deteriorating.
SFP voltage (VOLTAGE)	The voltage supplied to the SFP transceiver in millivolts (mV). If this value exceeds the threshold, the SFP transceiver is deteriorating.
SFP temperature (SFP_TEMP)	The temperature of the SFP transceiver in degrees Celsius. A high temperature indicates that the SFP transceiver may be in danger of damage.

Port health and CRC monitoring

There are two types of CRC errors that can be logged on a switch; taken together they can assist in determining which link introduced the error into the fabric. The two types are plain CRCs, which have bad end-of-frame (EOF) markers and CRCs with good EOF (crc g_eof) markers. When a crc g_eof error is detected on a port, it indicates that the transmitter or path from the sending side may be a possible source. When a complete frame containing a CRC error is first detected, the error is logged, and the good EOF (EOFn) is replaced with a bad EOF marker (EOFni). Because Brocade switches forward all packets to their endpoints, changing the EOF marker allows the packet to continue but not be counted.

For thresholding and fencing purposes, only frames with CRC errors and good end-of-frame markers are counted. This enables you to know exactly how many errors were originated in a specific link.

Back-end Health

The Back-end Health category enables monitor the health of the back-end switch ports for CRC and Link reset error rates, invalid transmission words, BAD_OS, and frame length (either too long or truncated).

The following table lists the monitored parameters in this category.

TABLE 12 Back-end Health category parameters

Monitored parameter	Description
CRC	The number of cyclic redundancy check errors.
Link reset	The number of times a link reset error occurs on a back-end port.
ITW	The number of times an invalid transmission word error occurs on a port. This means that a word did not transmit successfully, resulting in encoding errors. Invalid word messages usually indicate a hardware problem.
BAD_OS	Invalid ordered sets outside the frame.
Frame too long	The frame is longer than expected (greater than 2148 bytes).
Frame truncated	The frame is too short (less than 36 bytes).

For more information on back-end health monitoring, refer to Back-end port monitoring on page 102.

FRU Health

The FRU Health category enables you to define rules for field-replaceable units (FRUs).

The following table lists the monitored parameters in this category. Possible states for all FRU measures are faulty, inserted, on, off, and out.

TABLE 13 FRU Health category parameters

Monitored parameter	Description
Power supplies (PS_STATE)	State of a power supply has changed.
Fans (FAN_STATE)	State of a fan has changed.
Blades (BLADE_STATE)	State of a slot has changed.
SFPs (SFP_STATE)	State of the SFP transceiver has changed.
WWN (WWN_STATE)	State of a WWN card has changed.

Security Violations

The Security Violations category monitors different security violations on the switch and takes action based on the configured thresholds and their actions.

The following table lists the monitored parameters in this category.

TABLE 14 Security Violations category parameters

Monitored parameter	Description
DCC violations (SEC_DCC)	An unauthorized device attempts to log in to a secure fabric.
HTTP violations (SEC_HTTP)	A browser access request reaches a secure switch from an unauthorized IP address.
Illegal command (SEC_CMD)	Commands permitted only to the primary Fibre Channel Switch (FCS) are executed on another switch.
Incompatible security DB (SEC_IDB)	Secure switches with different version stamps have been detected.
Login violations (SEC_LV)	Login violations which occur when a secure fabric detects a login failure.
Invalid Certifications (SEC_CERT)	Certificates are not valid.
SCC violations (SEC_SCC)	SCC violations which occur when an unauthorized switch tries to join a secure fabric. The WWN of the unauthorized switch appears in the ERRLOG.
SLAP failures (SEC_AUTH_FAIL)	SLAP failures which occur when packets try to pass from a non-secure switch to a secure fabric.
Telnet violations (SEC_TELNET)	Telnet violations which occur when a Telnet connection request reaches a secure switch from an unauthorized IP address.
TS out of sync (SEC_TS)	Time Server (TS) violations, which occur when an out-of-synchronization error has been detected.

Fabric State Changes

The Fabric State Changes category contains areas of potential inter-device problems, such as zone changes, fabric segmentation, E_Port down, fabric reconfiguration, domain ID changes, and fabric logins.

The following table lists all the monitored parameters in this category.

TABLE 15 Fabric State Changes category parameters

Monitored parameter	Description
Domain ID changes (DID_CHG)	Monitors forced domain ID changes. These occur when there is a conflict of domain IDs in a single fabric and the principal switch must assign another domain ID to a switch.
Fabric logins (FLOGI)	Activates when ports and devices initialize with the fabric.

TABLE 15 Fabric State Changes category parameters (continued)

Monitored parameter	Description
Fabric reconfigurations (FAB_CFG)	Tracks the number of fabric reconfigurations. These occur when the following events happen: • Two fabrics with the same domain ID are connected • Two fabrics are joined • An E_Port or VE_Port goes offline • A principal link segments from the fabric
E_Port downs (EPORT_DOWN)	Tracks the number of times that an E_Port or VE_Port goes down. E_Ports and VE_Ports go down each time you remove a cable or an SFP transceiver (where there are SFP transceiver failures or transient errors).
Segmentation changes (FAB_SEG)	Tracks the cumulative number of segmentation changes. Segmentation changes occur because of one of the following events occurs: • Zone conflicts • Domain conflicts • Incompatible link parameters During E_Port and VE_Port initialization, ports exchange link parameters, and incompatible parameters (uncommon) result in segmentation. • Segmentation of the principal link between two switches
Zone changes (ZONE_CHG)	Tracks the number of zone changes. Because zoning is a security provision, frequent zone changes may indicate a security breach or weakness. Zone change messages occur whenever there is a change in zone configurations.
Percentage of devices in a Layer 2 fabric (L2_DEVCNT_PER)	Monitors the percentage of imported devices in a Fibre Channel fabric relative to the total number of devices supported in the fabric, whether they are active or not. The switches in a pure Layer 2 fabric do not participate in the metaSAN.
Percentage of devices in a FCR-enabled backbone fabric (LSAN_DEVCNT_PER)	Monitors the percentage of active devices in a Fibre Channel router-enabled backbone fabric relative to the maximum number of devices permitted in the metaSAN. This percentage includes devices imported from any attached edge fabrics.
Used zone configuration size (ZONE_CFGSZ_PER)	Monitors the "used zone configuration" size relative to the maximum zone configuration size on the switch.
Number of FCRs in backbone fabric (BB_FCR_CNT)	Monitors the number of Fibre Channel routers configured in a backbone fabric.

Switch Resource

Switch Resource monitoring enables you to monitor your system's temperature, flash usage, memory usage, and CPU usage.

You can use Switch Resource monitors to perform the following tasks:

- Configure thresholds for MAPS event monitoring and reporting for the environment and resource classes. Environment thresholds enable temperature monitoring, and resource thresholds enable monitoring of flash memory.
- Configure memory or CPU usage parameters on the switch or display memory or CPU usage. Configuration options include setting usage thresholds which, if exceeded, trigger a set of specified MAPS alerts. You can set up the system monitor to poll at certain intervals and specify the number of retries required before MAPS takes action.

The following table lists the monitored parameters in this category.

TABLE 16 Switch Resource category parameters

Monitored parameter	Description
Temperature (TEMP)	The ambient temperature inside the switch in degrees Celsius. Temperature sensors monitor the switch in case the temperature rises to levels at which damage to the switch might occur.
Flash (FLASH_USAGE)	The available compact flash space, calculated by comparing the percentage of flash space consumed with the configured high threshold value.
CPU usage (CPU)	The percentage of CPU available, calculated by comparing the percentage of CPU consumed with the configured threshold value.
Memory (MEMORY_USAGE)	The available memory, calculated by comparing the percentage of memory consumed with the configured threshold value.
Management port (ETH_MGMT_PORT_STATE)	The status of the management port (Bond0).

Traffic Performance

The Traffic Performance category groups areas that monitor the metrics of flows that were created in Flow Vision and for which Flow Monitor has been enabled. Refer to Monitoring of Flow Vision flows with MAPS on page 107 for information about using MAPS to monitor flows. You can use traffic thresholds and alarms to determine traffic load and to reallocate resources appropriately.

The following table lists the monitored parameters in this category. These parameters are the metrics monitored on flows for which the Flow Monitor feature has been enabled. They are only applicable to Flow Monitor flows.

TABLE 17 Traffic Performance category parameters

Monitored parameter	Description
Transmitted frame count (TX_FCNT)	The number of frames transmitted from the flow source.
Received frame count (RX_FCNT)	The number of frames received by the flow destination.
Transmitted throughput (TX_THPUT)	The number of megabytes (MB) transmitted per second by the flow source.
Received throughput (RX_THPUT)	The number of megabytes (MB) received per second by the flow destination.
SCSI frames read (IO_RD)	The number of SCSI I/O read command frames recorded for the flow.
SCSI frames written (IO_WR)	The number of SCSI I/O write command frames recorded for the flow.
SCSI frames read (IO_RD_BYTES)	The number of SCSI I/O bytes read as recorded for the flow.
SCSI frames written (IO_WR_BYTES)	The number of SCSI I/O bytes written as recorded for the flow.

FCIP Health

The FCIP Health category enables you to define rules for FCIP health, including circuit state changes, circuit state utilization, and packet loss.

The following tables list the monitored parameters in this category. The first table lists those FCIP Health parameters monitored on all Brocade platforms.

TABLE 18 FCIP Health category parameters

Monitored parameter	Description
FCIP circuit state changes (CIR_STATE)	The state of the circuit has changed for one of the following reasons: • The circuit has gone offline. • The circuit has come online. • The circuit is faulty.
FCIP circuit utilization (CIR_UTIL)	The percentage of circuit utilization in the configured time period (this can be minute, hour, or day).

TABLE 18 FCIP Health category parameters (continued)

Monitored parameter	Description
FCIP circuit packet loss (CIR_PKTLOSS)	The percentage of the total number of packets that have had to be retransmitted.
FCIP QoS utilization (UTIL)	The percentage of FCIP circuit QoS groups utilization.
FCIP packet loss (PKTLOSS)	The percentage of the total number of packets that have had to be retransmitted in each QoS level. This applies to each FCIP QoS group only.
FCIP circuit round trip time (RTT)	The circuit round-trip latency. This is an absolute value, and only applies to the circuit group.
FCIP circuit jitter (JITTER)	The amount of jitter in a circuit. This is a calculated percentage, and only applies to the circuit group.

The following FCIP parameters are only monitored on the Brocade 7840 extension switch. These parameters are in addition to the ones listed in the previous table.

TABLE 19 Brocade 7840 FCIP Health category parameters

Monitored parameter	Description
FCIP tunnel state change (STATE_CHG)	The count of FCIP tunnel state changes. This applies to the tunnel group only.
FCIP tunnel or tunnel QoS utilization (UTIL)	The percentage of FCIP utilization. This applies to both the tunnel and the tunnel QoS groups.

Refer to the Extension monitoring thresholds on page 149 for the default values.

Fabric Performance Impact

The Fabric Performance Impact category monitors the current condition of the latency seen on E_Ports and F_Ports over different time windows and uses that to determine the performance impact to the fabric and network. MAPS generates alerts when either the congestion levels or port latencies meet or exceed the specified thresholds. To achieve this, MAPS monitors ports for bandwidth utilization and the IO_PERF_IMPACT and IO_FRAME_LOSS bottleneck states. MAPS uses the IO_LATENCY_CLEAR state to show that one of the latency rules was triggered, but the latency has been cleared from the port.

The following table lists the monitored parameters in this category.

TABLE 20 Fabric Performance Impact category parameters

Monitored Parameter	Description
IO_FRAME_LOSS	When a timeout is seen on a port, the bottleneck state of that port is changed to "IO_FRAME_LOSS". This state will also be set if the Inter-Frame-Time (IFT) or Average R_RDY_DELAY is greater than or equal to 80ms.
IO_LATENCY_CLEAR	MAPS supports monitoring for the latency impact CLEAR state. That is, once the device latency is cleared from an F_Port, MAPS clears the latency record and sets the state of the port to "IO_LATENCY_CLEAR". This allows the system to resume monitoring for latency.
IO_PERF_IMPACT	When a port does not quickly clear the frames sent through it, this can cause a backup in the fabric. When MAPS detects that the backpressure from such a condition is significant enough, the bottleneck state of that port is changed to "IO_PERF_IMPACT".
BE_LATENCY_IMPACT	In all fabric edge switches connecting to Brocade Analytics Monitoring Platform with active VTAP flows, MAPS monitors the *CREDZ* back-end port on a mirrored traffic path.
Received bandwidth usage percentage (RX)	The percentage of port bandwidth being used by incoming (RX) traffic. For example, if the port speed is 10 Gbps and the port receives 5 Gb of data in one second, then the percentage of RX utilization is 50 percent (5 Gb*100/(10 Gb*1 second)). For a master trunk port, this indicates the RX percentage for the entire trunk.
Transmitted bandwidth usage percentage (TX)	The percentage of port bandwidth being used by outgoing (TX) traffic. For example, if the port speed is 10 Gbps and the port sends 5 Gb of data in one second, then the percentage of TX utilization is 50 percent (5 Gb*100/(10 Gb*1 second)). For a master trunk port, this indicates the TX percentage for the entire trunk.
Utilization (UTIL)	The percentage of individual port (or trunk) bandwidth being used at the time of the most recent poll.

For more information on Fabric Performance Impact monitoring, refer to Fabric performance impact monitoring using MAPS on page 113.

Switch Status Policy

The Switch Status Policy category lets you monitor the health of the switch by defining the number of types of errors that transition the overall switch state into a state that is not healthy. For example, you can specify a switch status policy so that if a switch has two port failures, it is considered to be in a marginal state, or if it has four failures, it is in a critical (down) state.

The following table lists the monitored parameters in this category and identifies the factors that affect their health.

NOTE
Not all switches support all monitors. Also, MAPS does not monitor CPU or memory usage with switch status policies, and you cannot configure this monitoring.

TABLE 21 Switch Status Policy category parameters

Monitored parameter	Description
Power Supplies (BAD_PWR)	Power supply thresholds detect absent or failed power supplies, and power supplies that are not in the correct slot for redundancy.
Temperatures (BAD_TEMP)	Temperature thresholds, faulty temperature sensors.
Fans (BAD_FAN)	Number of problematic fans, including missing fans and faulty fans.
Flash (FLASH_USAGE)	Flash thresholds.
Marginal Ports (MARG_PORTS)	Thresholds for physical ports, E_Ports, and F_Ports (both optical and copper). Whenever these thresholds are persistently high, the port is marginal.
Faulty Ports (FAULTY_PORTS)	Hardware-related port faults.
Missing SFPs (MISSING_SFP)	Ports that are missing SFP media.
Error Ports (ERR_PORTS)	Ports with errors.
WWN (WWN_DOWN)	Faulty WWN card (applies to modular switches only).
Core Blade (DOWN_CORE)	Faulty core blades (applies to modular switches only).
Faulty blades (FAULTY_BLADE)	Faulty blades (applies to modular switches only).
High Availability (HA_SYNC)	Switch does not have a redundant CP (this applies to modular switches only).

NOTE
Marginal ports, faulty ports, error ports, and missing SFP transceivers are calculated as a percentage of the physical ports (this calculation excludes logical ports, FCoE_Ports, and VE_Ports).

Security certificate monitoring

Fabric OS software provides support for many cryptographic services. The applications use these services for the secure exchange of information. Authentication and encryption of these applications depend on the certificate signature. Therefore, the monitoring of the certificates in MAPS is critical to the reliable functioning of the system.

MAPS supports the following monitoring systems:

- DAYS_TO_EXPIRE—checks the certificate expiry date and notifies the user when the threshold is reached.
- EXPIRED_CERTS—calculates the number of expired certificates, and sets the switch to marginal state if the threshold is reached.

Digital certificates are electronic credentials that are used to ascertain the online identities of individuals, computers, and other entities on a network. In Brocade switches, the certificate is used for authenticating the remote system before communicating and exchanging data

packets with a server. Web servers use the certificate during the encryption of data . The certificate binds the identity of a user, computer, or service to a public key by providing information about the subject of the certificate, the validity of the certificate, and applications and services that can use the certificate.

Version 3 certificates support the following fields that are supported since X.509 version 1:

- Version—Gives the version of the certificate.
- Signature algorithm—An algorithm identifier for the certificate issuer's signature.
- Serial Number—Provides a unique identifier for each certificate that a certificate authority (CA) issues.
- Subject—Provides the name of the computer, user, network device, or service that the CA issues the certificate to.
- Issuer—Provides a distinguished name for the CA that issued the certificate. The issuer name is commonly represented by using an X.500 or LDAP format.
- Valid From—Provides the date and time when the certificate becomes valid.
- Valid To—Provides the date and time when the certificate is no longer considered valid.
- Public Key—Contains the public key of the key pair that is associated with the certificate.

For MAPS to monitor a certificate, the "Valid To" date information is needed from the security library and is the attribute that needs to be tracked. This date is checked every day for every certificate and when any certificate is about to expire, the user is notified to take actions that have been configured in the system. When there are any expired certificates, the switch is in SW_MARGINAL state, and when there are no expired certificates, the switch is reset to HEALTHY state.

The following certificates can be imported to the system:

- HTTPS
- LDAP (TLS client)
- FCAP
- SYSLOG

NOTE
In MAPS alerts, common certs are read as FCAP certificates.

Default policy and rules

The default rules for the new monitoring systems are created in all the default policies. The following rules are added to all the default policies.

```
Rule name                 |Condition                          |Actions           |
--------------------------------------------------------------------------------
Conservative:             |                                   |RASLOG,SNMP,EMAIL,|
defCHASSISCERT_VALIDITY_15|ALL_CERTS(DAYS_TO_EXPIRE/NONE)<15)|SW_MARGINAL       |
defCHASSISCERTS_EXPIRED   |CHASSIS (EXPIRED_CERTS/NONE>0)     |                  |
--------------------------------------------------------------------------------|
Moderate:                 |                                   |RASLOG,SNMP,EMAIL,|
defCHASSISCERT_VALIDITY_20|ALL_CERTS(DAYS_TO_EXPIRE/NONE)<20)|SW_MARGINAL       |
defCHASSISCERTS_EXPIRED   |CHASSIS (EXPIRED_CERTS/NONE>0)     |                  |
--------------------------------------------------------------------------------|
Aggressive:               |                                   |RASLOG,SNMP,EMAIL,|
defCHASSISCERT_VALIDITY_30|ALL_CERTS(DAYS_TO_EXPIRE/NONE)<30)|SW_MARGINAL       |
defCHASSISCERTS_EXPIRED   |CHASSIS (EXPIRED_CERTS/NONE>0)     |                  |
--------------------------------------------------------------------------------
```

Certificate monitor rule creation

The rule creation for certificate monitoring system is similar to the rule creation for any other security monitoring system. You can enable default policies to monitor the certificates or create custom policies rules to monitor the certificates.

The following example defines a rule to send an alert when a certificate is about to expire.

```
 2015/08/03-20:51:01, [MAPS-1004], 1965, SLOT 6 FID 128, INFO, sw0, LDAP Certificate 1,
Condition=ALL_CERTS(DAYS_TO_EXPIRE<20), Current Value:[DAYS_TO_EXPIRE,15 days], RuleName=test_cert_rule_1,
Dashboard Category=Security Violations.
```

The following example defines a rule when one or more certificates are expired.

```
 2015/08/03-20:51:01, [MAPS-1021], 1968, SLOT 6 FID 128, WARNING, sw0, Switch,
Condition=CHASSIS(EXPIRED_CERTS>0), Current Value:[EXPIRED_CERTS,2 days] RuleName=test_cert_rule_2,
Dashboard Category=Security Violations.
```

MAPS Groups, Conditions, Rules, and Policies

- MAPS groups overview....41
- MAPS conditions....49
- MAPS rules overview....52
- MAPS policies overview....61

MAPS groups overview

A MAPS group is a collection of similar objects that you can then monitor using a common threshold.

MAPS provides predefined groups, or you can create a user-defined group and then use that group in rules, thus simplifying rule configuration and management. For example, you can create a group of UNIX ports, and then create specific rules for monitoring this group. As another example, to monitor your network, you can define Flow Vision flows on a switch that have different feature sets, and then import them into MAPS as groups.

Viewing group information

MAPS allows you to view the information for all logical groups collectively, or for a single specific group.

To view a summary of all the logical groups on a switch, enter **logicalgroup --show**. This command returns the group name, and whether the group is predefined. The output presents a table with columns that list characteristics for each group:

- The name of the group
- Whether it is a predefined group
- The type of items in the group (port, SFP, power supply, and so on)
- A list of all of the current members

The following example shows the output of **logicalgroup --show**.

```
switch:admin> logicalgroup --show

--------------------------------------------------------------------------------
Group Name               |Predefined |Type         |Member Count |Members
--------------------------------------------------------------------------------
ALL_100M_16GSWL_QSFP     |Yes        |Sfp          |4            |52-55
ALL_32GSWL_QSFP          |Yes        |Sfp          |0            |
ALL_TARGET_PORTS         |Yes        |Port         |4            |10-11,26-27
ALL_ASICS                |Yes        |Asic         |0            |
ALL_16GLWL_SFP           |Yes        |Sfp          |0            |
ALL_HOST_PORTS           |Yes        |Port         |7            |16-18,21-22,38-39
ALL_SFP                  |Yes        |Sfp          |27           |0,2-3,7,9-12,16-23
NON_E_F_PORTS            |Yes        |Port         |43           |1,4-6,8,13-15,28-31
ALL_10GSWL_SFP           |Yes        |Sfp          |0            |
SWITCH                   |Yes        |             |1            |0
CHASSIS                  |Yes        |             |1            |0
ALL_10GLWL_SFP           |Yes        |Sfp          |0            |
ALL_TS                   |Yes        |Temp. sensor|7            |0-6
ALL_F_PORTS              |Yes        |Port         |12           |10-12,26-27,38-39
ALL_FAN                  |Yes        |Fan          |2            |1-2
ALL_QUARANTINED_PORTS    |Yes        |Port         |0            |
ALL_16GSWL_SFP           |Yes        |Sfp          |21           |0,2,7,9-11,38-39,46
```

```
ALL_QSFP                |Yes        |Sfp          |0            |
ALL_CERTS               |Yes        |Certificate  |0            |
ALL_OTHER_F_PORTS       |Yes        |Port         |1            |12
ALL_E_PORTS             |Yes        |Port         |9            |0,2-3,7,9,32,52,54-55
ALL_WWN                 |Yes        |WWN          |1            |1
ALL_PORTS               |Yes        |Port         |64           |0-63
ALL_D_PORTS             |Yes        |Port         |0            |
ALL_OTHER_SFP           |Yes        |Sfp          |1            |12
ALL_PS                  |Yes        |Power Supply|2            |1-2
ALL_FLASH               |Yes        |Flash        |1            |0
ALL_PIDS                |Yes        |Pid          |12           |All Pids monitored
ALL_25Km_16GLWL_SFP     |Yes        |Sfp          |0            |
ALL_32GSWL_SFP          |Yes        |Sfp          |1            |3
ALL_32GLWL_SFP          |Yes        |Sfp          |0            |
ALL_32GSWL_QSFP         |Yes        |Sfp          |0            |
io_mon_                 |No         |Flow         |1            |Monitored Flow
```

To view details of a specific logical group on a switch, enter **logicalgroup --show** *group_name*. This provides exactly same information as that of **logicalgroup --show** but for the specified group only. The following example shows the output of **logicalgroup --show ALL_TS**.

```
switch:admin> logicalgroup --show ALL_TS
--------------------------------------------------------------------------------
Group Name          |Predefined |Type           |Member Count  |Members
--------------------------------------------------------------------------------
ALL_TS               Yes         Temp Sensor     4              0-3
```

You can also use this command to display the state of flows from a MAPS perspective. The state of a flow is shown in the output in the "Members" column. The following example shows the output of **logicalgroup --show fpm1** for the active Flow Vision flow "fpm1" that has been imported into, and being monitored through, MAPS.

```
switch:admin> logicalgroup -show fpm1
--------------------------------------------------------------------------------
Group Name          |Predefined |Type           |Member Count  |Members
--------------------------------------------------------------------------------
fpm1                 No          Flow            1              Monitored Flow
```

The following example shows the output of **logicalgroup --show fpm2**. In this example, the flow "fpm2" was imported into MAPS, but was subsequently deleted in Flow Vision. MAPS is not monitoring this flow, but it is maintained as a zero member group. If the flow is recreated in Flow Vision and you want to resume monitoring this flow, you must reimport the flow using the **mapsconfig --import** *flow_name* **-force** command. Refer to the *Fabric OS Command Reference* for more information on using the **mapsconfig** or **logicalgroup** commands.

```
switch:admin> logicalgroup --show fpm2
--------------------------------------------------------------------------------
Group Name          |Predefined |Type          |Member Count|Members
--------------------------------------------------------------------------------
fpm2                 No          Flow           0            Not Monitored
                                                              (Stale Flow)
```

Predefined groups

MAPS provides several predefined groups. You cannot delete any of these groups. You can add and remove members from the "PORTS" groups, and you can change the predefined threshold values for any predefined group.

> NOTE
> On switches configured as Access Gateways, F_Ports are categorized and displayed only in the ALL_F_PORT group. They are not categorized in the ALL_HOST_PORTS, ALL_TARGET_PORTS, or ALL_OTHER_F_PORTS groups.

The following table lists these predefined groups organized by object type.

TABLE 22 Predefined MAPS groups

Predefined group name	Object type	Description
ALL_PORTS	FC Port	All ports in the logical switch.
ALL_BE_PORTS	N/A	All back-end ports in the physical switch.
ALL_D_PORTS	FC Port	All D_Ports in the logical switch.
ALL_E_PORTS	FC Port	All E_Ports and EX_Ports in the logical switch. This includes all the ports in E_Port and EX_Port trunks as well. This includes AE ports as well.
ALL_F_PORTS	FC Port	All F_Ports in the logical switch. This includes all the ports in F_Port trunks as well.
ALL_OTHER_F_PORTS	FC Port	All F_Ports in the logical switch which are neither Host nor Target ports.
NON_E_F_PORTS	FC Port	All ports in the logical switch which are neither E_Ports nor F_Ports.
ALL_QUARANTINED_PORTS	FC Port	All ports in the logical switch which have been quarantined for slow-drain performance.
ALL_SFP	SFP	All small form-factor pluggable (SFP) transceivers.
ALL_25Km_16GLWL_SFP	SFP	All 25 kilometer-capable 16-Gbps long wavelength (LWL) small form-factor pluggable (SFP) transceivers in the logical switch.
ALL_32GLWL_SFP	SFP	All 32-Gbps long wavelength (LWL) small form-factor pluggable (SFP) transceivers in the logical switch.
ALL_32GSWL_SFP	SFP	All 32-Gbps short wavelength (SWL) small form-factor pluggable (SFP) transceivers in the logical switch.
ALL_32GSWL_QSFP	QSFP	All 4 X 32-Gbps short wavelength (SWL) quad small form-factor pluggable (QSFP) transceivers in the logical switch.
ALL_10GSWL_SFP	SFP	All 10-Gbps Short Wavelength (SWL) SFP transceivers on FC Ports in the logical switch.
ALL_10GLWL_SFP	SFP	All 10-Gbps Long Wavelength (LWL) SFP transceivers on FC Ports in the logical switch.
ALL_16GSWL_SFP	SFP	All 16-Gbps SWL SFP transceivers in the logical switch.
ALL_16GLWL_SFP	SFP	All 16-Gbps LWL SFP transceivers in the logical switch.
ALL_OTHER_SFP	SFP	All SFP transceivers that do *not* belong to one of the following groups: • ALL_100M_16GSWL_QSFP • ALL_10GSWL_SFP • ALL_10GLWL_SFP • ALL_16GSWL_SFP • ALL_16GLWL_SFP • ALL_QSFP
ALL_QSFP	SFP	All quad small form-factor pluggable (QSFP) transceivers in the logical switch.
ALL_2K_QSFP	SFP	All 2 kilometer-capable 16-Gbps quad small form-factor pluggable (QSFP) transceivers used for the Inter-Chassis Link (ICL) connections in the logical switch.
ALL_SLOTS	Slot	All slots present in the chassis.
ALL_SW_BLADES	Blade	All port and application blades in the chassis.
ALL_CORE_BLADES	Blade	All core blades in the chassis.
ALL_FAN	Fan	All fans in the system.
ALL_FLASH	Flash	The flash memory card in the system.
ALL_PS	Power Supply	All power supplies in the system.
ALL_TS	Temperature Sensor	All temperature sensors in the system.
ALL_WWN	WWN	All WWN cards in the chassis.

TABLE 22 Predefined MAPS groups (continued)

Predefined group name	Object type	Description
SWITCH	Switch	Default group used for defining rules on parameters that are global for the whole switch level, for example, security violations or fabric health.
CHASSIS	Chassis	Default group used for defining rules on parameters that are global for the whole chassis, for example, CPU or flash.
ALL_EXT_GE_PORTS	GE ports	All GE ports in the chassis.
ALL_CERTS	Certificates	All imported certificates.
ALL_CIRCUIT_IP_HIGH_QOS	IP QoS members of a circuit	All IP QoS members of a circuit with high levels of traffic.
ALL_CIRCUIT_IP_MED_QOS	IP QoS members of a circuit	All IP QoS members of a circuit with medium levels of traffic.
ALL_CIRCUIT_IP_LOW_QOS	IP QoS members of a circuit	All IP QoS members of a circuit with low levels of traffic.
ALL_TUNNELS	Tunnels	All tunnels configured on the system.
ALL_TUNNEL_F_QOS	?????	??????
ALL_TUNNEL_IP_HIGH_QOS	IP QoS members of a circuit	All IP QoS members of a tunnel with high levels of traffic.
ALL_TUNNEL_IP_MED_QOS	IP QoS members of a tunnel	All IP QoS members of a tunnel with medium levels of traffic.
ALL_TUNNEL_IP_LOW_QOS	IP QoS members of a tunnel	All IP QoS members of a tunnel with low levels of traffic.

User-defined groups

User-defined groups allow you to specify groups defined by characteristics you select.

In many cases, you may need groups of elements that are more suited for your environment than the predefined groups. For example, small form-factor pluggable (SFP) transceivers from a specific vendor can have different specifications than SFP transceivers from another vendor. When monitoring the transceivers, you may want to create a separate group of SFP transceivers for each vendor. In another scenario, some ports may be more critical than others, and so can be monitored using different thresholds than other ports.

You can define membership in a group either statically or dynamically. For a group using a static definition, the membership is explicit and only changes if you redefine the group. For a group using a dynamic definition, membership is determined by meeting a filter value. When the value is met, the port or device is added to the group, and is included in any monitoring. When the value is not met, the port or device is removed from the group, and is not included in any monitoring data for that group.

The following items should be kept in mind when working with user-defined groups:

- Dynamic groups are only used to group ports.
- The device node WWN information is fetched from the FDMI database, group membership is validated against this database.
- On an Access Gateway device, if you create a group with the feature specified as "device node WWN", then the ports to which the devices are connected will be part of the group.
- On a switch connected to Access Gateway, the ports connected to Access Gateway are not grouped in a user-defined group.
- A port or device can be a member of multiple groups.
- A maximum of 64 user-defined groups and imported flows combined is permitted per logical switch.
- All operations on a dynamic group are similar to those for static groups.

- Group names are not case sensitive; My_Group and my_group are considered to be the same.

Creating a static user-defined group

MAPS allows you to create a monitorable group defined using a static definition, in which the membership is explicit and only changes if you redefine the group.

As an example of a static definition, you could define a group called MY_CRITICAL_PORTS and specify its members as "2/1-10,2/15,3/1-20". In this case, the group has a fixed membership, and to add or remove a member from the group you would have to use the **logicalGroup** command and specify what you want to do (add or remove a member).

To create a static group containing a specific set of ports, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **logicalGroup --create** *group_name* **-type port -members** "*member_list*".

   You can specify either a single port, or specify multiple ports as either individual IDs separated by commas, or a range where the IDs are separated by a hyphen.
3. Optional: Enter **logicalGroup --show** *group_name* **-details** to view the group membership.

The following example creates a group named MY_CRITICAL_PORTS whose membership is defined as the ports 2/1-10,2/15,3/1-20.

```
switch:admin> logicalgroup --create MY_CRITICAL_PORTS -type port -members "2/1-10,2/15,3/1-20"
```

For more information on the **logicalGroup** command, refer to the *Fabric OS Command Reference*.

Modifying a static user-defined group

MAPS allows you to modify the membership of a static user-defined group (that is, one with a fixed membership).

To change which ports are in a static user-defined group, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Use the following commands to add or delete specific ports from the group.
   - To explicitly add ports to the group, enter **logicalGroup --addmember** *group_name* **-members** *member_list*.
   - To explicitly remove ports from the group, enter **logicalGroup --delmember** *group_name* **-members** *member_list*.

   You need to specify every element in the command. When adding or removing ports, you can specify either a single port (2/15), or specify multiple ports as either individual IDs separated by commas (2/15, 5/16), or a range of ports with the IDs separated by hyphens (2/15-16/15).
3. Optional: Enter **logicalGroup --show** *group_name* **-details** to view the group membership.

The following example removes the port 2/15 from the MY_CRITICAL_PORTS group:

```
switch:admin> logicalgroup --delmember MY_CRITICAL_PORTS -members 2/15
```

For more information on the **logicalGroup** command, refer to the *Fabric OS Command Reference*.

Creating a dynamic user-defined group

By using a dynamic definition, you can create a group that can be monitored, with membership determined by meeting a filter value. When the value is matched, the port or device is added to the group, and it is included in any monitoring. When the value is not matched, the port or device is removed from the group, and it is not included in any monitoring data for that group.

As an example of a dynamic definition, you could specify a port name or an attached device node WWN, and all ports which match the port name or device node WWN will be automatically included in this group. As soon as a port meets the criteria, it is automatically added to the group. As soon as it ceases to meet the criteria, it is removed from the group. The characters in the following table are used to identify the feature characteristics (port name or device node WWN) that you want to specify to identify the group.

TABLE 23 Group-definition operators

Character	Meaning	Explanation
*	Match any set of characters in the position indicated by the asterisk.	Defining the port name as **brcdhost*** will include any port name starting with brcdhost, such as brcdhost1, brcdhostnew, and so on.
?	Match any single character in the position indicated by the question mark.	Defining the port name as **brcdhost?** will include any port name that has exactly one character following "brcdhost," such as brcdhost1, brcdhostn, and so on. However, brcdhostnew will not match this criterion.
[*expression*]	Match any character defined by the expression inside the square brackets; that is, one character from the set specified in the expression. For example, [1-4] will match for values of 1, 2, 3, or 4.	Defining the port name as **brcdhost[1-3]** will include only the port names brcdhost1, brcdhost2, and brcdhost3.
!	Match the string following, and exclude any ports that match. You must include the entire term in single quotation marks (').	Defining the port name as **'!brcdhost'** will include all the port names except for those that begin with brcdhost.

To create a dynamic group of all the ports that are connected to devices that have a node WWN starting with 30:08:00:05, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **logicalgroup --create** *group_name* **-type** *type* **-feature** *feature_type* **-pattern** *pattern*.

   For *feature_type*, either port names or WWNs can be used, not both. Quotation marks around the *pattern* value are required. If ! is specified in the pattern it must be within single quotation marks ('!'). You can only specify one feature as part of a group definition.
3. Optional: Enter **logicalgroup --show** *group_name* **-details** to view the group membership.

The following example creates a group named "GroupWithWwn_30:08:00:05" that has a membership defined as ports connected to a device with a node WWN that starts with 30:08:00:05.

```
switch:admin> logicalgroup --create GroupWithWwn_30:08:00:05 -type port -feature nodewwn -pattern
"30:08:00:05*"
```

Alternatively, the following example creates a group that has a membership defined as ports with a port name that begins with "brcdhost." The only difference from the previous example is that the feature is defined as "portname" rather than "nodewwn".

```
switch:admin> logicalgroup --create GroupWithNode_brcdhost -type port -feature portname -pattern "brcdhost*"
```

For more information on the **logicalgroup** command, refer to the *Fabric OS Command Reference*.

Modifying a dynamic user-defined group

MAPS allows you to change the definition pattern used to specify a dynamic user-defined group after you have created it.

To modify a dynamic user-defined group after you have created it, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **logicalGroup --update** *group_name* **-feature** *feature_name* **-pattern** *pattern*.

NOTE
The values for group_name and feature_name must match existing group and feature names. You can only specify one feature as part of a group definition.

3. Use the following commands to add or delete specific ports from the group. (You can also use this command to modify the group membership of pre-defined groups.)
   - To explicitly add ports to the group, enter **logicalGroup --addmember** *group_name* **-members** *member_list*.
   - To explicitly remove ports from the group, enter **logicalGroup --delmember** *group_name* **-members** *member_list*.

   You need to specify every element in the command. When adding or removing ports, you can specify either a single port (2/15), or specify multiple ports as either individual IDs separated by commas (2/15, 2/16, 3/15), or a range of ports with the IDs separated by hyphens (2/15-16, 3/15).
4. Optional: Enter **logicalGroup --show** *group_name* **-details** to view the group membership.

The following example changes the node WWN of the attached devices in Group_001 to start with 30:08:01.

```
switch:admin> logicalgroup --update Group_001 -feature nodewwn -pattern "30:08:01*"
```

For more information on the **logicalGroup** command, refer to the *Fabric OS Command Reference*.

Restoring a group to its default membership

MAPS allows you to restore the membership of any modified MAPS group back to its default. This can be done to predefined groups and dynamic user-defined groups. This command does not work on groups with a static definition.

To restore the membership of a modified MAPS group, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **logicalgroup --restore** *group_name*. This restores the group membership to its default.
3. Optional: Enter **logicalgroup --show** *group_name* **-details** to view the group membership.

The following example restores all the deleted members and removes the added members of the GOBLIN_PORTS group. First it shows the detailed view of the modified GOBLIN_PORTS group, then restores the membership of the group and then it shows the post-restore group details. Notice the changes in the MemberCount, Members, Added Members, and Deleted Members fields between the two listings.

```
switch:admin> logicalgroup --show GOBLIN_PORTS -detail

GroupName          : GOBLIN_PORTS
Predefined         : No
Type               : Port
MemberCount        : 11
Members            : 1-2,12-20
Added Members      : 2,20
Deleted Members    : 10-11
Feature            : PORTNAME
Pattern            : port1*

switch:admin> logicalgroup --restore GOBLIN_PORTS

switch:admin> logicalgroup --show GOBLIN_PORTS -detail

GroupName          : GOBLIN_PORTS
Predefined         : No
Type               : Port
MemberCount        : 11
Members            : 1,10-19
Added Members      :
Deleted Members    :
Feature            : PORTNAME
Pattern            : port1*
```

Cloning a group

MAPS allows you to clone any predefined, static, or dynamic user-defined group.

To clone a group, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **logicalgroup --clone** *existing_group_name* **-name** *new_group_name*.
   You can now modify the new group.

The following example clones the predefined group "ALL_TARGET_PORTS" as "ALL_TARGET_PORTS-LR_5".

```
switch:admin> logicalgroup --clone ALL_TARGET_PORTS -name ALL_TARGET_PORTS-LR_5
```

Deleting groups

The **logicalgroup --delete** *group_name* command allows you to remove any logical group other than the predefined groups.

You cannot delete a group that is used by any rules. Adding the **-force** option to this command overrides the default behavior and forces the deletion of all the rules that are configured with the given group and then deletes the group. If a logical group is present in user-defined rules, the **-force** option deletes all the rules that are configured with the given group and then deletes the group.

The following example shows that the user-defined group GOBLIN_PORTS exists, deletes the group, and then shows that the group has been deleted.

```
switch:admin> logicalgroup --show
--------------------------------------------------------------------------------
Group Name             |Predefined |Type          |Member Count |Members
--------------------------------------------------------------------------------
ALL_PORTS              |Yes        |Port          |48           |6/0-15,7/0-31
ALL_SFP                |Yes        |SFP           |11           |7/8-14,7/24-27
ALL_PS                 |Yes        |PowerSupply   |2            |0-1
  :                      :           :              :             :
  :                      :           :              :             :
GOBLIN_PORTS           |No         |Port          |10           |1/1-5,3/7-9,3/12
SFPGroup               |No         |SFP           |10           |1/1-5,3/7-9,3/12
ALL_QUARANTINED_PORTS  |Yes        |Port                |2            |8/0,8/4

switch:admin> logicalgroup --delete GOBLIN_PORTS
switch:admin> logicalgroup --show
--------------------------------------------------------------------------------
Group Name             |Predefined |Type          |Member Count |Members
--------------------------------------------------------------------------------
ALL_PORTS              |Yes        |Port          |48           |6/0-15,7/0-31
ALL_SFP                |Yes        |SFP           |11           |7/8-14,7/24-27
ALL_PS                 |Yes        |PowerSupply   |2            |0-1
  :                      :           :              :             :
  :                      :           :              :             :
SFPGroup               |No         |SFP           |10           |1/1-5,3/7-9,3/12
ALL_QUARANTINED_PORTS  |Yes        |Port                |2            |8/0,8/4
```

MAPS conditions

A MAPS condition includes a monitoring system, a timebase, and a threshold. If the condition is evaluated as true, the actions specified in the rule are triggered. The condition depends on the element that is to be monitored.

For example, if you specified that a rule should be triggered if the CRC counter increment in the last minute is greater than 10, then the threshold value is 10 and the timebase is the preceding minute. In this rule, the condition is the combination of the two; that is, the CRC value must be greater than the threshold value of 10 AND this threshold must be exceeded during the minute timebase. If the counter reaches 11 within that minute, the rule would trigger.

> NOTE
> MAPS conditions are applied on a per-port basis, not switch- or fabric-wide. For example, 20 ports that each get 1 CRC counter would not trigger a “greater than 10” rule.

Threshold values

Thresholds are the values at which potential problems might occur. When configuring a rule, you can specify a threshold value that, when exceeded, triggers an action. For example, if you had specified a rule to log a RASLog entry if the CRC counter is greater than 10, then when the counter reaches 11, the rule is triggered and a RASLog entry is logged.

Timebase

The **–timebase** value specifies the time interval between samples, and affects the comparison of sensor-based data with user-defined threshold values.

You can set the timebase to the following durations:

Duration	Description

day	Samples used for comparison are one day apart.
hour	Samples used for comparison are one hour apart.
minute	Samples used for comparison are one minute apart.
none	A comparison is made between the real-time value and the configured threshold value.

- **day** specifies that the samples will be compared once a day.
- **hour** specifies that the samples are compared every hour.
- **minute** specifies that the samples are compared every minute.
- **none** is used for comparisons where the timebase is not applicable.

The following example sets the timebase for the rule "BiWeekly" to be every two weeks (14 days).

```
switch:admin> mapsRule --config BiWeekly -group ALL_F_PORTS -monitor LOSS_SIGNAL -op eq -timebase day -
value 14 -action raslog,snmp
```

Supported timebases

The following table identifies which monitors support which timebases.

TABLE 25 Monitors and supported timebases

Monitor Name	Day	Hour	Minute	None
Domain ID change	Yes	Yes	Yes	No
Fabric logins	Yes	Yes	Yes	No
Fabric reconfigurations	Yes	Yes	Yes	No
E_Ports down	Yes	Yes	Yes	No
Segmentation changes	Yes	Yes	Yes	No
Zone changes	Yes	Yes	Yes	No
L2 Device Count	Yes	Yes	Yes	No
LSAN Device Count	Yes	Yes	Yes	No
Zone Configuration size	No	No	No	Yes
FCR Count	Yes	Yes	Yes	No
Tunnel (STATE_CHG)	Yes	Yes	Yes	No
Tunnel QoS (UTIL)	Yes	Yes	Yes	No
QoS packet loss percentage (PKTLOSS)	Yes	Yes	Yes	No
Circuit state (CIR_STATE)	Yes	Yes	Yes	No
Circuit utilization percentage (CIR_UTIL)	Yes	Yes	Yes	No
Circuit packet loss percentage (CIR_PKTLOSS)	Yes	Yes	Yes	No
Circuit round-trip times (RTT)	Yes	Yes	Yes	No
Circuit jitter (JITTER)	Yes	Yes	Yes	No
Power Supply (PS_STATE)	No	No	No	No
Fan (FAN_STATE)	No	No	No	No
Blade (BLADE_STATE)	No	No	No	No
SFP (SFP_STATE)	No	No	No	No
WWN (WWN)	No	No	No	No
CRC Errors	Yes	Yes	Yes	No

TABLE 25 Monitors and supported timebases (continued)

Monitor Name	Day	Hour	Minute	None
Invalid Transmit Words	Yes	Yes	Yes	No
Sync Loss	Yes	Yes	Yes	No
Link Failure	Yes	Yes	Yes	No
Loss of Signal	Yes	Yes	Yes	No
Protocol Errors	Yes	Yes	Yes	No
Link Reset	Yes	Yes	Yes	No
C3 Time outs	Yes	Yes	Yes	No
State change	Yes	Yes	Yes	No
SFP Current	No	No	No	Yes
SFP Receive Power	No	No	No	Yes
SFP Transmit Power	No	No	No	Yes
SFP Voltage	No	No	No	Yes
SFP Temperature	No	No	No	Yes
SFP Power On Hours	No	No	No	Yes
Flash memory percentage used (FLASH_USAGE)	No	No	No	Yes
CPU percentage used (CPU)	No	No	No	Yes
Memory percentage used (MEMORY_USAGE)	No	No	No	Yes
Ethernet management port state (ETH_MGMT_PORT_STATE)	No	No	No	Yes
Temperature Sensor (TEMP)	No	No	No	Yes
DCC violations	Yes	Yes	Yes	No
HTTP violation	Yes	Yes	Yes	No
Illegal command	Yes	Yes	Yes	No
Incompatible security DB	Yes	Yes	Yes	No
Login violations	Yes	Yes	Yes	No
Invalid certifications	Yes	Yes	Yes	No
No-FCS	Yes	Yes	Yes	No
SCC violations	Yes	Yes	Yes	No
SLAP failures	Yes	Yes	Yes	No
Telnet violations	Yes	Yes	Yes	No
TS out of sync	Yes	Yes	Yes	No
Current (CURRENT)	No	No	No	No
Receive Power (RXP)	No	No	No	No
Transmit Power (TXP)	No	No	No	No
Voltage (VOLTAGE)	No	No	No	No
Temperature (TEMP)	No	No	No	No
Fabric Performance Impact (DEV_LATENCY_IMPACT)	No	No	No	Yes
Receive Bandwidth usage percentage (RX)	Yes	Yes	Yes	No
Transmit Bandwidth usage percentage (TX)	Yes	Yes	Yes	No
Trunk Utilization percentage (UTIL)	Yes	Yes	Yes	No
Absent or faulty power supply (BAD_PWR)	No	No	No	Yes
Temperature sensors outside range (BAD_TEMP)	No	No	No	Yes

TABLE 25 Monitors and supported timebases (continued)

Monitor Name	Day	Hour	Minute	None
Absent or faulty fans (BAD_FAN)	No	No	No	Yes
Flash usage (FLASH_USAGE)	No	No	No	Yes
Percentage of marginal ports (MARG_PORTS)	No	No	No	Yes
Percentage of error ports (ERR_PORTS)	No	No	No	Yes
Percentage of faulty ports (FAULTY_PORTS)	No	No	No	Yes
Faulty blades (FAULTY_BLADE)	No	No	No	Yes
Faulty WWN (WWN_DOWN)	No	No	No	Yes
Core blade monitoring (DOWN_CORE)	No	No	No	Yes
HA Sync (HA_SYNC)	No	No	No	Yes
Receive throughput	Yes	Yes	Yes	No
Transmit Frame Count	Yes	Yes	Yes	No
Receive Frame Count	Yes	Yes	Yes	No
Transmit throughput	Yes	Yes	Yes	No
IO Read Command Count	Yes	Yes	Yes	No
IO Write Command Count	Yes	Yes	Yes	No
IO Read Data	Yes	Yes	Yes	No
IO Write Data	Yes	Yes	Yes	No
Throughput Degradation	Yes	Yes	Yes	Yes

MAPS rules overview

A MAPS rule associates a condition with actions that need to be taken when the condition is evaluated to be true and the specified rule is triggered. A MAPS rule can exist outside of a MAPS policy, but are only evaluated by MAPS when the rule is part of the active policy.

Each rule specifies the following items:

- A group of objects to be evaluated. Refer to MAPS groups overview on page 41 for additional information.
- The condition being monitored. Each rule specifies a single condition. A condition includes a timebase and a threshold. Refer to MAPS conditions on page 49 for additional information.
- The actions to take if the condition is evaluated to be true. Refer to MAPS rule actions on page 52 for additional information.

The combination of actions, conditions, and groups allows you to create a rule for almost any scenario required for your environment.

MAPS rule actions

When you create, modify, or clone a rule using the **mapsrule --create**, **--config**, or **--clone** commands, you associate an action for MAPS to take if the condition defined in the rule evaluates to "true." Each rule can have one or more actions associated with it. For example, you can configure a rule to log a RASLog message and fence the port if the number of CRC errors on any E_Port is greater than 20 per minute.

MAPS provides the following actions for rules:

- E-mail alerts on page 54
- FICON alerts on page 55

- Port fencing and port decommissioning on page 56
- RASLog messages on page 59
- SFP marginal on page 60
- Slow Drain Device Quarantine on page 61
- MAPS SNMP traps on page 56
- Switch critical on page 61
- Switch marginal on page 61
- Port toggling on page 59

For each action, you can define a "quiet time" for most rules in order to reduce the number of alert messages generated. Refer to Quieting a rule on page 72 for details.

For rules that are state-bound (such as those for temperature sensing monitoring such as IN_RANGE or OUT_OF_RANGE), the rule is only triggered when the condition transitions from one state to another. These types of rules will not be triggered if the condition remains in a single state, even if that state is not a normal operating state, such as PS_Faulty. .

The global action settings on a switch take precedence over the actions defined in the rules. For example, if the global action settings allow RASLog alerts, but do not allow port fencing, then if the CRC threshold is reached, a RASLog message would be issued but the port would not be fenced. To enable global actions, use the **mapsconfig --actions** command. For more details, refer to Enabling or disabling rule actions at a global level on page 53. Refer to the *Fabric OS Command Reference* for further details on using the **mapsconfig** and **mapsrule** commands.

Enabling or disabling rule actions at a global level

Allowable actions on a switch can be specified globally, and supersede any actions specified in individual rules. Enabling and disabling actions at a global level allows you to configure rules with stricter actions, such as port fencing, but disable the action globally until you can test the configured thresholds. After validating the thresholds, you can enable an action (such as port decommissioning) globally without having to change all of the rules.

In Fabric OS 8.0.0 and later, you can configure all actions, even if the switch does not have a license. However, without the license, MAPS only performs email, RASLog, and SNMP actions.

> ATTENTION
> For MAPS to trigger an action, the action must be explicitly enabled using the **mapsconfig --actions** command.

To enable or disable actions at a global level, complete the following steps.

1. Enter **mapsconfig --show** to display the actions that are currently allowed on the switch.
2. Enter **mapsconfig --actions** and specify all of the actions that you want to allow on the switch, for example, **mapsconfig --actions** *action1*, *action2*, *action3* ... (up to the complete set of actions).

> NOTE
> If you are changing the list of active actions, you need to specify all the actions to be active. For example, if you are adding RASLog notifications to a switch that already has e-mail notifications enabled, you must specify both "email" and "RASLog" as actions in the **mapsconfig** command.

To disable all actions, enter **mapsconfig --actions none**. The keyword **none** cannot be combined with any other action.

The following example shows that RASLog, e-mail, and fence notifications are not currently active actions on the switch, and then shows them being added to the list of allowed actions.

> NOTE
> Starting with 8.0.1, SW_CRITICAL and SW_MARGINAL notifications are enabled by default in a switch and they cannot be disabled by the user.

```
Updated for 8.0.1 (03-30-16)to show that SW_CRITICAL and SW_MARGINAL are
already configured. Examples provided by Steve McKie.
switch:admin> mapsconfig --show
Configured Notifications:       SW_CRITICAL,SW_MARGINAL
Mail Recipient:                 Not Configured
Paused members :
===============
PORT :
CIRCUIT :
SFP :


switch:admin> mapsconfig --actions raslog,email,fence
2016/03/28-19:05:29, [MAPS-1130], 406, SLOT 2 FID 2, INFO, switch_2,
Actions raslog,email,fence configured.


switch:admin> mapsconfig --show
Configured Notifications:       RASLOG,EMAIL,FENCE,SW_CRITICAL,SW_MARGINAL
Mail Recipient:                 Not Configured
Paused members :
===============
PORT :
CIRCUIT :
SFP :
```

E-mail alerts

An e-mail alert action sends information about the event to one or more specified e-mail addresses. The e-mail alert specifies the threshold and describes the event, much like an error message.

To configure the e-mail recipients, use the **mapsConfig --emailcfg** command. Multiple destination e-mail addresses are possible; they must be separated using a comma between each individual address and each address must be a complete e-mail address. For example, abc@brocade.com is a valid e-mail address; abc@brocade is not. Refer to Sending alerts using e-mail on page 75 for more information.

To clear all configured e-mail addresses, enter **mapsconfig --emailcfg -address none**. All configured e-mail addresses will be erased.

Enhancements to e-mail alert content

In Fabric OS 8.0.1, e-mail alerts have been enhanced to include information that is more meaningful, including the monitor port name and rephrasing of other content so it helps you understand the error condition or violation in the switch and take action accordingly.

Examples of e-mail alert enhancements

The follow example shows the enhancements for e-mail alerts from threshold-based rules. The enhanced information is labeled "Subject," "Group," and "Current Value." :

TABLE 26 E-mail alerts from threshold-based rules

Previous e-mail alert	Enhanced e-mail alert
U-Port 38 triggered LOSS_SIGNAL rule - defALL_HOST_PORTSLOSS_SIGNAL_0Switch Time: Feb 11 19:43:30 Affected Entity: slot1 port28, F-Port 1/28 Rule Name: defNON_E_F_PORTSLOSS_SIGNAL_0 Condition: ALL_OTHER_F_PORTS(TX/hour>90.00) Dashboard Category: Port Health Switch Name: dcx_178 Switch WWN: 10:00:00:05:1e:47:64:00 Switch IP: 10.38.18.178 Fabric Name: xyz VFID: 128	Subject: 2015/02/11-19:43:30: Rule for monitor “loss of signal” has been triggered for port member 1/28. Affected Entity: slot1 port28, F-Port 1/28 Rule Name: defNON_E_F_PORTSLOSS_SIGNAL_0 Condition: ALL_OTHER_F_PORTS(TX/hour>90.00) Group: ALL_OTHER_F_PORTS Current Value: 1 LOS Dashboard Category: Port Health Switch Name: dcx_178 Switch WWN: 10:00:00:05:1e:47:64:00 Switch IP: 10.38.18.178 Fabric Name: xyz VFID: 128

TABLE 27 E-mail alerts from state-change rules

Previous e-mail alert	Enhanced e-mail alert
F-Port 9/47 triggered DEV_LATENCY_IMPACT rule - defALL_PORTS_IO_PERF_IMPACTSwitch Time: Oct 29: 07:02:05 Affected Entity: slot9 port47, F-Port 9/47 Rule Name: defALL_PORTS_IO_PERF_IMPACT Condition: ALL_PORTS(DEV_LATENCY_IMPACT==IO_PERF_IMPACT) Dashboard Category: Fabric Performance Impact Switch Name: ALLEGIANCE_SWAT Switch WWN: 10:00:00:27:f8:f0:32:70 Switch IP: 10.17.59.30 Fabric Name: FV_Krishna VFID: 128	Subject: 2015/10/29-07:02:05:422719: Rule for monitor “Device Latency Impact” has been triggered for port member 9/47. Affected Entity: slot9 port47, F-Port 9/47 Rule Name: defALL_PORTS_IO_PERF_IMPACT Condition: ALL_PORTS(DEV_LATENCY_IMPACT==IO_PERF_IMPACT) Group: ALL_PORTS Current Value: IO_PERF_IMPACT, (97.5% of 1 secs) Dashboard Category: Fabric Performance Impact Switch Name: ALLEGIANCE_SWAT Switch WWN: 10:00:00:27:f8:f0:32:70 Switch IP: 10.17.59.30 Fabric Name: FV_Krishna VFID: 128

NOTE

Member information is only added to the message subject for the rule violations in which the group name is neither a switch group nor a chassis group.

FICON alerts

FICON notification support is available as an action from MAPS.

The FICON management service uses the MAPS events to create a health summary report. Rules with a FICON notification action are part of all four default policies. In the active policy, if FICON notification is specified for any triggered events, MAPS sends a notification with the following event information to the FICON management service:

- Triggered event rule name
- Object and its type on which the event was triggered
- Severity of the event

- Condition on which the event was triggered
- Monitoring service details and the measured value

MAPS SNMP traps

When specific events occur on a switch, SNMP generates a message (called a "trap") that notifies a management station.

A MAPS SNMP trap forwards the following information to an SNMP management station:

- Name of the element whose counter registered an event
- Area and index number of the threshold that the counter crossed
- Event type
- Value of the counter that exceeded the threshold
- State of the element that triggered the alarm
- Source of the trap

In environments where you have a high number of messages coming from a variety of switches, you might want to receive them in a single location and view them using a graphical user interface (GUI). In this type of scenario, Simple Network Management Protocol (SNMP) notifications could be the most efficient notification method, because you can avoid having to log in to each switch individually as you would have to do for error log notifications.

In order to get the event notifications, you must configure the SNMP software to receive the trap information from the network device, and you must configure the SNMP agent's IP address on the switch to send the trap to the management station. You can configure SNMP traps receiver using the **snmpconfig** command. For additional information on configuring the SNMP agent using **snmpconfig**, refer to the *Fabric OS Command Reference*.

SNMP MIB support

MAPS requires SNMP management information base (MIB) support on the device for management information collection. For additional information on SNMP MIB support, refer to the *MIB Reference Manual*.

SNMP quiet time support

MAPS supports quiet time only for RASLog and EMAIL actions. The following example shows quiet time support in a RASLog.

- Quiet time time-out: one minute.
- Number of times the rule triggered: 2
- Last rule execution time: Tue Jan 11 12:56:26 2016

```
Switch, Condition=SWITCH(SEC_LV/min>=0), Current Value:[SEC_LV,2 Violations], Rule
SWITC_LVGE0_M_RxxTxxxxxxx triggered 2 times in 1 min and last trigger time Tue Jan 11 12:56:26 2016,
Dashboard Category=Security Violations
```

Port fencing and port decommissioning

MAPS supports port fencing for both E_Ports and F_Ports. MAPS also supports port decommissioning for E_Ports. However, decommissioning for F_Ports can only be done with MAPS in conjunction with Brocade Network Advisor. These actions automatically take ports offline when configured thresholds in a given rule are exceeded. Port fencing immediately takes ports offline, which might cause loss of traffic. Port decommissioning takes a port offline without loss of traffic. Both are disabled by default. Port

decommissioning and port fencing can only be configured for the port health monitoring systems for which decommissioning is supported.

Port decommissioning cannot be configured by itself in a MAPS rule or action. It requires port fencing to be enabled in the same rule. If you attempt to create a MAPS rule or action that has port decommissioning without port fencing, the rule or action will be rejected. MAPS can be configured to have only port fencing enabled in a rule; if this is the case, the port will be taken offline immediately.

MAPS supports port fencing and port decommissioning actions for rules that monitor CRC, ITW, PE, LR, STATE_CHG, or C3TXTO errors from physical ports, such as E_Ports, F_Ports, or U_Ports. Otherwise, for circuits, MAPS supports only port fencing actions for rules that monitor changes of state (STATE_CHG). Refer to the Port Health monitoring thresholds on page 150 tables for these rules.

Be aware that if multiple rules for the same counter are configured with different thresholds, then both port fencing and port decommissioning should be configured for the rule with the highest threshold monitored. For example, if you configure one rule with a CRC threshold value "greater than 10 per minute" and you configure a second rule with a CRC threshold value "greater than 20 per minute", you should configure port fencing and port decommissioning as the action for the rule with the 20 per minute threshold, as configuring it for the 10 per minute rule will block the other rule from being triggered.

Port decommissioning for E_Ports and F_Ports

For E_Ports, if port decommissioning fails, MAPS will fence the port. Switches themselves can decommission E_Ports through MAPS. In this case, when port decommissioning is triggered on an E_Port, the neighboring switches will perform a handshake so that traffic is re-routed before the port is disabled. Be aware that there are multiple reasons that the port-decommissioning operation between two E_Ports could fail; for example, if the link that fails is the last link between the two switches. To see which parameters can trigger port fencing and port decommissioning, refer to Port Health monitoring thresholds on page 150.

For F_Ports, port decommissioning will only work if BNA is actively monitoring the switch. BNA can decommission F_Ports based on specified criteria (refer to Port Health monitoring thresholds on page 150.) MAPS notifications are integrated with BNA, which in turn must coordinate with the switch and the end device to orchestrate the port decommissioning. If BNA is not configured on a switch, MAPS will fence the F_Port.

For more information on port fencing, port decommissioning, and related failure codes, refer to the *Fabric OS Administrator's Guide*.

Configuring port decommissioning

Port decommissioning is a two-part process. You configure port decommissioning along with port fencing in the MAPS actions configuration, and then you configure it as an action in a MAPS rule.

To enable port decommissioning, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Create a rule or action as follows:
    - Enter **mapsconfig --actions fence,decom** to create an action for the entire switch.
    - Use the **mapsrule --create** *new_rule_name* **-group** *group_name* **-monitor** *monitor_value* **-timebase** *time_unit* **-op** *comparison_operator* **-value** *comp_op_value* **-action fence,decom** command to create a rule.

The following example enables port fencing and port decommissioning for a switch and then displays the confirmation.

```
switch246:FID128:admin> mapsconfig --actions fence,decom
switch246:admin> mapsconfig --show
Configured Notifications:       FENCE,DECOM
Mail Recipient:                 Not Configured
Paused members :
===============
PORT :
CIRCUIT :
SFP :
```

The following example makes port fencing and port decommissioning part of a rule and then displays the confirmation.

```
switch246:FID128:admin> mapsrule --create crc_decom -group ALL_E_PORTS -monitor CRC -timebase min -op g -
value 3 -action raslog,fence,decom
switch246:admin> mapsrule --show crc_decom
Rule Data:
----------
RuleName: crc_decom
Condition: ALL_E_PORTS(CRC/min>3)
Actions: raslog,fence,decom
Associated Policies:
```

Port decommissioning and firmware downgrades

- If there are any default policy rules present with port decommissioning configured, the firmware downgrade is not blocked, because in this case, the decommissioning rules are mapped to the port fencing rules of the previous version of Fabric OS. That is, a default MAPS rule from the current version of Fabric OS with port commissioning specified will remain mapped to the same rule but without port decommissioning as an action when the switch is downgraded to a previous version of Fabric OS.
- Currently, the decommission action is present for the port monitoring rules in dflt_aggressive_policy. When the switch is rebooted using a previous version of Fabric OS, the default rules in dflt_aggressive_policy which had port decommissioning specified will have port fencing specified.
- User-defined rules in the active policy are not checked to see if they have port decommissioning configured, because user-defined rules in the active policy are present only in memory and are erased as soon as a different policy is enabled, whether in they are in the current version of Fabric OS or any earlier version.

Enabling port fencing

Port fencing in MAPS can be either an action that is part of the overall switch configuration, or part of a specific rule. If it is part of the overall switch configuration, it will happen any time the port fails, while if it is part of a rule, the port will be fenced if that rule is triggered. Multiple rules can have port fencing as an action; it will happen if any of them are triggered.

To enable port fencing, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Create a rule or action as follows:
   - To set up a port fencing action for the entire switch, enter **mapsConfig --actions fence**.
   - To create a rule for port fencing, enter **mapsRule --create** *new_rule_name* **-group** *group_name* **-monitor** *monitor_value* **-timebase** *time_unit* **-op** *comparison_operator* **-value** *comp_op_value* **-action fence**.

The following example enables port fencing on a switch and then displays the confirmation.

```
switch1234:admin> mapsconfig --actions raslog,fence
switch1234:admin> mapsconfig --show

Configured Notifications:       RASLOG,FENCE
Mail Recipient:                 Not Configured
Paused members :
===============
PORT :
CIRCUIT :
SFP :
```

The following example makes port fencing part of a rule and then displays the confirmation.

```
switch1234:admin> mapsrule --create crc_fence_Eport -group ALL_E_PORTS -monitor CRC -timebase min -op g -
value 3 -action raslog,fence
switch:admin> mapsrule --show crc_fence_Eport

Rule Data:
----------
RuleName: crc_fence_Eport
Condition: ALL_E_PORTS(CRC/min>3)
Actions: raslog,fence
Associated Policies:
```

Port toggling

The "Toggle" action temporarily disables a port and then re-enables it, allowing the port to reset and recover from some device based issues. If the issue does not get resolved, the port toggling action will suspend the port for a longer period of time, forcing the port traffic to switch over to a different path if one is available.

> NOTE
> The Toggle action and the SDDQ action are mutually exclusive.

Refer to Port toggling support on page 120 for a more complete discussion of this action.

RASLog messages

Following an event, MAPS adds an entry to the switch event log for an individual switch. The RASLog stores event information but does not actively send alerts. You can use the **errShow** command to view the RASLog.

MAPS triggers RASLog messages MAPS-1001 to MAPS-1004 when the condition in a rule is true for regular counters or when the errors are above the threshold value. Depending on the state and the condition set, RASLog generates INFO, WARNING, CRITICAL, or ERROR messages.

TABLE 28 RASLog message category for non-state-based monitoring systems

Condition description	RASLog message category	Example
A rule with ">", ">=" or "==" condition	Generates a WARNING (MAPS-1003) message.	LOSS_SIGNAL monitoring system **Exception:** Class 3 Transmission Timeouts (C3TX_TO), where the ">" and ">=" condition generates an ERROR (MAPS-1002) message and a " ==" condition generates a WARNING (MAPS-1003) message.
A rule with "<" or "<=" condition	Generates an INFO (MAPS-1004) message.	

TABLE 29 RASLog message category for state-based monitoring systems

Condition description	RASLog message category	Example
A rule with " ==" or "!=" condition	Generates a WARNING (MAPS-1003) message.	LOSS_SIGNAL monitoring system **Exception:** FPI monitoring for the IO_PERF_IMPACT state, where the "==" and "!=" generates a WARNING (MAPS-1003) message and a "==" for IO_FRAME_LOSS state generates a CRITICAL (MAPS-1001) message.

In Fabric OS 8.0.1 release, MAPS provides the port name information as part of RASLog.

```
2015/06/25-21:11:43, [MAPS-1003], 239, FID 128, WARNING, odin82, PortName, F-Port 0,
Condition=ALL_OTHER_F_PORTS(LF/min>5), Current Value:[LF,100], RuleName=defALL_OTHER_F_PORTSLF_5,
Dashboard Category=Port Health.
```

Refer to the *Fabric OS Message Reference* for a complete listing and explanation of MAPS-related RASLog messages.

SFP marginal

The **SPF_MARGINAL** action sets the state of the affected small form-factor **SPF_MARGINAL** pluggable (SFP) transceiver in the MAPS dashboard to "Green" or "Yellow." The word "Green" indicates that the transceiver is operating normally; the word "Yellow" indicates it is operating outside the normal range.

This action is valid only in the context of Advanced SFP groups.

> NOTE
> This applies only to Brocade-branded SFPs having speeds greater than or equal to 10G.

Example of health status output before triggering SFP_MARGINAL

A MAPS rule can contain the **SFP_MARGINAL** action. Before the action is triggered, the health status of the SFP, as shown by the **sfpshow -health** command, is not affected. By default, the status displays "Green."

The following example shows the output of the **sfpshow -health** command *before* the **SFP_MARGINAL** action in a MAPS rule has been triggered. Notice that the health status of the first two ports is "Green" in the first example, but it changes to "Yellow" in the second example.

```
switch:admin> sfpshow -health |grep 32_G
Port03: id (sw) Vndr: BROCADE Ser.No: JAA11521007R1 Speed: 8,16,32_Gbps Health: Green
Port19: id (sw) Vndr: BROCADE Ser.No: JAA1152100181 Speed: 8,16,32_Gbps Health: Green
Port60: id 3/0 (sw) Vndr: BROCADE Ser.No: ZTA1151600009 Speed: 32_Gbps Health: Green
Port61: id 3/1 (sw) Vndr: BROCADE Ser.No: ZTA1151600009 Speed: 32_Gbps Health: Green
Port62: id 3/2 (sw) Vndr: BROCADE Ser.No: ZTA1151600009 Speed: 32_Gbps Health: Green
Port63: id 3/3 (sw) Vndr: BROCADE Ser.No: ZTA1151600009 Speed: 32_Gbps Health: Green
switch:admin>
```

Example of health status output after triggering SFP_MARGINAL

Whenever a MAPS rule is triggered that contains **SFP_MARGINAL** as an action, the health status shown by the **sfpshow -health** command is affected.

The following example shows the output of the **sfpshow -health** command *after* the **SFP_MARGINAL** action in a MAPS rule has been triggered for Port 03 and Port 19. Notice that the health status of the these two ports is now "Yellow."

```
switch:admin> sfpshow -health |grep 32_G
Port03: id (sw) Vndr: BROCADE Ser.No: JAA11521007R1 Speed: 8,16,32_Gbps Health: Yellow
```

```
Port19: id (sw) Vndr: BROCADE Ser.No: JAA1152100181 Speed: 8,16,32_Gbps Health: Yellow
Port60: id 3/0 (sw) Vndr: BROCADE Ser.No: ZTA1151600009 Speed: 32_Gbps Health: Green
Port61: id 3/1 (sw) Vndr: BROCADE Ser.No: ZTA1151600009 Speed: 32_Gbps Health: Green
Port62: id 3/2 (sw) Vndr: BROCADE Ser.No: ZTA1151600009 Speed: 32_Gbps Health: Green
Port63: id 3/3 (sw) Vndr: BROCADE Ser.No: ZTA1151600009 Speed: 32_Gbps Health: Green
switch:admin>
```

Slow Drain Device Quarantine

The Slow Drain Device Quarantine (SDDQ) action moves the traffic destined to a port affected by device based latency to a low-priority virtual channel. This action does not disable the port, but reduces the effect of its latency on other flows in the fabric.

SDDQ actions can be configured to only monitor rules, for example the DEV_LATENCY_IMPACT state of IO_PERF_IMPACT and IO_FRAME_LOSS.

> NOTE
> The SDDQ action and the Toggle action are mutually exclusive.

Refer to Slow Drain Device quarantining on page 121 for a more complete discussion of this action.

Switch critical

The “switch critical” action sets the state of the affected switch in the MAPS dashboard display to SW_CRITICAL. This action does not bring the switch down, but only affects what is displayed in the dashboard. This action is valid only in the context of Switch Status Policy-related rules.

> NOTE
> You do not need to configure the SW_CRITICAL action with **mapsconfig**; it is already pre-defined. However, when you create a rule, you must specify this action.

Switch marginal

The “switch marginal” action sets the state of the affected switch in the MAPS dashboard to SW_MARGINAL. This action does not affect the actual state of the switch, but only affects what is displayed in the dashboard. This action is valid only in the context of Switch Status Policy-related rules.

> NOTE
> You do not need to configure the SW_MARGINAL action with **mapsconfig**; it is already pre-defined. However, when you create a rule, you must specify this action.

MAPS policies overview

A MAPS policy is a set of rules. When you enable a policy, all of the rules in the policy are in effect. Refer to MAPS rules overview on page 52 for more information about MAPS rules.

A switch can have multiple policies. For example, you can have a policy for everyday use and you can have another policy for when you are running backups or performing switch maintenance.

The following restrictions apply to policies:

- Only one policy can be active at a time.
- When you enable a policy, it becomes the active policy and the rules in the active policy take effect.
- One policy must always be active on the switch.

- You can have an active policy with no rules, but you must have an active policy.
- You cannot disable the active policy. You can only change the active policy by enabling a different policy.

Viewing policy values

You can display the values for a policy by using the **mapspolicy --show** *policy_name* **|grep** *group_name* command.

The following example displays all the thresholds for host ports in the My_all_hosts_policy.

```
switch:admin>  mapspolicy --show My_all_hosts_policy | grep HOST
defALL_HOST_PORTSC3TXTO_10     |ALL_HOST_PORTS(C3TXTO/MIN>10)      |FENCE,SNMP,EMAIL |
defALL_HOST_PORTSC3TXTO_3      |ALL_HOST_PORTS(C3TXTO/MIN>3)       |RASLOG,SNMP,EMAIL|
defALL_HOST_PORTSCRC_10        |ALL_HOST_PORTS(CRC/MIN>10)         |RASLOG,SNMP,EMAIL|
defALL_HOST_PORTSCRC_20        |ALL_HOST_PORTS(CRC/MIN>20)         |FENCE,DECOM,SNMP |
defALL_HOST_PORTSITW_21        |ALL_HOST_PORTS(ITW/MIN>21)         |RASLOG,SNMP,EMAIL|
defALL_HOST_PORTSITW_40        |ALL_HOST_PORTS(ITW/MIN>40)         |DECOM,SNMP,EMAIL |
defALL_HOST_PORTSLF_3          |ALL_HOST_PORTS(LF/MIN>3)           |RASLOG,SNMP,EMAIL|
defALL_HOST_PORTSLOSS_SIGNAL_3 |ALL_HOST_PORTS(LOSS_SIGNAL/MIN>3)  |RASLOG,SNMP,EMAIL|
defALL_HOST_PORTSLOSS_SYNC_3   |ALL_HOST_PORTS(LOSS_SYNC/MIN>3)    |RASLOG,SNMP,EMAIL|
defALL_HOST_PORTSLR_10         |ALL_HOST_PORTS(LR/MIN>10)          |FENCE,DECOM,SNMP |
defALL_HOST_PORTSLR_5          |ALL_HOST_PORTS(LR/MIN>5)           |RASLOG,SNMP,EMAIL|
defALL_HOST_PORTSPE_3          |ALL_HOST_PORTS(PE/MIN>3)           |RASLOG,SNMP,EMAIL|
defALL_HOST_PORTSPE_7          |ALL_HOST_PORTS(PE/MIN>7)           |FENCE,DECOM,SNMP |
defALL_HOST_PORTSRX_75         |ALL_HOST_PORTS(RX/HOUR>75)         |RASLOG,SNMP,EMAIL|
defALL_HOST_PORTSSTATE_CHG_10  |ALL_HOST_PORTS(STATE_CHG/MIN>10)   |FENCE,DECOM,SNMP |
defALL_HOST_PORTSSTATE_CHG_5   |ALL_HOST_PORTS(STATE_CHG/MIN>5)    |RASLOG,SNMP,EMAIL|
defALL_HOST_PORTSTX_75         |ALL_HOST_PORTS(TX/HOUR>75)         |RASLOG,SNMP,EMAIL|
defALL_HOST_PORTSUTIL_75       |ALL_HOST_PORTS(UTIL/HOUR>75)       |RASLOG,SNMP,EMAIL|
```

Predefined policies

MAPS provides four predefined policies that you can neither modify nor delete.

The predefined policies are as follows:

- dflt_conservative_policy

  This policy contains rules with more lenient thresholds that allow a buffer and do not immediately trigger actions. Use this policy in environments where the elements are resilient and can accommodate errors. This policy can be used in Tape target setups.

- dflt_moderate_policy

  This policy contains rules with thresholds values between the aggressive and conservative policies.

- dflt_aggressive_policy

  This policy contains rules with very strict thresholds. Use this policy if you need a pristine fabric (for example, FICON fabrics).

- dflt_base_policy

  NOTE
  If you have installed a Fabric Vision license, then you should use the conservative, aggressive, or moderate policies. Use the base policy only for basic monitoring, similar to using MAPS without a license. Refer to Features that do not require a Fabric Vision license on page 24 for details.)

This policy contains rules that monitor the unlicensed features which were made available earlier through Fabric Watch. Refer to Features that do not require a Fabric Vision license on page 24 for a description of these features.

Although you cannot modify these predefined policies, you can create a policy based on these policies that you can modify. You can create a policy based on these policies. For more information, refer to the following links:

- Modifying a default policy on page 67
- Creating a policy on page 65
- User-defined policies on page 63

MAPS automatically monitors the management port (Eth0 or Bond0), because the rule for Ethernet port monitoring is present in all four default policies. While the default policies cannot be modified, the management port monitoring rules can be removed from cloned policies.

For System z and FICON environments, Brocade recommends that you start with the Aggressive policy. For Open Systems environments and other environments, Brocade recommends that you start with the Moderate policy. For the IO Analytics switch, the default policy is dflt_conservative_policy.

Default MAPS policy rules

Each of the predefined default policies has its own rule set.

To view the rules for a policy, enter **mapsPolicy --show** followed by the name of the policy.

User-defined policies

MAPS allows you to define your own policies. You can create a policy and add rules to it, or you can clone one of the default policies and modify the cloned policy.

Refer to Working with MAPS policies on page 64 for information on working with user-defined policies.

Fabric Watch legacy policies

When you migrate from Fabric Watch to MAPS, the following three policies are automatically created if you have used **mapsConfig --fwconvert**. If you do not use this command, then these policies are not created.

- fw_custom_policy

  This policy contains all of the monitoring rules based on the custom thresholds configured in Fabric Watch.

- fw_default_policy

  This policy contains all of the monitoring rules based on the default thresholds configured in Fabric Watch.

- fw_active_policy

  This policy contains all of the monitoring rules based on the active thresholds in Fabric Watch at the time of the conversion.

These policies are treated as user-defined policies. You can modify them by adding and deleting rules, and you can delete them.

The following factors also apply to Fabric Watch conversions:

- Converted active Fabric Watch policies reference either custom or default Fabric Watch rules.
- No custom rules are created if the "custom" thresholds are the same as the default thresholds. Instead, the default Fabric Watch rule will be referenced in the fw_custom_policy.
- Converted rules are prefixed with “fw_def_*name*” or “fw_cust_*name*”. The value for *name* is a string based on the Fabric Watch class, the area, threshold criteria (above high or below low), and the threshold number. This is the same pattern that MAPS rules use.

Working with MAPS policies

The following sections discuss viewing, creating, enabling, and modifying MAPS policies.

Viewing policy information

MAPS allows you to view the policies on a switch. You can use this command to show all policies, only a particular policy, or a summary.

To view the MAPS policies on a switch, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Choose from the following options:
   - To view a summary of all the policies on the switch, enter **mapspolicy --show -summary**.
   - To view the features of all the policies on the switch, enter **mapspolicy --show -all**.
   - To view the features of a specific policy on the switch, enter **mapspolicy --show** *policy_name*.

The following example shows the result of using the **--show -summary** option.

```
switch:admin> mapspolicy --show -summary

        Policy Name                    Number of Rules
-------------------------------------------------------------
dflt_aggressive_policy           :               204
dflt_conservative_policy         :               206
dflt_moderate_policy             :               206
dflt_base_policy                 :                20
fw_default_policy                :               109
fw_custom_policy                 :               109
fw_active_policy                 :               109
Active Policy is 'dflt_base_policy'.
```

The following example shows the result of entering **mapspolicy--show dflt_base_policy** , the active policy.

```
switch:admin>admin> mapspolicy --show dflt_base_policy

Policy Name: dflt_base_policy
Rule Name                          |Condition                                  |Actions                |
-----------------------------------------------------------------------------------------------------
defALL_TSTEMP_OUT_OF_RANGE         |ALL_TS(TEMP/NONE==OUT_OF_RANGE)            |RASLOG,SNMP,EMAIL      |
defCHASSISFLASH_USAGE_90           |CHASSIS(FLASH_USAGE/NONE>=90)              |RASLOG,SNMP,EMAIL      |
defCHASSISMEMORY_USAGE_75          |CHASSIS(MEMORY_USAGE/NONE>=75)             |RASLOG,SNMP,EMAIL      |
defCHASSISCPU_80                   |CHASSIS(CPU/NONE>=80)                      |RASLOG,SNMP,EMAIL      |
defCHASSISBAD_TEMP_MARG            |CHASSIS(BAD_TEMP/NONE>=1)                  |SW_MARGINAL,SNMP,EMAIL |
defCHASSISBAD_TEMP_CRIT            |CHASSIS(BAD_TEMP/NONE>=2)                  |SW_CRITICAL,SNMP,EMAIL |
defCHASSISBAD_PWR_CRIT             |CHASSIS(BAD_PWR/NONE>=2)                   |SW_CRITICAL,SNMP,EMAIL |
defCHASSISBAD_FAN_MARG             |CHASSIS(BAD_FAN/NONE>=1)                   |SW_MARGINAL,SNMP,EMAIL |
defCHASSISBAD_FAN_CRIT             |CHASSIS(BAD_FAN/NONE>=2)                   |SW_CRITICAL,SNMP,EMAIL |
defALL_PSPS_STATE_FAULTY           |ALL_PS(PS_STATE/NONE==FAULTY)              |RASLOG,SNMP,EMAIL      |
defALL_PSPS_STATE_ON               |ALL_PS(PS_STATE/NONE==ON)                  |RASLOG,SNMP,EMAIL      |
defALL_PSPS_STATE_OUT              |ALL_PS(PS_STATE/NONE==OUT)                 |RASLOG,SNMP,EMAIL      |
defALL_FANFAN_STATE_FAULTY         |ALL_FAN(FAN_STATE/NONE==FAULTY)            |RASLOG,SNMP,EMAIL      |
defALL_FANFAN_STATE_ON             |ALL_FAN(FAN_STATE/NONE==ON)                |RASLOG,SNMP,EMAIL      |
defALL_FANFAN_STATE_OUT            |ALL_FAN(FAN_STATE/NONE==OUT)               |RASLOG,SNMP,EMAIL      |
*defALL_PORTSSFP_STATE_FAULTY      |ALL_PORTS(SFP_STATE/NONE==FAULTY)          |RASLOG,SNMP,EMAIL      |
*defALL_PORTSSFP_STATE_OUT         |ALL_PORTS(SFP_STATE/NONE==OUT)             |RASLOG,SNMP,EMAIL      |
*defALL_PORTSSFP_STATE_IN          |ALL_PORTS(SFP_STATE/NONE==IN)              |RASLOG,SNMP,EMAIL      |
defCHASSISETH_MGMT_PORT_STATE_DOWN |CHASSIS(ETH_MGMT_PORT_STATE/NONE==DOWN)    |RASLOG,SNMP,EMAIL      |
defCHASSISETH_MGMT_PORT_STATE_UP   |CHASSIS(ETH_MGMT_PORT_STATE/NONE==UP)      |RASLOG,SNMP,EMAIL      |
Active Policy is 'dflt_base_policy'.
Unmonitored Rules are prefixed with "*"
```

The following example shows an excerpted result of using the **--show -all** option. The entire listing is too long (over 900 lines) to include.

```
switch:admin> mapspolicy --show -all

Rule Name                          |Condition                                  |Actions                |
-----------------------------------------------------------------------------------------------------
dflt_aggressive_policy:
defNON_E_F_PORTSCRC_0              |NON_E_F_PORTS(CRC/MIN>0)                   |RASLOG,SNMP,EMAIL      |
defNON_E_F_PORTSCRC_2              |NON_E_F_PORTS(CRC/MIN>2)                   |FENCE,SNMP,EMAIL       |
defNON_E_F_PORTSITW_15             |NON_E_F_PORTS(ITW/MIN>15)                  |RASLOG,SNMP,EMAIL      |
… [+201 lines]
dflt_conservative_policy:
-----------------------------------------------------------------------------
defNON_E_F_PORTSCRC_21             |NON_E_F_PORTS(CRC/MIN>21)                  |RASLOG,SNMP,EMAIL      |
defNON_E_F_PORTSCRC_40             |NON_E_F_PORTS(CRC/MIN>40)                  |FENCE,SNMP,EMAIL       |
defNON_E_F_PORTSITW_41             |NON_E_F_PORTS(ITW/MIN>41)                  |RASLOG,SNMP,EMAIL      |
… [+203 lines]
dflt_moderate_policy:
-----------------------------------------------------------------------------
defNON_E_F_PORTSCRC_10             |NON_E_F_PORTS(CRC/MIN>10)                  |RASLOG,SNMP,EMAIL      |
defNON_E_F_PORTSCRC_20             |NON_E_F_PORTS(CRC/MIN>20)                  |FENCE,SNMP,EMAIL       |
defNON_E_F_PORTSITW_21             |NON_E_F_PORTS(ITW/MIN>21)                  |RASLOG,SNMP,EMAIL      |
… [+204 lines]
dflt_base_policy:
-----------------------------------------------------------------------------
defALL_TSTEMP_OUT_OF_RANGE         |ALL_TS(TEMP/NONE==OUT_OF_RANGE)            |RASLOG,SNMP,EMAIL      |
defCHASSISFLASH_USAGE_90           |CHASSIS(FLASH_USAGE/NONE>=90)              |RASLOG,SNMP,EMAIL      |
defCHASSISMEMORY_USAGE_75          |CHASSIS(MEMORY_USAGE/NONE>=75)             |RASLOG,SNMP,EMAIL      |
… [+17 lines]

Active Policy is 'dflt_moderate_policy'.
Unmonitored Rules are prefixed with "*"
```

The following example shows the result of the **mapspolicy --show dflt_base_policy** command with the **-concise** option; this displays legends instead of complete action names in the output.

```
switch:admin> mapspolicy --show –dflt_base_policy -concise

Rule Name                          |Condition                                  |Actions     |
-----------------------------------------------------------------------------------------------
defALL_TSTEMP_OUT_OF_RANGE         |ALL_TS(TEMP/NONE==OUT_OF_RANGE)            |RS,SN,EM    |
defCHASSISFLASH_USAGE_90           |CHASSIS(FLASH_USAGE/NONE>=90)              |RS,SN,EM    |
defCHASSISMEMORY_USAGE_75          |CHASSIS(MEMORY_USAGE/NONE>=75)             |RS,SN,EM    |
… [+17 lines]

Legend:
RS:RASLOG  EML:EMAIL  SN:SNMP  PF:FENCE  SWD:SW_DOWN  SWM:SW_MARGINAL  SFPM:SFP_MARGINAL  PD:DECOM
FM:FMS  PT:TOGGLE  SQ:SDDQ
```

Creating a policy

In many cases, you must have multiple different policies available. For example, you can apply a different set of rules when maintenance operations are in progress from those that are in place for normal operations. Fabric OS allows you to create multiple policies beforehand and then easily switch between policies when necessary.

> NOTE
> When you create a policy, the policy is automatically saved, but not enabled. The policy is not enabled unless you explicitly enable it. Policy names are not case-sensitive; My_Policy and my_policy are considered to be the same.

To create policies and then add rules to them, complete the following steps.

1. Create a new policy or clone a policy from one of your existing policies.
   - To create a new policy, enter **mapsPolicy --create** *policy_name* to create a policy.
   - To clone an existing policy, enter **mapsPolicy --clone** *policy_name* **-name** *clone_policy_name*.

2. Create or modify rules to configure the required thresholds in the new policy.
    - To create a rule, enter **mapsRule --create** *rule_name* **-group** *group_name* **-monitor** *ms name* **-timebase** *timebase* **-op** *op_value* **-value** *value* **-action** *action* **-policy** *policy_name*.
    - To clone an existing rule, enter **mapsRule --clone** *rule_name* **-name** *clone_rule_name*.
    - To modify existing rules, enter **mapsRule --config** *rule_name parameters*.

The following example creates a policy by cloning another policy, and then adds a rule to the new policy.

```
switch:admin> mapspolicy --clone defpol -name backup_pol

switch:admin> mapsrule --create chassiscpu -monitor CPU -group chassis -op ge -value 70 -action raslog -
policy backup_pol
```

Enabling a policy

Only one policy can be enabled at a time, and it must be enabled before it takes effect.

> NOTE
> If the active policy is changed, or if the rules in the active policy are changed, the active policy must be re-enabled for the changes to take effect.

To enable a policy, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **mapspolicy --enable** followed by the name of the policy you want to enable.
   The previously enabled policy is automatically disabled and the specified policy is then enabled. There is no confirmation of the change.

The following example enables the "dflt_aggressive_policy" policy.

```
switch:admin> mapspolicy --enable dflt_aggressive_policy
```

Modifying a user-defined policy

It is possible to modify existing policies. For example, you may need to modify a policy if elements in the fabric change or if threshold configurations needed to be modified to catch certain error conditions.

To modify a policy and its associated rules, complete the following steps.

1. Modify the rules in the policy based on your requirements.

   You cannot modify the default rules, but you can add rules to and delete rules from the policy, and you can create rules and add them to the policy.
    - Use **mapspolicy** to add rules to and delete rules from the policy.
    - Use **mapsrule** to modify rules or to create rules and add them to the policy.
2. Optional: Even if the policy is the active policy, you must re-enable the policy using the **mapspolicy --enable** *policy_name* command for the changes to take effect. Adding a new rule or changing an existing rule in the active policy does not take effect until you re-enable the policy.

The following example adds a rule to the policy named daily_policy, displays the policy, and then re-enables the policy so the change can become active.

```
switch:admin> mapspolicy --addrule daily_policy -rulename check_crc

switch:admin> mapspolicy --show daily_policy

Policy Name: daily_policy
Rule List  :
                check_crc
                defALL_E_PORTSITW_21
                defALL_E_PORTSITW_40
                myCHASSISFLASH_USAGE_90
Active Policy is 'daily_policy'

switch:admin> mapspolicy --enable daily_policy
```

Modifying a default policy

You cannot modify any of the predefined MAPS policies, but you can clone one to create a new policy, and then modify that new policy.

To create and activate a modified version of a default policy, complete the following steps.

1. Create a copy of the default policy.

   ```
   switch:admin> mapspolicy --clone dflt_conservative_policy -name my_policy
   ```

2. Modify the rules in the policy based on your requirements.

   You cannot modify the default rules, but you can add rules to and delete rules from the policy, and you can create or clone rules and add them to the policy.

   Use **mapsPolicy** to add and delete rules to and from the policy. Use **mapsRule** to create rules and add them to the policy.

3. Enable the policy.

   ```
   switch:admin> mapspolicy --enable my_policy
   ```

   The previously enabled policy is disabled, and the specified policy is enabled.

The following example clones the default policy, deletes two rules, and modifies a rule to send an e-mail message in addition to a RASLog entry.

```
switch:admin> mapspolicy --clone dflt_conservative_policy -name rule_policy

switch:admin> mapspolicy --delrule rule_policy -rulename defCHASSISFLASH_USAGE_90

switch:admin> mapspolicy --delrule rule_policy -rulename defCHASSISMEMORY_USAGE_75

switch:admin> mapsrule --clone myCHASSISFLASH_USAGE_90 -monitor flash_usage -group chassis -timebase none -
op ge -value 90 -action raslog,email -policy rule_policy

switch:admin> mapspolicy --enable rule_policy
```

Automatic creation of MAPS rules and policies

MAPS automatically generates a set of monitoring rules, groups, and policies, which are stored in a configuration file on each Brocade device.

Following rules are added to automate MAPS configuration file:

- Common monitoring rules: applicable to all the platforms.

  - `defSWITCHSEC_DCC_4, defSWITCHSEC_FCS_0, defSWITCHSEC_FCS_2, defSWITCHSEC_FCS_4, defSWITCHSEC_HTTP_0`

- Chassis monitoring rules: applicable to only chassis platforms.
  - `defALL_WWNWWN_FAULTY, defALL_WWNWWN_ON, defALL_WWNWWN_OUT, defCHASSISDOWN_CORE_1, defCHASSISDOWN_CORE_2, defCHASSISFAULTY_BLADE_1, defCHASSISHA_SYNC_0, defCHASSISWWN_DOWN_1`
- Fixed-port switch monitoring rules: applicable to only fixed-port platforms.
  - `defCHASSISBAD_PWR_MARG`
- ASIC monitoring rules: applicable to specific ASIC platforms.
- FCIP monitoring rules: applicable to only FCIP platforms.
  - `defALL_CIRCUITS_JITTER_PER_, defALL_CIRCUITS_JITTER_PER_15, defALL_CIRCUITS_JITTER_PER_20, defALL_CIRCUITS_RTT_250`
- Extension monitoring rules: applicable to only extension platforms.
  - `defALL_CIRCUITS_IP_JITTER_PER_, defALL_CIRCUITS_IP_JITTER_PER_15, defALL_CIRCUITS_IP_JITTER_PER_20`
- BE ports monitoring rules: applicable to only platforms which supports BE ports.

Some rules get replaced with other rules. For example, if two rules having the same functionality exist on two platforms, then one rule is replaced with the other, so that both platforms have the same rule names.

Working with MAPS rules and actions

MAPS allows you to view, create, modify, and delete rules, and enable or disable actions.

Viewing MAPS rules

MAPS allows you to display a listing of all the MAPS rules on a switch or the details of a single MAPS rule.

To view the MAPS rules on a switch, perform the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Choose from the following options:
   - To view all the MAPS rules on the switch, enter **mapsrule --show -all**. This displays all the rules on the switch, listing the rule name, the actions in the rule, and the threshold condition that triggers the rule.
   - To view the details of a specific MAPS rule on the switch, enter **mapsrule --show** *rule_name*. This displays the rule name, the actions in the rule, the threshold condition that triggers the rule, and the names of any policies associated with the rule. If the rule is not associated with any policies, nothing is shown for the associated policies.

The following example shows all rules on the switch. Notice that the policies are not shown in the output.

```
switch:admin> mapsrule --show -all
-------------------------------------------------------------------
RuleName              Action                  Condition
-------------------------------------------------------------------
Rule1                 Raslog, Fence, SNMP     Switch(SEC_IDB/Min>0)
Rule2                 Raslog                  Switch(SEC_IDB/Hour>1)
NewRule1              Raslog, Fence, SNMP     Switch(SEC_IDB/Min>0)
NewRule2              Raslog, Fence, SNMP     Switch(SEC_IDB/Hour>1)
```

The following example shows the policy names associated with the rule name "Rule1".

```
switch:admin> mapsrule --show Rule1
Rule Data:
----------
RuleName: Rule1
Action: Raslog, Fence, SNMP
Condition: Switch(SEC_IDB/Min>0)
Associated Policies: daily_policy, crc_policy
```

The following example shows the result of using the **mapsrule --show -all** command with the **-concise** option; this displays abbreviations instead of complete action names in the output. It also displays a legend explaining each abbreviated action name.

```
switch:admin> mapsrule --show -all -concise
Rule Name                      |Condition                                |Actions         |
-------------------------------------------------------------------------------------------
defALL_32GSWL_SFPSFP_TEMP_85   |ALL_32GSWL_SFP(SFP_TEMP/NONE>=85)   |SFPM,RS,SN,EML  |
defALL_32GSWL_SFPSFP_TEMP_n5   |ALL_32GSWL_SFP(SFP_TEMP/NONE<=-5)   |SFPM,RS,SN,EML  |
defALL_32GSWL_SFPTXP_1259      |ALL_32GSWL_SFP(TXP/NONE>=1259)      |SFPM,RS,SN,EML  |
defALL_32GSWL_SFPVOLTAGE_3000  |ALL_32GSWL_SFP(VOLTAGE/NONE<=3000)  |SFPM,RS,SN,EML  |
defALL_32GSWL_SFPVOLTAGE_3600  |ALL_32GSWL_SFP(VOLTAGE/NONE>=3600)  |SFPM,RS,SN,EML  |
defALL_ASICS_VTAP_IOPS_250K    |ALL_ASICS(VTAP_IOPS/SEC>250000)     |RS,SN,EML       |
defALL_D_PORTSCRC_1            |ALL_D_PORTS(CRC/MIN>1)              |RS,SN,EML       |
defALL_D_PORTSCRC_2            |ALL_D_PORTS(CRC/MIN>2)              |RS,SN,EML       |
defALL_D_PORTSCRC_3            |ALL_D_PORTS(CRC/MIN>3)              |RS,SN,EML       |
defALL_D_PORTSCRC_D1000        |ALL_D_PORTS(CRC/DAY>1000)           |RS,SN,EML       |
defALL_D_PORTSCRC_D1500        |ALL_D_PORTS(CRC/DAY>1500)           |RS,SN,EML       |
defALL_D_PORTSCRC_D500         |ALL_D_PORTS(CRC/DAY>500)            |RS,SN,EML       |
defALL_D_PORTSCRC_H30          |ALL_D_PORTS(CRC/HOUR>30)            |RS,SN,EML       |
defALL_D_PORTSCRC_H60          |ALL_D_PORTS(CRC/HOUR>60)            |RS,SN,EML       |
defALL_D_PORTSCRC_H90          |ALL_D_PORTS(CRC/HOUR>90)            |RS,SN,EML       |


Legend:
RS:RASLOG  SN:SNMP  EML:EMAIL  PF:FENCE  PL:PORTLOG  PD:DECOM  FMS:FMS  PT:TOGGLE
SDDQ:SDDQ  SWD:SW_CRITICAL
```

Creating a rule

Each MAPS rule monitors a single condition. When you create a rule, you can choose to add it to a policy.

To create a policy rule, complete the following steps.

1. Enter **mapsrule --create** *rule_name* followed by the rule parameters, and optionally the policy you want to assign it to. Rule names are not case sensitive; My_Rule and my_rule are considered to be the same.
2. Optional: Enter **mapsrule --show** *rule_name* to display the rule.
3. Optional: If you added the rule to the active policy, you must re-enable the policy for the rule to take effect by entering **mapspolicy --enable policy** *policy_name*.

> NOTE
> If you are specifying a group, the group must already exist.

Example of creating a rule to generate a RASLog message

The following example creates a rule to generate a RASLog message if the CRC counter for a group of critical ports is greater than 10 in an hour. This rule is added to the daily_policy, and the daily_policy is re-enabled for the rule to take effect.

```
switch:admin> mapsrule --create C3_timeout_rule -monitor C3TXTO -group ALL_E_PORTS
                            -value 20 -timebase day -op ge -action raslog
                            -policy User_Watch_policy

switch:admin> mapsrule --show C3_timeout_rule
Rule Data:
----------
RuleName: C3_timeout_rule
Condition: ALL_E_PORTS(C3TXTO/day>=20)
Actions:  raslog
Associated Policies: User_Watch_policy
```

Example of creating a rule for a flow

To accommodate creating a rule for a flow, **mapsrule** accepts a flow name as a value for the **-group** parameter. The following example illustrates the structure.

> NOTE
> Before you can create a rule for a flow, you must import it using the **mapsconfig --import** command.

```
switch:admin> switch:admin> mapsconfig --import io_mon_4

switch:admin> switch:admin> mapsrule --create MON_io_mon_4_rx_thrput -monitor RX_THPUT
                            -group io_mon_4 -value 0 -timebase min -op ge
                            -action raslog,email,snmp -policy User_Watch_policy
```

Modifying a MAPS rule

You can modify only user-defined MAPS rules. You cannot modify the default MAPS rules.

To modify a user-defined MAPS rule, complete the following steps.

1. Enter **mapsrule --show** *rule_name* to display the rules, so you can identify the rule you want to modify.
2. Enter **mapsrule --config** followed by the parameters you are changing to modify the rule.

   > NOTE
   > You only need to specify the parameters you are changing. Any parameters you do not specify are not changed.
3. Optional: Enter **mapsrule --show** to display the updated rule.
4. If the rule is included in the active policy you must re-enable the policy using **mapspolicy --enable policy** *policy_name* for the modified rule to take effect.

Changing one parameter

The following example changes the timebase for a rule from minutes to hours.

```
switch:admin> mapsrule --show check_crc
Rule Data:
----------
RuleName: check_crc
Condition: critical_ports(crc/minute>5)
Actions: raslog
Policies Associated: daily_policy

switch:admin> mapsrule --config check_crc -timebase hour

switch:admin> mapsrule --show check_crc
Rule Data:
----------
RuleName: check_crc
Condition: critical_ports(crc/hour>5)
Actions: raslog
Policies Associated: daily_policy
```

Changing multiple parameters

The following example modifies the rule "check_crc2" to generate a RASLog message and an e-mail message if the CRC counter for a group of critical ports is greater than 15 in an hour (rather than 10 in a minute). This rule is part of the active policy, so the policy is re-enabled for the change to take effect.

```
switch:admin> mapsrule --show check_crc2
Rule Data:
----------
RuleName: check_crc2
Condition: critical_ports(crc/minute>10)
Actions: RASLOG
Policies Associated: daily_policy

switch:admin> mapsrule --config check_crc2 -timebase hour -op g -value 15 -action raslog,email -policy
daily_policy

switch:admin> mapsrule --show check_crc2
Rule Data:
----------
RuleName: check_crc2
Condition: critical_ports(crc/hour>15)
Actions: RASLOG, EMAIL
Mail Recipient: admin@mycompany.com
Policies Associated: daily_policy

switch:admin> mapspolicy --enable daily_policy
```

Cloning a rule

You can clone both default and user-defined MAPS rules.

To clone a MAPS rule, complete the following steps.

1. Use the **mapsrule --show** *rule_name* command to display the rule you want to clone.
2. Use the **mapsrule --clone** *oldRuleName* **-rulename** *newRuleName* command to duplicate the rule.

   NOTE

   If no parameters other than **--clone** *oldRuleName* **-rulename** *newRuleName* are specified, an exact copy of the original rule is created. Otherwise, specify the parameters you want to change in the new rule.

   For more information on this command and all its parameters, refer to the *Fabric OS Command Reference*.

Creating an exact clone

The following example shows an existing rule, creates an exact clone of that rule and renames it, and then displays the new rule.

```
switch:admin> mapsrule --show defALL_HOST_PORTSCRC_20
Rule Data:
----------
RuleName: defALL_HOST_PORTSCRC_20
Condition: ALL_HOST_PORTS(CRC/MIN>20)
Actions: FENCE,DECOM,SNMP,EMAIL
Associated Policies: dflt_moderate_policy, User_created_policy


switch:admin> mapsrule --clone defALL_HOST_PORTSCRC_20 -rulename ANY_HOST_CRC_20


switch:admin> mapsrule --show ANY_HOST_CRC_20
Rule Data:
----------
RuleName: ANY_HOST_CRC_20
Condition: ALL_HOST_PORTS(CRC/MIN>20)
Actions: FENCE,DECOM,SNMP,EMAIL
Associated Policies:
```

Cloning a rule and changing its values

When you clone a rule, you can also specify the parameters you want to be different from the old rule in the new rule. The following example clones an existing rule and changes the group being monitored for the new rule. It then displays the new rule.

```
switch:admin>  mapsrule --clone defALL_HOST_PORTSCRC_20
                 -rulename Check_CRC_on_Eng_Ports_20 -group Eng_ports

switch:admin>  mapsrule --show Check_CRC_on_Eng_Ports_20
Rule Data:
----------
RuleName: Check_CRC_on_Eng_Ports_20
Condition: Eng_ports(CRC/MIN>20)
Actions: FENCE,DECOM,SNMP,EMAIL
Associated Policies:
```

Cloning a rule and changing its timebase

The following example creates a clone of a rule, changing the timebase to an hour, and then displays the new rule.

```
switch:admin> mapsrule --clone Check_CRC_on_Eng_Ports_20
                    -rulename Check_CRC_on_Eng_ports_hour -timebase hour


switch:admin> mapsrule --show Check_CRC_on_Eng_ports_hour
Rule Data:
----------
RuleName: Check_CRC_on_Eng_ports_hour
Condition: Eng_ports(CRC/hour>20)
Actions: FENCE,DECOM,SNMP,EMAIL
Associated Policies:
```

Quieting a rule

MAPS supports the concept of "quiet time" as an optional rule parameter for the **mapsrule** command. Including **-qt** *time* in a rule keeps MAPS from sending another alert based on the same rule for the length of time specified after it sends the initial alert. The parameter is supported only for SNMP traps, RASLog, or e-mail actions.

By default, MAPS continuously sends alerts if the triggering condition persists. For example, if the voltage of an SFP drops and remains below 2960 mV, then MAPS will send an alert based on the defALL_OTHER_SFPVOLTAGE_2960 rule every time the rule is checked. This might interfere with monitoring for other faults and can rapidly fill an e-mail box. If a rule for which quiet time has been set is

triggered, MAPS performs the configured and enabled actions for the rule the first time. Afterwards, if the rule is triggered again within the quiet time period, MAPS does not perform any of the actions until the quiet time expires. At that time, MAPS sends an update alert. This alert is the same as the initial alert, but it includes information about the number of times the rule was triggered in the interim. After MAPS sends the update, it resets the quiet time timer. The following example shows the additional information that is sent as part of an update notification.

```
2014/11/26-01:56:00, [MAPS-1005], 2588, FID 128, WARNING, sw0, D-Port 1, Condition=ALL_PORTS(STATE_CHG/
min>=2), Current Value:[STATE_CHG,12], RuleName=st_chg triggered 2 times in 180 duration and last trigger
time Wed Nov 26 01:53:24 2014, Dashboard Category=Port Health
```

Quiet time does not work with rules that define port fencing, port decommissioning, or I/O statistics. You can configure quiet time with the port fencing action, but MAPS will ignore the quiet time parameter in this case. Quiet time is not supported for monitoring the following values:

• IO_RD • IO_RD_BYTES • IO_WR • IO_WR_BYTES • IT_FLOW • RX_FCNT • RX_THPUT • TX_FCNT • TX_THPUT	• RD_PENDING_IO_LT_8K • RD_PENDING_IO_8_64K • RD_PENDING_IO_64_512K • RD_PENDING_IO_GE_512K • RD_STATUS_TIME_LT_8K • RD_STATUS_TIME_8_64K • RD_STATUS_TIME_64_512K • RD_STATUS_TIME_GE_512K • RD_1stDATA_TIME_LT_8K • RD_1stDATA_TIME_8_64K • RD_1stDATA_TIME_64_512K • RD_1stDATA_TIME_GE_512K	• WR_PENDING_IO_LT_8K • WR_PENDING_IO_8_64K • WR_PENDING_IO_64_512K • WR_PENDING_IO_GE_512K • WR_STATUS_TIME_LT_8K • WR_STATUS_TIME_8_64K • WR_STATUS_TIME_64_512K • WR_STATUS_TIME_GE_512K • WR_1stDATA_TIME_LT_8K • WR_1stDATA_TIME_8_64K • WR_1stDATA_TIME_64_512K • WR_1stDATA_TIME_GE_512K

The length of a quiet time period is configured in seconds, with the minimum value determined by the timebase defined in the rule. For all timebases other than NONE, the minimum quiet time is decided by the timebase present in the rule. For example, if the rule has "minute" specified as the timebase, then the minimum value for quiet time is one minute. There is no predefined upper limit. If the rule timebase is "NONE", then the quiet time value is different for different monitored systems, as shown in the following table.

TABLE 30 Minimum quiet time values for monitoring systems that support a time base of NONE

Monitoring system	Minimum quiet time (seconds)
DEV_LATENCY_IMPACT	60
DEV_NPIV_LOGINS	3600
FLASH_USAGE	60
CPU	120
MEMORY_USAGE	120
TEMP	60
PS_STATE	15
FAN_STATE	15
BLADE_STATE	15
SFP_TEMP	360
VOLTAGE	360

TABLE 30 Minimum quiet time values for monitoring systems that support a time base of NONE (continued)

Monitoring system	Minimum quiet time (seconds)
CURRENT	360
RXP	360
TXP	360
PWR_HRS	360
BAD_TEMP	60
BAD_PWR	60
BAD_FAN	60

The following example shows a rule that includes a 120-second quiet time period.

```
switch:admin>    mapsrule --config toggle_crc_rule -group ALL_PORTS -monitor CRC -timebase MIN -op ge -
value 0 -action raslog,snmp -qt 120
```

The command **mapsTrap** contains OIDs to indicate the quiet time parameters. If **mapsRuleTriggerCount** OID value is greater than 1, then **mapsTrap** is generated at the expiration of quiet time.

```
mapsRuleTriggerCount   OBJECT-TYPE
SYNTAX                 Unsigned32
MAX-ACCESS             read-only
STATUS                 current
DESCRIPTION            "Number of times rule trigger in quiet time duration"
::= { mapsConfig 15 }

mapsLastRuleExecTime OBJECT-TYPE
SYNTAX               DateAndTime
MAX-ACCESS           read-only
STATUS               current
DESCRIPTION          "Last rule execution time"
::= { mapsConfig 16 }


mapsQuietTime      OBJECT-TYPE
SYNTAX             Unsigned32
UNITS              "seconds"
MAX-ACCESS         read-only
STATUS             current
DESCRIPTION       "Quiet time configured in the rule"
::= { mapsConfig 17 }
```

Rule deletion

A rule must be removed from every policy that references it before it can be deleted.

Although you can use the **mapsrule --delete** *rule_name* command to delete individual instances of a user-defined rule, you must remove the rule individually from every policy that uses the rule before you can finally delete the rule itself. This could require a lot of tedious work if the rule has been added to many policies. To simplify the process, adding the **-force** keyword to the command allows you to delete the named user-defined rule from every policy that uses the rule before deleting the rule itself.

> NOTE
> There is a difference between using the **-force** keyword to delete a rule and using it to delete a group. When you delete a rule using this option, the rule is first removed from all policies, and then the rule itself is deleted. When you delete a group, first the rule referencing the specified group is deleted and, if the rule is part of any policies, it is deleted from those policies). Then the group is deleted. Refer to Deleting groups on page 48 for information on deleting groups.

The following example shows that the rule port_test_rule35 exists in test_policy_1. The examples show the rule being deleted from that policy using the **-force** keyword, and then it shows a verification that the rule has been deleted from the policy.

```
switch:admin> mapspolicy --show test_policy_1
Policy Name: xyz_test60
Rule List                       Action      Condition
-------------------------------------------------------------------
def_port_test_rule35         RASLOG      ALL_PORTS(CRC/min>300)
def_port_test_rule50         RASLOG      ALL_PORTS(CRC/min>650)
def_port_test_rule80         RASLOG      ALL_PORTS(CRC/min>850)
Active Policy is 'dflt_conservative_policy'.

switch:admin> mapsrule --delete port_test_rule35 -force
Execution is successful.
2014/02/02-17:55:38, [MAPS-1101], 255, FID 128, INFO, sw0, Rule port_test_35 is deleted.

switch:admin> mapspolicy --show test_policy_1
Policy Name: xyz_test60
Rule List               Action      Condition
port_test_rule50        RASLOG         ALL_PORTS(CRC/min>650)
port_test_rule80        RASLOG         ALL_PORTS(CRC/min>850)
```

Sending alerts using e-mail

E-mail alerts allow you to be notified immediately when MAPS detects that an error has occurred. There is a limit of five e-mail addresses per alert, and the maximum length for each individual e-mail address is 128 characters.

To configure MAPS to send an alert using e-mail, complete the following steps.

1. Configure and validate the e-mail server. Refer to Configuring e-mail server information on page 76 for information on specifying the e-mail server to be used.
2. Enter the **mapsconfig --emailcfg** command to set the e-mail parameters.

   To send an alert to multiple e-mail addresses, separate the addresses using a comma.

   NOTE

   You can also send a test e-mail alert. Refer to E-mail alert testing on page 76 for additional information.

Specifying e-mail address for alerts

The following example specifies the e-mail address for e-mail alerts on the switch, and then displays the settings. It assumes that you have already correctly configured and validated the e-mail server.

```
switch:admin> mapsconfig --emailcfg -address admin1@mycompany.com

switch:admin> mapsconfig --show
Configured Notifications:       RASLOG,EMAIL,FENCE,SW_CRITICAL,SW_MARGINAL
Mail Recipient:                 admin1@mycompany.com
Network Monitoring:             Enabled
Paused members :
===============
PORT :
CIRCUIT :
SFP :
```

Specifying multiple e-mail addresses for alerts

The following example specifies multiple e-mail addresses for e-mail alerts on the switch, and then displays the settings. It assumes that you have already correctly configured and validated the e-mail server.

```
switch:admin> mapsconfig --emailcfg -address admin1@mycompany.com, admin2@mycompany.com

switch:admin> mapsconfig --show
Configured Notifications:       RASLOG,EMAIL,FENCE,SW_CRITICAL
Mail Recipient:                 admin1@mycompany.com,
                                admin2@mycompany.com
Paused members :
PORT :
CIRCUIT :
SFP :
```

Clearing the configured e-mail address

To clear the configured e-mail addresses, enter **mapsconfig --emailcfg -address none**. All configured e-mail addresses will be erased.

E-mail alert testing

You can send a test e-mail message to check that you have the correct e-mail server configuration. You can use any combination of default and custom subject or message for your test e-mail message.

To verify that the MAPS e-mail feature is correctly configured, enter **mapsConfig --testmail** *optional_customizations* command. You can customize the subject and message as described in the following table.

TABLE 31 Test e-mail command parameters

Command option	Details
--testmail	MAPS sends the default test e-mail with the default subject "MAPS Welcome mail" and message text "Test mail from switch".
--testmail -subject *subject*	MAPS sends the test e-mail with the subject you provided and the default message text.
--testmail -message *message*	MAPS sends the test e-mail with the default subject and the message text you provided.
--testmail -subject *subject* **-message** *message*	MAPS sends the test e-mail with the subject and message text you provided.

For more information on this command, refer to the *Fabric OS Command Reference*.

Configuring e-mail server information

Fabric OS allows you to specify the e-mail server used to send e-mail alerts. The e-mail configuration is global at the chassis level and is common for all logical switches in the chassis.

> NOTE
> To send e-mail, the domain name system (DNS) server configuration has to be specified. Refer to the *Fabric OS Command Reference* for information on using the **dnsconfig** command.

The relay host is a smart relay server which is used to filter e-mail messages coming from outside world to the switch. If the relay host is not configured, all the e-mails from and to the switch will be handled by the DNS mail server. If a relay host is configured all the e-mails are routed through the relay host to the switch, reducing the load on the DNS mail server.

To specify the e-mail server used to send e-mail alerts, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **relayconfig --config -rla_ip** *relay IP address* **-rla_dname** "*relay domain name*". The quotation marks are required.

There is no confirmation of this action.

3. Optional: Enter **relayconfig --show**.
   This displays the configured e-mail server host address and domain name.

The following example configures the relay host address and relay domain name for the switch, and then displays it.

```
switch:admin> relayconfig --config -rla_ip 10.70.212.168 -rla_dname "mail.brocade.com"

switch:admin> relayconfig --show
Relay Host:                    10.70.212.168
Relay Domain Name:           mail.brocade.com
```

For additional information on the relay host and the **relayconfig** command, refer to the *Fabric OS Command Reference*.

Viewing configured e-mail server information

Fabric OS allows you to view the e-mail server host address and domain name configured for MAPS.

To view the e-mail server host address and domain name configured for MAPS, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Optional: Enter **relayConfig --show**.
   This displays the configured e-mail server host address and domain name.

The following example displays the configured relay host address and relay domain name for the switch.

```
switch:admin> relayconfig --show
Relay Host:                    10.70.212.168
Relay Domain Name:           mail.brocade.com
```

For additional information on the relay host and the **relayConfig** command, refer to the *Fabric OS Command Reference*.

Deleting e-mail server configuration

Fabric OS allows you to remove the e-mail server configuration used by MAPS.

To remove the e-mail server host address and domain name configured for MAPS, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **relayConfig --delete**.
   There is no confirmation of this action.
3. Optional: Enter **relayConfig --show** to confirm the deletion.

The following example deletes the configured relay host address and relay domain name for the switch, and then shows that these items have been deleted.

```
switch:admin> relayconfig --delete

switch:admin> relayconfig --show
Relay Host:
Relay Domain Name:
```

For additional information on the relay host and the **relayConfig** command, refer to the *Fabric OS Command Reference*.

MAPS Dashboard

- MAPS dashboard overview 79
- Viewing the MAPS dashboard 83
- Clearing MAPS dashboard data 95

MAPS dashboard overview

The Monitoring and Alerting Policy Suite (MAPS) dashboard provides a summary view of the switch health status that allows you to easily determine whether everything is working according to policy or whether you need to investigate further.

MAPS dashboard sections

The MAPS dashboard output is divided into the following main sections. A history section is displayed if you enter **mapsdb --show all**.

Dashboard high-level information section

The dashboard high-level information section displays basic dashboard data: the time the dashboard was started, the name of the active policy, any fenced, decommissioned, or quarantined ports, the list of FCIP circuits that are fenced, and top PIDs.

The following output extract the use of the command to display high-level dashboard information:

```
switch:admin> mapsdb --show all

1 Dashboard Information:
========================
DB start time:                  Fri Jan 08 18:38:12 2016
Active policy:                  test_xy1
Configured Notifications:       SW_CRITICAL,SW_MARGINAL,RASLOG,FENCE
Fenced Ports :                  None
Decommissioned Ports :          None
Fenced circuits :               38/0,38/1,38/2,38/3,38/4,38/5,38/6,38/7
Quarantined Ports :             3/32, 4/12
Top PIDs <pid(it-flows)>:       0x69b0c0(8) 0x697b00(4)

     (output truncated)
```

Switch Health Report section

The Switch Health Report section displays the current switch policy status and lists any factors contributing to that status as defined by the Switch Health Report rules in the active policy.

The following output extract shows a sample Switch Health Report section; revealing that the switch status is HEALTHY.

```
2 Switch Health Report:
=======================

Current Switch Policy Status: HEALTHY

     (output truncated)
```

Refer to Switch Status Policy on page 38 for more details on switch policies.

Summary Report section

The Summary Report section has two subsections, the Category report and the Rules Affecting Health report. The Category report subsection collects and summarizes the various switch statistics monitored by MAPS into multiple categories, and displays the current status of each category since midnight, and the status of each category for the past seven days. If a rule violation has caused a change in the status of a category, rule-related information is displayed in the Rules Affecting Health subsection, broken out by category.

The following categories are monitored by MAPS:

- Port Health on page 32
- Back-end Health on page 33
- FRU Health on page 34
- Security Violations on page 34
- Fabric State Changes on page 34
- Switch Resource on page 35
- Traffic Performance on page 36
- Fabric Performance Impact on page 37
-

The following output extract shows a sample Summary Report section.

```
3.1 Summary Report:
===================

Category                 |Today                    |Last 7 days               |
-------------------------------------------------------------------------------
Port Health              |In operating range       |No Errors                 |
BE Port Health           |In operating range       |No Errors                 |
GE Port Health           |No Errors                |No Errors                 |
Fru Health               |Out of operating range   |In operating range        |
Security Violations      |No Errors                |No Errors                 |
Fabric State Changes     |Out of operating range   |No Errors                 |
Switch Resource          |In operating range       |In operating range        |
Traffic Performance      |In operating range       |In operating range        |
FCIP Health              |Out of operating range   |No Errors                 |
Fabric Performance Impact|Out of operating range   |In operating range        |
```

When a category contains an "out-of-range" error, the dashboard displays a table showing the rules triggered in that category since midnight. This allows you to see more precisely where the problem is occurring. Each category in the table contains the following information:

- The number of times rules were triggered in each category
- The rules that were triggered
- The time when the rule was triggered
- The entities (ports, circuits, and others) that triggered the rule
- The values of these entities when the rule was triggered

For each category, the dashboard stores the following information for each hour since midnight:

- The five most recent distinct rule violations that occurred.
- For each rule, the five most recent entities on which the rules were triggered.

- Although a rule might be triggered multiple times within a given hour, only the timestamp of the latest violation is stored.
- However, each violation of a rule individually is reflected in the rule count for that category and the repeat count for that rule.

For example, if the same rule was triggered 12 times in one hour, the repeat count value (shown as `Repeat Count` in the following example) for that rule will be 12, but only the timestamp for the last occurrence is displayed. In addition, the last five distinct entities on which this rule was triggered are stored (and these could be stored from different instances of the rule's violation). Alternatively, if a rule was triggered 12 times since midnight, but each violation happened in a different hour, then each violation is logged separately in the dashboard.

The following output extract shows a sample Rules Affecting Health section. The column headings in the example have been edited slightly so as to allow the example to display clearly.

```
3.2 Rules Affecting Health:
===========================

Category       |Repeat|Rule Name             |Execution Time   |Object     |Triggered   |
(Rule Count)   |Count |                      |                 |           |Value(Units)|
-------------------------------------------------------------------------------------------
Fru Health(8)  |8     |defALL_PORTSSFP_      |02/04/16 18:43:43|U-Port 6/31|FAULTY      |
               |      |STATE_FAULTY          |                 |           |            |
               |      |                      |                 |U-Port 6/30|FAULTY      |
               |      |                      |                 |U-Port 6/29|FAULTY      |
               |      |                      |                 |U-Port 6/28|FAULTY      |
Fabric State   |2     |defSWITCHEPORT_DOWN_1|02/04/16 20:21:01|Switch     |2 Ports     |
Changes (4)    |      |                      |                 |           |            |
               |      |                      |                 |Switch     |2 Ports     |
               |1     |defSWITCHEPORT_DOWN_1|02/04/16 19:17:12|Switch     |2 Ports     |
               |1     |defSWITCHFLOGI_4      |02/04/16 18:43:42|Switch     |7 Logins    |
FCIP Health(2)|2     |defALL_TUNNELSSTATE_  |02/04/16 19:23:48|Tunnel 8/19|1           |
               |      |CHG_0                 |                 |           |            |
               |      |                      |                 |Tunnel 8/18|1           |
               |      |                      |                 |Tunnel 8/19|1           |
               |      |                      |                 |Tunnel 8/18|1           |
Fabric         |2     |defALL_PORTS_IO_      |02/04/16 19:08:01|E-Port 3/32|IO_LATENCY_ |
Performance    |      |LATENCY_CLEAR         |                 |           |CLEAR       |
Impact(4)      |      |                      |                 |           |            |
               |      |                      |                 |E-Port 3/29|IO_LATENCY_ |
               |      |                      |                 |           |CLEAR       |
               |      |                      |                 |E-Port 3/29|IO_LATENCY_ |
               |      |                      |                 |           |CLEAR       |
               |2     |defALL_PORTS_IO_      |02/04/16 19:07:01|E-Port 3/32|IO_PERF_    |
               |      |PERF_IMPACT           |                 |           |IMPACT      |
               |      |                      |                 |           |            |
               |      |                      |                 |E-Port 3/29|IO_PERF_    |
               |      |                      |                 |           |IMPACT      |
               |      |                      |                 |E-Port 3/29|IO_PERF_    |
               |      |                      |                 |           |IMPACT      |
```

History Data section (optional)

When displayed, the History Data section provides information on how the switch has been behaving regardless of whether rules were triggered. It contains only port-related statistics, and is the raw counter information recorded since the previous midnight.

The historical data log stores the last seven days on which errors were recorded (not the last seven calendar days, but the last seven days, irrespective of any interval between these days). If a day has no errors, that day is not included in the count or the results. Using this information, you can get an idea of the errors seen on the switch even though none of the rules might have been violated. If you see potential issues, you can reconfigure the appropriate rule thresholds to specifically fit the switch based on the actual behavior of traffic on the switch. For more information on historical data, refer to Viewing historical data on page 89.

The following output extract shows a sample History Data section.

```
(output truncated)
4.1 Front end port History Data:
================================
Stats(Units)         Current   07/21/14  07/14/14    --/--/--    --/--/--   --/--/--   --/--/--
                     Port(val) Port(val) Port(val)
-------------------------------------------------------------------------------------------
CRC(CRCs)            1/13(20)   -          -          -          -          -          -
ITW(ITWs)            -          1/13(612)  -          -          -          -          -
LOSS_SYNC(SyncLoss)  -          -          -          -          -          -          -
LF                   -          -          -          -          -          -          -
LOSS_SIGNAL(LOS)     -          -          -          -          -          -          -
PE(Errors)           -          -          -          -          -          -          -
STATE_CHG            -          -          -          -          -          -          -
C3TXTO(Timeouts)     -          -          -          -          -          -          -
RX(%)                -          -          -          -          -          -          -
TX(%)                -          -          -          -          -          -          -
UTIL(%)              -          -          -          -          -          -          -
BN_SECS(Seconds)     -          -          -          -          -          -          -

4.2 Backend port History Data:
================================
Stats(Units)         Current   07/21/14   07/14/14   --/--/--   --/--/--   --/--/--   --/--/--
                     Port(val) Port(val)  Port(val)
-------------------------------------------------------------------------------------------
CRC(CRCs)            6/8(50)     -          -          -          -          -          -
```

Notes on dashboard data

The following information should be kept in mind when examining dashboard data.

- The following dashboard state conditions can be displayed:
  - No Errors: Displayed if there are no errors for the switch ports, security, fabric, or FCIP health; for example, if no port has had an error since midnight.
  - In operating range: Displayed if there are no errors, or if there were errors but no rule was triggered.
  - Out of operating range: Displayed if at least one error triggered a rule belonging to the category in which this state message appears.
- RX, TX, UTIL errors are not displayed in the History Data section unless port errors are recorded for that day.
- The "Rule Count" value is the absolute number of different violations in that category since the previous midnight. The "Repeat Count" is the number of times a rule has been violated in the hour, for example, between 10:00:00 and 10:59:59.
- By default, only the last five violations are displayed for each category. However, entering **mapsdb --show all** causes the dashboard to display all the rule violations currently stored along with additional historical data.

MAPS dashboard display options

The **mapsdb** command allows you to the MAPS dashboard for a specific period of time. You can use various options of the **mapsdb** command to display data gathered since midnight, for any one-hour period since midnight, or for the last seven days on which errors were recorded.

Refer to the *Fabric OS Command Reference* for detailed instructions on using the **mapsdb** command options to configure the dashboard.

> NOTE
> If the MAPS license is not active, then only the results for unlicensed features will be displayed in the MAPS dashboard. Refer to MAPS commands that do not require a Fabric Vision license on page 24 for a list of these features.

Viewing the MAPS dashboard

The MAPS dashboard allows you to monitor the switch status. There are three primary views: a summary view, a detailed view (which includes historical data), and a history-only view.

To view the status of the switch as seen by MAPS, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **mapsdb --show** followed by the scope parameter: all, history, or details. Entering details allows you to specify either a specific day or a specific hour of the current day.

The following example shows a typical result of entering **mapsdb --show all**.

```
switch:admin> mapsdb --show all
1 Dashboard Information:
========================
DB start time:                  Tue Jan 19 18:38:12 2016
Active policy:                  test_xy1
Configured Notifications:       SW_CRITICAL,SW_MARGINAL, RASLOG,FENCE
Fenced Ports :                  None
Decommissioned Ports :          None
Fenced circuits :               38/0,38/1,38/2,38/3,38/4,38/5,38/6,38/7
Quarantined Ports :             None


2 Switch Health Report:
=======================
Current Switch Policy Status: HEALTHY


3.1 Summary Report:
===================
Category                 |Today                      |Last 7 days               |
-------------------------------------------------------------------------------
Port Health              |In operating range         |In operating range        |
BE Port Health           |No Errors                  |In operating range        |
Fru Health               |In operating range         |In operating range        |
Security Violations      |No Errors                  |No Errors                 |
Fabric State Changes     |No Errors                  |In operating range        |
Switch Resource          |In operating range         |In operating range        |
Traffic Performance      |In operating range         |In operating range        |
FCIP Health              |No Errors                  |No Errors                 |
Fabric Performance Impact|In operating range         |In operating range        |


3.2 Rules Affecting Health:
===========================
Category(Rule Count)|RepeatCount|Rule Name       |Execution Time   |Object     |Triggered Value       |
                                                                                  (Units)        |
-------------------------------------------------------------------------------------------------------
Switch Resource (1) |1          |defCHASSISCPU_80|03/02/15 12:10:01|Chassis    |99.00 %               |

4 History Data:
===============
Stats(Units)        Current     03/03/15    03/02/15    --/--/--    --/--/--    --/--/--    --/--/--
                    Port(val)   Port(val)   Port(val)
-------------------------------------------------------------------------------------------------------
CRC(CRCs)           -           -           -           -           -           -           -
ITW(ITWs)           -           -           -           -           -           -           -
LOSS_SYNC(SyncLoss) -           -           -           -           -           -           -
LF                  -           -           7/13(65)    -           -           -           -
                    -           -           7/14(1)     -           -           -           -
                    -           -           7/15(1)     -           -           -           -
LOSS_SIGNAL(LOS)    7/4(51)     7/4(52)     7/4(44)     -           -           -           -
                    7/5(51)     7/5(52)     7/5(44)     -           -           -           -
                    7/6(51)     7/6(52)     7/6(44)     -           -           -           -
                    7/7(51)     7/7(52)     7/7(44)     -           -           -           -
                    -           -           7/14(1)     -           -           -           -
PE(Errors)          -           -           -           -           -           -           -
STATE_CHG           -           -           7/14(2)     -           -           -           -
                    -           -           7/15(2)     -           -           -           -
                    -           -           7/13(1)     -           -           -           -
LR                  -           -           7/13(55)    -           -           -           -
                    -           -           7/14(3)     -           -           -           -
                    -           -           7/15(3)     -           -           -           -
C3TXTO(Timeouts)    -           -           -           -           -           -           -
RX(%)               2/11(38.46) 7/31(5.87)  2/11(7.39)  -           -           -           -
                    7/31(37.38) -           7/31(6.82)  -           -           -           -
                    7/0(29.11)  -           7/0(4.81)   -           -           -           -
                    7/8(2.79)   -           -           -           -           -           -
TX(%)               2/11(38.46) 2/11(6.20)  2/11(7.68)  -           -           -           -
                    7/31(37.39) 7/0(4.25)   7/31(6.54)  -           -           -           -
```

```
                    7/0(29.10)  7/8(1.89)   7/0(5.00)   -           -           -           -
                    7/8(16.74)  -           7/8(4.02)   -           -           -           -
UTIL(%)             2/11(38.46) 2/11(3.38)  2/11(7.53)  -           -           -           -
                    7/31(37.39) 7/31(3.20)  7/31(6.68)  -           -           -           -
                    7/0(29.11)  7/0(2.33)   7/0(4.90)   -           -           -           -
                    7/8(9.77)   7/8(1.01)   7/8(2.21)   -           -           -           -
BN_SECS(Seconds)    -           -           -           -           -           -           -

5 History Data for Backend ports:
=================================
Stats(Units)        Current     03/03/15    03/02/15    --/--/--    --/--/--    --/--/--    --/--/--
-------------------------------------------------------------------------------------------------------
CRC(CRCs)           -           -           -           -           -           -           -
ITW(ITWs)           -           -           -           -           -           -           -
LR                  -           -           -           -           -           -           -
BAD_OS(Errors)      -           -           2/21(42709) -           -           -           -
                    -           -           2/11(41447) -           -           -           -
                    -           -           2/4(36958)  -           -           -           -
                    -           -           2/24(36175) -           -           -           -
                    -           -           2/18(11153) -           -           -           -
FRM_LONG(Errors)    -           -           -           -           -           -           -
FRM_TRUNC(Errors)   -           -           -           -           -           -           -
```

Refer to MAPS monitoring categories on page 31 for explanations of the categories listed in the dashboard output.

Viewing a summary switch status report

A summary view provides health status at a high level and includes enough information for you to investigate further if necessary.

To view a summary switch status report, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **mapsdb --show** with no other parameters to display the summary status.

The following example displays the general status of the switch (MARGINAL) and lists the overall status of the monitoring categories for the current day (measured since midnight) and for the last seven days. If any of the categories are shown as being "Out of range", the last five rules that caused this status are listed. If a monitoring rule is triggered, the corresponding RASLog message appears under Rules Affecting Health of the dashboard. Be aware that the example column headings have been edited slightly so as to allow it to display clearly.

```
switch:admin> mapsdb --show

1 Dashboard Information:
========================

DB start time:                  Thu Feb  4 19:17:13 2016
Active policy:                  dflt_aggressive_policy
Configured Notifications:       RASLOG,EMAIL,FENCE
Fenced Ports :                  5/60,5/62
Decommissioned Ports :          None
Fenced circuits :               None
Quarantined Ports :             None
Top PIDs <pid(it-flows)>:       0x69b0c0(8) 0x697b00(4)

2 Switch Health Report:
=======================

Current Switch Policy Status: CRITICAL
Contributing Factors:
---------------------
*BAD_PWR (CRITICAL).
*BAD_FAN (MARGINAL).


3.1 Summary Report:
===================

Category                 |Today                      |Last 7 days               |
-------------------------------------------------------------------------------
Port Health              |Out of operating range     |No Errors                 |
BE Port Health           |No Errors                  |No Errors                 |
GE Port Health           |In operating range         |No Errors                 |
Fru Health               |Out of operating range     |In operating range        |
Security Violations      |No Errors                  |No Errors                 |
Fabric State Changes     |Out of operating range     |No Errors                 |
Switch Resource          |In operating range         |In operating range        |
Traffic Performance      |In operating range         |In operating range        |
FCIP Health              |No Errors                  |No Errors                 |
Fabric Performance Impact|In operating range         |In operating range        |


3.2 Rules Affecting Health:
===========================

Category     |Repeat|Rule Name      |Execution Time    |Object         |Triggered   |
(Rule Count) |Count |               |                  |               |Value(Units)|
-----------------------------------------------------------------------------------
Fru Health(2)|2     |defALL_PSPS_   |02/04/16 21:32:16 |Power Supply 3 |FAULTY      |
             |      |STATE_FAULTY   |                  |               |            |
             |      |               |                  |Power Supply 4 |FAULTY      |
```

Sub-flow rule violation summaries

In the MAPS dashboard you can view a summary of all sub-flows that have rule violations.

When a rule is triggered, the corresponding RASLog rule trigger appears in the "Rules Affecting Health" sub-section of the dashboard as part of the Traffic Performance category. In this category, the five flows or sub-flows with the highest number of violations since the previous midnight are listed.

The naming convention for "Object" in sub-flows has the format: `Flow (flow_name:sub-flow parameters)`, where *flow_name* is the name of the imported flow.

The following extract provides an illustration of violations of the "thruputflow_thput_10" rule. Some of the lines in the output have been split at the backslash (\) to allow the example to display clearly.

```
switch:admin> mapsdb --show
.
.
.
3.2. Rules Affecting Health:
================================
Category(Rule Count)     |Repeat Count|Rule Name            |Execution Time| \
Traffic Performance(10)  |5           | thruputflow_thput_10|2/21/13 1:30:6| \
                                                             2/21/13 1:30:6| \
                                                             2/21/13 1:28:6| \
                                                             2/21/13 1:26:6| \
                                                             2/21/13 1:24:6| \

    \|Object                                        |Trigger Value(Units)
    \|Flow (thruputflow:SID=011000,DID=011200,Tx=10)| 860 MBps
    \|Flow (thruputflow:SID=012000,DID=011200,Tx=10)| 707 MBps
    \|Flow (thruputflow:SID=012100,DID=011200,Tx=10)| 812 MBps
    \|Flow (thruputflow:SID=012200,DID=011200,Tx=10)| 753 MBps
    \|Flow (thruputflow:SID=012300,DID=011200,Tx=10)| 736 MBps
     (output truncated)
```

- For *learning* flows, in addition to the name of the flow being monitored by the rule, the source and destination values for each individual sub-flow that violated the threshold are included in the RASLog entry. These values replace the learning parameters specified in the flow definition. The specific type of values (such as SID, DID, SFID, DFID, Rx, Tx and so on) are derived from the flow definition. In the following example, "(SID=039c00,DID=040700,Rx=10)" is the flow identifier for the learned flow "flows_to_did" (which was defined using **"*"** for the source and destination devices).

  ```
  2014/04/07-07:20:01, [MAPS-1003], 11131, SLOT 4 | FID 128, WARNING, SWAT_TUHIN_PLUTO, Flow
  (flows_to_did:SID=039c00,DID=040700,Rx=10), Condition= flows_to_did (TX_FCNT/hour>=10), Current
  Value:[TX_FCNT,698366979], RuleName=flow2, Dashboard Category=Traffic Performance.
  ```

- For *static* flows, the name of the flow is provided as part of the RASLog. In the following example, "max_thruput_flow" is the name of the problematic flow.

  ```
  2013/12/21-11:50:00, [MAPS-1003], 1225, FID 128, WARNING, sw0, Flow (max_thruput_flow),
  Condition=max_thruput_flow(TX_FCNT/min>=10), Current Value:[TX_FCNT,42654538],
  RuleName=thruputflow_thput_10, Dashboard Category=Traffic Performance.
  ```

Viewing a detailed switch status report

The detailed switch status displays historical data for port performance errors in addition to the summary view.

To view a detailed switch status report, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **mapsdb --show all** to display the detailed status.

The following example shows the detailed switch status. The status includes the summary switch status, plus port performance data for the current day (measured since midnight). If a monitoring rule is triggered, the corresponding RASLog message appears under the summary section of the dashboard. The column headings in the example have been edited slightly so as to allow the example to display clearly.

```
FID128:root> mapsdb --show all

1 Dashboard Information:
========================

DB start time:                  Thu Feb  4 19:17:13 2016
Active policy:                  dflt_aggressive_policy
Configured Notifications:       RASLOG,EMAIL,FENCE
Fenced Ports :                  5/60,5/62
Decommissioned Ports :          None
Fenced circuits :               None
Quarantined Ports :             None
Top PIDs <pid(it-flows)>:       0x69b0c0(8) 0x697b00(4)

2 Switch Health Report:
=======================

Current Switch Policy Status: HEALTHY


3.1 Summary Report:
===================

Category                 |Today                      |Last 7 days               |
-------------------------------------------------------------------------------
Port Health              |Out of operating range     |No Errors                 |
BE Port Health           |No Errors                  |No Errors                 |
GE Port Health           |In operating range         |No Errors                 |
Fru Health               |Out of operating range     |In operating range        |
Security Violations      |No Errors                  |No Errors                 |
Fabric State Changes     |Out of operating range     |No Errors                 |
Switch Resource          |In operating range         |In operating range        |
Traffic Performance      |In operating range         |In operating range        |
FCIP Health              |No Errors                  |No Errors                 |
Fabric Performance Impact|In operating range         |In operating range        |


3.2 Rules Affecting Health:
===========================

Category    |Repeat|Rule Name          |Execution Time   |Object          |Triggered   |
(Rule Count)|Count |                   |                 |                |Value(Units)|
-------------------------------------------------------------------------------------
Port        |1     |defNON_E_F_        |02/04/16 21:27:37|U-Port 5/60     |5           |
Health(8)   |      |PORTSSTATE_CHG_4   |                 |                |            |
            |      |                   |                 |                |            |
            |      |                   |                 |U-Port 5/62     |5           |
            |2     |defNON_E_F_        |02/04/16 21:28:19|U-Port 5/18     |4           |
            |      |PORTSSTATE_CHG_2   |                 |                |            |
            |      |                   |                 |U-Port 5/16     |4           |
            |      |                   |                 |U-Port 5/60     |4           |
            |      |                   |                 |U-Port 5/62     |4           |
            |1     |defALL_E_          |02/04/16 21:27:31|E-Port 5/18     |3           |
            |      |PORTSSTATE_CHG_2   |                 |                |            |
            |      |                   |                 |E-Port 5/16     |3           |
            |2     |defNON_E_F_        |02/04/16 21:28:37|U-Port 5/61     |1           |
            |      |PORTSLF_0          |                 |                |            |
            |      |                   |                 |U-Port 5/62     |2           |
            |      |                   |                 |U-Port 5/60     |2           |
            |1     |defNON_E_F_        |02/04/16 21:26:07|U-Port 1/23     |1 LOS       |
            |      |PORTSLOSS_SIGNAL_0 |                 |                |            |
            |      |                   |                 |U-Port 1/20     |1 LOS       |
Fru         |2     |defALL_PSPS_       |02/04/16 21:34:21|Power Supply 4  |ON          |
Health(4)   |      |STATE_ON           |                 |                |            |
            |      |                   |                 |Power Supply 3  |ON          |
```

```
            |2     |defALL_PSPS       |02/04/16 21:32:16|Power Supply 3 |FAULTY      |
            |      |STATE_FAULTY      |                 |               |            |
            |      |                  |                 |Power Supply 4 |FAULTY      |
Fabric State|5     |defSWITCHEPORT_   |02/04/16 21:29:02|Switch         |6 Ports     |
Changes(8)  |      |DOWN_1            |                 |               |            |
            |      |                  |                 |Switch         |2 Ports     |
            |      |                  |                 |Switch         |8 Ports     |
            |      |                  |                 |Switch         |4 Ports     |
            |      |                  |                 |Switch         |2 Ports     |
            |3     |defSWITCHEPORT_   |02/04/16 20:58:31|Switch         |2 Ports     |
            |      |DOWN_1            |                 |               |            |
            |      |                  |                 |Switch         |2 Ports     |
            |      |                  |                 |Switch         |2 Ports     |

4  History Data:
==================================

Stats(Units)        Current   01/28/16   01/21/16   --/--/--   --/--/--   --/--/--
-------------------------------------------------------------------------------
CRC(CRCs)           13(20)     -          -          -          -          -
ITW(ITWs)           -          13(612)    -          -          -          -
LOSS_SYNC(SyncLoss) -          -          -          -          -          -
LF                  -          -          -          -          -          -
LOSS_SIGNAL(LOS)    12(4)      12(4)      13(5)      -          -          -
                    -          13(4)      12(4)      -          -          -
                    -          14(4)      14(4)      -          -          -
PE(Errors)          -          -          -          -          -          -
STATE_CHG           12(5)      12(5)      12(9)      -          -          -
                    -          13(5)      13(9)      -          -          -
                    -          14(5)      14(9)      -          -          -
LR                  -          13(6)      12(10)     -          -          -
                    -          12(4)      13(10)     -          -          -
                    -          14(4)      14(10)     -          -          -
C3TXTO(Timeouts)    -          -          -          -          -          -
RX(%)               -          -          -          -          -          -
TX(%)               -          -          -          -          -          -
UTIL(%)             -          -          -          -          -          -
BN_SECS(Seconds)    -          -          -          -          -          -

5  History Data for back-end ports:
==================================

Stats(Units)        Current   01/28/16   01/21/16   --/--/--   --/--/--   --/--/--
-------------------------------------------------------------------------------
CRC(CRCs)           2/1/0(15)   -          -          -          -          -
LOSS_SYNC(SyncLoss) 2/1/0(1)  3/3/1(2)    3/3/1(2)      -          -          -
BAD_OS(Errors)      -           -          -          -          -          -
FRM_LONG(Errors)    -           -          -          -          -          -
FRM_TRUNC(Errors)   -           -          -          -          -          -

6 History Data for Gig Ethernet ports:
========================================

Stats(Units)        Current   01/28/16   01/21/16   --/--/--    --/--/--    --/--/--
-------------------------------------------------------------------------------
GE_CRC(CRCs)        -           -          -          -          -          -
GE_INV_LEN(Errors)  -           -          -          -          -          -
GE_LOS_OF_SIG(LOS)  10/ge2(66)  -          -          -          -          -
                    10/ge0(3)   -          -          -          -          -
                    10/ge10(3)  -          -          -          -          -
                    10/ge3(2)   -          -          -          -          -
                    10/ge4(2)   -          -          -          -          -
```

Viewing historical data

To view what has happened on a switch since the previous midnight, enter **mapsdb --show history** to view a summarized status history of the switch for this period, including both front-end ports and back-end ports (if present). History data for back-end ports is collected for a period of seven days and it is displayed in the "Backend port History Data" section.

NOTE

The output of the **mapsdb --show history** command differs depending on the platform on which you run it. On fixed-port switches, ports are shown in port index format; on chassis-based platforms, ports are shown in slot/port format. The values are expressed in kilos (k), Million (m), and Giga (g) units.

To view a summarized history of the switch status, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **mapsdb --show history**.

The following example displays all stored historical port performance data for 20 million CRCs.

```
switch:admin> mapsdb --show history

1 History Data:
================================
Stats(Units)        Current   01/28/16   01/21/16   --/--/--   --/--/--   --/--/--
-------------------------------------------------------------------------------
CRC(CRCs)           13(20 m)   -          -           -          -          -
ITW(ITWs)           -          13(612)    -           -          -          -
LOSS_SYNC(SyncLoss) -          -          -           -          -          -
LF                  -          -          -           -          -          -
LOSS_SIGNAL(LOS)    12(4m)     12(4)      13(5)       -          -          -
                    -          13(4)      12(4)       -          -          -
                    -          14(4)      14(4)       -          -          -
PE(Errors)          -          -          -           -          -          -
STATE_CHG           12(5m)     12(5)      12(9)       -          -          -
                    -          13(5)      13(9)       -          -          -
                    -          14(5)      14(9)       -          -          -
LR                  -          13(6)      12(10)      -          -          -
                    -          12(4)      13(10)      -          -          -
                    -          14(4)      14(10)      -          -          -
C3TXTO(Timeouts)    -          -          -           -          -          -
RX(%)               -          -          -           -          -          -
TX(%)               -          -          -           -          -          -
UTIL(%)             -          -          -           -          -          -
BN_SECS(Seconds)    -          -          -           -          -          -

2 History Data for back-end ports:
================================

Stats(Units)        Current   01/28/16    01/21/16   --/--/--   --/--/--   --/--/--
-------------------------------------------------------------------------------
CRC(CRCs)           2/1/0(15m)  -          -           -          -          -
LOSS_SYNC(SyncLoss) 2/1/0(1 m) 3/3/1(2 m) 3/3/1(2 m)  -          -          -
LR                  -           -          -          -          -          -
BAD_OS(Errors)      -           -          -          -          -          -
FRM_LONG(Errors)    -           -          -          -          -          -
FRM_TRUNC(Errors)   -           -          -          -          -          -

3 History Data for Gig Ethernet ports:
=======================================
Stats(Units)        Current    01/28/16    01/21/16  --/--/--  --/--/--    --/--/--
-------------------------------------------------------------------------------

GE_CRC(CRCs)        -           -           -          -         -          -
GE_INV_LEN(Errors)  -           -           -          -         -          -
GE_LOS_OF_SIG(LOS) 10/ge2(445) 10/ge2(129) 10/ge5(2)   -         -          -
                    -          10/ge0(3)   10/ge7(2)   -         -          -
                    -          10/ge10(3)  10/ge11(2)  -         -          -
                    -          10/ge3(2)   -           -         -          -
                    -          10/ge7(2)   -           -         -          -
                    -          10/ge11(5)  -           -         -          -
```

Viewing data for a specific time window

Detailed historical data provides the status of the switch for a specific time window. This is useful if, for example, users are reporting problems on a specific day or time. The same port-display patterns apply to viewing detailed historical data as for ordinary historical data.

To view detailed historical data about a switch, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Specify either the day or the hour of the current day you want to view:
   - To specify the day, enter **mapsdb --show details -day** *mm/dd/yyyy*.
   - To specify the hour, enter **mapsdb --show details -hr** *hh*.

   The following example displays historical port performance data for November 29, 2014 on a chassis-based platform. Because the health status of the current switch policy is CRITICAL, the sections "Contributing Factors" and "Rules Affecting Health" are displayed. If the current switch policy status was HEALTHY, neither of these sections would be displayed. The column headings in the example have been edited slightly so as to allow the example to display clearly.

```
switch:admin> mapsdb --show details -day 2/5/2016

1 Dashboard Information:
========================

DB start time:                  Thu Feb  4 19:17:13 2016
Active policy:                  dflt_aggressive_policy
Configured Notifications:       RASLOG,EMAIL,FENCE,SW_CRITICAL,SW_MARGINAL
Fenced Ports :                  5/60,5/62
Decommissioned Ports :          None
Fenced circuits :               None
Quarantined Ports :             None
Top PIDs <pid(it-flows)>:       0x69b0c0(8) 0x697b00(4)

2 Switch Health Report:
========================

Current Switch Policy Status: HEALTHY


3.1 Summary Report:
===================

Category                 |Today                      |Last 7 days               |
-------------------------------------------------------------------------------
Port Health              |In operating range         |In operating range        |
BE Port Health           |No Errors                  |No Errors                 |
GE Port Health           |In operating range         |In operating range        |
Fru Health               |In operating range         |In operating range        |
Security Violations      |Out of operating range     |No Errors                 |
Fabric State Changes     |In operating range         |In operating range        |
Switch Resource          |In operating range         |In operating range        |
Traffic Performance      |In operating range         |In operating range        |
FCIP Health              |No Errors                  |No Errors                 |
Fabric Performance Impact|Out of operating range     |In operating range        |


3.2 Rules Affecting Health:
===========================

Category     |Repeat|Rule Name           |Execution Time   |Object      |Triggered    |
(Rule Count) |Count |                    |                 |            |Value(units) |
-------------------------------------------------------------------------------------
Security     |1     |defSWITCHSEC_       |02/05/16 07:15:02|Switch      |1 Violations |
Violations(2)|      |TELNET_0            |                 |            |             |
             |1     |defSWITCHSEC_LV_0   |02/05/16 07:15:02|Switch      |1 Violations |
```

```
Fabric       |1     |defALL_PORTS_       |02/05/16 06:53:02|E-Port 12/27|IO_LATENCY_  |
Performance  |      |LATENCY_CLEAR       |                  |            |CLEAR        |
Impact (2)   |      |                    |                  |            |             |
             |1     |defALL_PORTS_IO_    |02/05/16 06:52:02|E-Port 12/27|IO_FRAME_LOSS|
             |      |FRAME_LOSS          |                  |            |             |

4 History Data:
===============

Stats(Units)        Current     02/04/16    --/--/--    --/--/--    --/--/--
--------------------------------------------------------------------------------

CRC(CRCs)           -           -           -           -           -
ITW(ITWs)           5/60(255)   -           -           -           -
LOSS_SYNC(SyncLoss) -           -           -           -           -
LF                  5/60(2)     4/17(4)     -           -           -
                    5/62(2)     4/18(4)     -           -           -
                    -           5/16(3)     -           -           -
                    -           5/18(3)     -           -           -
                    -           5/61(3)     -           -           -
                    -           5/19(2)     -           -           -
                    -           5/17(2)     -           -           -
                    -           5/63(2)     -           -           -
                    -           1/20(1)     -           -           -
                    -           1/21(1)     -           -           -
                    -           1/22(1)     -           -           -
                    -           1/23(1)     -           -           -
LOSS_SIGNAL(LOS)    -           1/20(2)     -           -           -
                    -           1/23(2)     -           -           -
                    -           1/22(1)     -           -           -
                    -           1/21(1)     -           -           -
PE(Errors)          -           -           -           -           -
STATE_CHG           5/60(5)     4/18(10)    -           -           -
                    5/62(5)     4/17(8)     -           -           -
                    4/59(2)     5/16(6)     -           -           -
                    -           5/18(6)     -           -           -
                    -           5/17(4)     -           -           -
                    -           5/19(4)     -           -           -
                    -           5/61(4)     -           -           -
                    -           5/63(4)     -           -           -
                    -           1/20(2)     -           -           -
                    -           1/21(2)     -           -           -
                    -           1/22(2)     -           -           -
                    -           1/23(2)     -           -           -
LR                  5/60(7)     4/18(17)    -           -           -
                    5/62(7)     4/17(16)    -           -           -
                    4/59(1)     5/16(10)    -           -           -
                    -           5/18(10)    -           -           -
                    -           5/61(10)    -           -           -
                    -           5/19(8)     -           -           -
                    -           5/17(8)     -           -           -
                    -           5/63(8)     -           -           -
                    -           1/20(4)     -           -           -
                    -           1/21(4)     -           -           -
                    -           1/22(4)     -           -           -
                    -           1/23(4)     -           -           -
C3TXTO(Timeouts)    12/27(1)    -           -           -           -
RX(%)               4/17(14.02) 1/19(2.59)  -           -           -
                    1/19(12.79) 4/17(2.12)  -           -           -
                    12/25(3.93) 4/18(1.35)  -           -           -
                    12/27(3.91) -           -           -           -
                    12/26(3.87) -           -           -           -
                    4/18(2.63)  -           -           -           -
                    8/35(1.68)  -           -           -           -
                    8/32(1.68)  -           -           -           -
                    8/34(1.68)  -           -           -           -
                    5/61(1.58)  -           -           -           -
                    5/63(1.58)  -           -           -           -
                    5/19(1.58)  -           -           -           -
                    5/18(1.58)  -           -           -           -
                    8/33(1.58)  -           -           -           -
                    5/17(1.58)  -           -           -           -
```

```
                    5/16(1.58)  -           -           -           -
                    5/1(1.57)   -           -           -           -
                    5/3(1.57)   -           -           -           -
                    5/2(1.56)   -           -           -           -
TX(%)               1/23(10.73) 1/19(2.15)  -           -           -
                    1/19(10.66) 1/23(1.54)  -           -           -
                    12/24(4.66) 4/17(1.02)  -           -           -
                    12/25(4.39) 4/18(1.01)  -           -           -
                    4/18(3.78)  -           -           -           -
                    4/17(3.78)  -           -           -           -
                    12/26(3.41) -           -           -           -
                    12/27(3.14) -           -           -           -
                    5/16(1.97)  -           -           -           -
                    5/17(1.97)  -           -           -           -
                    5/18(1.96)  -           -           -           -
                    5/2(1.96)   -           -           -           -
                    5/63(1.96)  -           -           -           -
                    8/33(1.80)  -           -           -           -
                    5/1(1.77)   -           -           -           -
                    8/34(1.30)  -           -           -           -
                    8/32(1.25)  -           -           -           -
                    5/3(1.21)   -           -           -           -
                    8/35(1.21)  -           -           -           -
                    5/61(1.20)  -           -           -           -
                    5/19(1.17)  -           -           -           -
UTIL(%)             1/19(11.73) 1/19(2.37)  -           -           -
                    4/17(8.90)  4/17(1.57)  -           -           -
                    1/23(5.36)  4/18(1.18)  -           -           -
                    12/24(4.27) -           -           -           -
                    12/25(4.16) -           -           -           -
                    12/26(3.64) -           -           -           -
                    12/27(3.52) -           -           -           -
                    4/18(3.21)  -           -           -           -
                    5/16(1.78)  -           -           -           -
                    5/17(1.77)  -           -           -           -
                    5/18(1.77)  -           -           -           -
                    5/63(1.77)  -           -           -           -
                    5/2(1.76)   -           -           -           -
                    8/33(1.69)  -           -           -           -
                    5/1(1.67)   -           -           -           -
                    8/34(1.49)  -           -           -           -
                    8/32(1.46)  -           -           -           -
                    8/35(1.45)  -           -           -           -
                    5/61(1.39)  -           -           -           -
                    5/3(1.39)   -           -           -           -
                    5/19(1.38)  -           -           -           -
BN_SECS(Seconds)    -           -           -           -           -

5 History Data for Backend ports:
=================================
Stats(Units)        Current     02/04/16    --/--/--    --/--/--    --/--/--
--------------------------------------------------------------------------------

CRC(CRCs)           -           -           -           -           -
ITW(ITWs)           -           -           -           -           -
LR                  -           -           -           -           -
BAD_OS(Errors)      -           -           -           -           -
FRM_LONG(Errors)    -           -           -           -           -
FRM_TRUNC(Errors)   -           -           -           -           -

6 History Data for Gig Ethernet ports:
======================================
Stats(Units)        Current     02/04/16    --/--/--    --/--/--    --/--/--
--------------------------------------------------------------------------------

GE_CRC(CRCs)        -           -           -           -           -
GE_INV_LEN(Errors)  -           -           -           -           -
GE_LOS_OF_SIG(LOS)  10/ge2(444) 10/ge2(129) -           -           -
                    -           10/ge0(3)   -           -           -
                    -           10/ge10(3)  -           -           -
                    -           10/ge3(2)   -           -           -
                    -           10/ge4(2)   -           -           -
```

```
                -           10/ge5(2)   -           -           -
                -           10/ge6(2)   -           -           -
                -           10/ge7(2)   -           -           -
                -           10/ge8(2)   -           -           -
                -           10/ge9(2)   -           -           -
                -           10/ge1(2)   -           -           -
                -           10/ge11(2)  -           -           -
```

The following example displays historical port performance data for four hours on a chassis-based platform. The History Data section has been trimmed so that the output will display correctly here. Normally there would be additional days of data.

```
switch:admin> mapsdb --show details -hr 4

1 Dashboard Information:
========================

DB start time:                  Thu Feb  4 19:17:13 2016
Active policy:                  dflt_aggressive_policy
Configured Notifications:       RASLOG,EMAIL,FENCE,SW_CRITICAL,SW_MARGINAL
Fenced Ports :                  5/60,5/62
Decommissioned Ports :          None
Fenced circuits :               None
Quarantined Ports :             None
Top PIDs <pid(it-flows)>:       0x69b0c0(8) 0x697b00(4)

2 Switch Health Report:
========================

Current Switch Policy Status: HEALTHY


3.1 Summary Report:
===================

Category                 |Today                      |Last 7 days               |
-------------------------------------------------------------------------------
Port Health              |In operating range         |In operating range        |
BE Port Health           |No Errors                  |No Errors                 |
GE Port Health           |In operating range         |In operating range        |
Fru Health               |In operating range         |In operating range        |
Security Violations      |In operating range         |No Errors                 |
Fabric State Changes     |In operating range         |In operating range        |
Switch Resource          |In operating range         |In operating range        |
Traffic Performance      |In operating range         |In operating range        |
FCIP Health              |No Errors                  |No Errors                 |
Fabric Performance Impact|In operating range         |In operating range        |


3.2 Rules Affecting Health:
===========================
Category    |Repeat|Rule Name            |Execution Time   |Object    |Triggered    |
(Rule Count)|Count |                     |                 |          |Value(Units) |
-------------------------------------------------------------------------------------
Port        |14    |defALL_E_PORTSITW_80 |02/04/15 19:17:13|E-Port 12 |12397261 ITWs|
Health(28)  |      |                     |                 |          |             |
            |      |                     |                 |E-Port 12 |12145947 ITWs|
            |      |                     |                 |E-Port 12 |12231844 ITWs|
            |      |                     |                 |E-Port 12 |12439476 ITWs|
            |      |                     |                 |E-Port 12 |12716699 ITWs|
            |14    |defALL_E_PORTSITW_41 |02/04/15 19:17:13|E-Port 12 |12397261 ITWs|
            |      |                     |                 |E-Port 12 |12145947 ITWs|
            |      |                     |                 |E-Port 12 |12231844 ITWs|
            |      |                     |                 |E-Port 12 |12439476 ITWs|
            |      |                     |                 |E-Port 12 |12716699 ITWs|


NOTE NOTE NOTE:   NEED TO ADD SECTIONS 4, 5, and 6 here.
```

Clearing MAPS dashboard data

To delete the stored data from the MAPS dashboard, enter **mapsdb --clear**. This command is useful if you want to see only the data logged after you have made a change to a switch (or a rule). The dashboard is also cleared if either a reboot or an HA failover happens.

To clear the stored dashboard data from a switch, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **mapsdb --clear** and specify the level of data (all, history, or summary) you want to remove from the display.

   When the dashboard is cleared, a RASLog message is generated. For more details on RASLog messages in MAPS, refer to the *Fabric OS Message Reference*.

   NOTE

   The **mapsdb --clear** command does not clear the current day's history data (that is, the first column of the history data). To clear the first column, enter **mapsdb --slotstatsclear**.

The following example clears only the dashboard summary data.

```
switch:admin> mapsdb --clear -summary
```

Port Monitoring Using MAPS

- Monitoring groups of ports using the same conditions....97
- Port monitoring using port names....97
- Port monitoring using device WWNs....98
- Adding a port to an existing static group....98
- Adding missing ports to a dynamic group....99
- Removing ports from a group....99
- D_Port monitoring....100
- Back-end port monitoring....102
- Port monitoring and pausing....103
- Gigabit Ethernet port monitoring....104

Monitoring groups of ports using the same conditions

You can create groups of ports that need to be modified using the same conditions. Then, you can use these groups to easily monitor the ports using a single set of rules and thresholds. MAPS refers to these as "logical groups."

Often on a switch there are sets of ports that behave in a similar manner and have a different behavior from other sets of ports. For example, the behavior of ports connected to UNIX hosts and servers is different from the behavior of ports connected to Windows hosts and servers. To easily monitor these similar sets of ports using the same rules, you can create a group and apply rules to the group.

To create a group and apply rules to the group, complete the following steps.

1. Create a logical group of similar ports.
2. Create rules using this logical group and add them to the active policy.
3. Enable the policy.

> NOTE
> You must enable the policy even if it is the active policy. Adding a rule to the active policy does not take effect until you re-enable the policy.

The following example creates the logical group "unix_ports" in the first line, creates a rule "unixHiCrc" using this logical group and adds them to the active policy "my_policy" in the second, and enables the policy in the third.

```
switch:FID6:admin> logicalgroup --create unix_ports -type port -members "1,3,17,21"
switch:FID6:admin> mapsrule --create unixHiCrc -monitor crc -group unix_ports -timebase min -op g -value 50
-action raslog -policy my_policy
switch:FID6:admin> mapspolicy --enable my_policy
```

Port monitoring using port names

Fabric OS allows you to monitor ports based on their assigned names.

Because the port name is an editable attribute of a port, you can name ports based on the device to which they are connected. You can then group the ports based on their port names. For example, if ports 1 to 10 are connected to devices from the ABC organization, you can name these ports ABC_port1, ABC_port2, and so on through ABC_port10. You can then define a group named "ABC_Ports" with a membership determined by having a port name that begins with "ABC_port". The following example defines a group based on this port name pattern. There is no limit on the number of ports that can be in a group.

```
switch246:FID128:admin> logicalgroup --create ABC_Ports -type port -feature portName -pattern ABC_port*
```

For more information on creating dynamic user-defined groups, refer to User-defined groups on page 44.

Port monitoring using device WWNs

Fabric OS allows you to monitor ports that are connected to a device that has device World Wide Name (WWN) that follows a certain pattern. This WWN pattern can then be used as part of the criteria for identifying a group. There is no limit on the number of ports that can be in a group.

One use of this might be for monitoring all ports on devices from a specific manufacturer. Because the WWN of a device contains information about the vendor, you can use this information to group devices based on this information, and then monitor them as a distinct group. For example, if you have a set of devices from vendor WXYZ with a WWN beginning 30:08:00:05, you can define a group named “WXYZ_Devs” with a membership determined by having a WWN that begins with “30:08:00:05”.

> NOTE
> The device node WWN information is fetched from the FDMI database, and group membership is validated against this database.

The following example defines a group based on this device WWN pattern.

```
switch1246:FID128:admin> logicalgroup --create WXYZ_Devs -type port -feature nodewwn -pattern 30:08:00:05*
```

For further information on creating dynamic user-defined groups, refer to User-defined groups on page 44.

Adding a port to an existing static group

If a new element, such as a host, target, or small form-factor pluggable (SFP) transceiver is added to the switch , you can monitor the ports in that element using existing rules for similar elements by adding it to an existing group, or creating a new group that uses an existing rule.

A port can be added to a static group or to dynamic groups, both user-defined and predefined.

For this type of monitoring, elements that are added manually to a group remain in the group whether they are online or offline.

To add a port to an existing group, complete the following steps. The added element is automatically monitored using the existing rules that have been set up for the group as long as the rules are in the active policy. You do not need to re-enable the active policy.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **logicalgroup --addmember** *group_name* **-member** *member_list*

   The element you want to add must be the same type as those already in the group (port, circuit, or SFP transceiver).

   You can specify either a single port, or specify multiple ports as either individual IDs separated by commas, or a range where the IDs are separated by a hyphen.
3. Optional: Enter **logicalgroup --show** *group_name* to see the members of the named group.

The following example adds the ports 31 and 41 to the critical_ports group.

```
switch:admin> logicalgroup --addmember critical_ports -members "31,41"
```

Listing this group produces the following output.

```
switch:admin> logicalgroup --show critical_ports
-------------------------------------------------------------------
Group Name     |Predefined |Type |Member Count |Members
-------------------------------------------------------------------
critical_ports |No         |Port |5            |10,15,25,31,41
```

Adding missing ports to a dynamic group

You can add ports to a predefined group (for example, ALL_HOST or ALL_TARGET) or user-defined dynamic group that might not have been included automatically.

For dynamic groups, you can specify any of the following:

- A single port
- Multiple ports separated by commas
- A range in which the IDs are separated by commas

You can create dynamic groups using either port names or WWNs, but you cannot use both in a single group definition. After a dynamic group is created, you can add ports to the same group using the same patterns as when the group was created. Quotation marks around the *member_list* value are optional. The operation is very similar to adding ports to a static group. However, the following items should be kept in mind for this monitoring:

- There is no validation of manual additions to a group; for example, if you add port 17 as part of an F_Port group, that port is added to the group even if it is not actually an F_Port.
- You can add a port to a predefined port group, but not to the group ALL_QUARANTINED_PORTS.

  NOTE
  The same restrictions as described in Adding a port to an existing static group on page 98 apply.

1. Enter **logicalgroup --show** *group_name*.
2. Enter **logicalgroup --addmember** *group_name* **-member** *member_list* to add the specified port to the named group.
3. Optional: Enter **logicalgroup --show** *group_name* to confirm the addition.

The following example shows these steps for the group ALL_HOST_PORTS, first showing that port 5 is not part of the group, then adding it to the group, then showing that it has been added to the group.

```
switch:admin> logicalgroup --show ALL_HOST_PORTS
------------------------------------------------------------------
Group Name     |Predefined |Type |Member Count |Members
------------------------------------------------------------------
ALL_HOST_PORTS |Yes        |Port |2            |0,15

switch:admin> logicalgroup --addmember ALL_HOST_PORTS -mem 5

switch:admin> logicalgroup --show ALL_HOST_PORTS
------------------------------------------------------------------
Group Name     |Predefined |Type |Member Count |Members
------------------------------------------------------------------
ALL_HOST_PORTS |Yes        |Port |3            |0,5,15
```

Removing ports from a group

In a similar way that you would add ports to either predefined or user-defined dynamic groups, you can also remove the ports from either group type. This is useful for devices that erroneously identify themselves as both host and target or for one-off exceptions when you want to remove a port that satisfies the specified pattern used for a user-defined dynamic group but you do no want it to be part of the group. .

To remove a port from a group, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.

2. Enter **logicalGroup --delmember** *group_name* **-members** *member_list*.

   You can specify either a single port, or specify multiple ports as either individual IDs separated by commas, or a range where the IDs are separated by a hyphen.

3. Optional: Enter **logicalGroup --show** *group_name* to confirm that the named ports are no longer part of the group.

The following example removes port 5 from the ALL_TARGET_PORTS group, and then shows that it is no longer a member of that group.

```
switch:admin> logicalgroup --delmember ALL_TARGET_PORTS -members "5"

switch:admin> logicalgroup --show ALL_TARGET_PORTS
--------------------------------------------------------------------------------
Group Name       |Predefined |Type |Member Count |Members
--------------------------------------------------------------------------------
ALL_TARGET_PORTS |Yes        |Port |5            |1,11,22,32,44
```

D_Port monitoring

In Fabric OS 7.3.0 and later, D_Ports can be monitored by MAPS using the group ALL_D_PORTS.

You can either configure a port as a D_Port using the CLI, or Fabric OS can dynamically convert a port to a D_Port. When a port is configured as a D_Port, MAPS automatically adds the port to the ALL_D_PORTS group, and starts monitoring the port.

> NOTE
>
> Refer to the "Diagnostic Port" chapter of the *Fabric OS Administrator's Guide* for more details about D_Port monitoring.

To simplify default monitoring, rules based on the ALL_D_PORTS group are already part of the default policies. To allow for short-running and long-running D_Port tests, the default policies in MAPS use D_Port rules that span multiple error thresholds spanning multiple timebases. If any of the rules are triggered, MAPS triggers the action configured for the rule, alerts the fabric service module and caches the data in the dashboard.

D_Port monitoring monitors all D_Port errors; however, the fabric service module is notified only for the following errors:

- CRC
- ITW — for the enc_out and enc_in invalid transmission words only
- LF
- LOSS_SYNC

> NOTE
> The MAPS DPORT_ITW rule for enc_out and enc_in is not applicable for the D_Port test.

The D_Port monitoring feature is only supported for 10 Gbps, 16 Gbps, and 32 Gbps SFPs/QSPFs and 8 Gbps LWL and ELWL ports on the following blades: CR16-4, CR16-8, FC8-32E, FC8-48E, FC16-32, FC16-48, and FC16-64.

> NOTE
> In versions of Fabric OS prior to 7.3, MAPS monitored D_Ports using the NON_E_F_PORTS group, but the default rules for this group did not provide the flexibility now available through the ALL_D_PORTS group.

The **mapsrule** command accepts the ALL_D_PORTS group, which can be used as shown in the following example.

```
mapsrule --create d_port_mon -group ALL_D_PORTS -monitor CRC -timebase min -op ge -value 1 -action raslog -
policy nil
```

Using the **mapsdb --show** command shows any error or rule violation during diagnostics tests on a D_Port.

```
switch:admin> mapsdb --show
1 Dashboard Information:
========================
DB start time:                  Wed Mar 26 10:02:38 2014
Active policy:                  dflt_moderate_policy
Configured Notifications:       SW_CRITICAL,SW_MARGINAL
Fenced Ports :                  None
Decommissioned Ports :          None
Quarantined  Ports :            None

2 Switch Health Report:
=======================
Current Switch Policy Status: MARGINAL
Contributing Factors:
---------------------
*BAD_PWR (MARGINAL).

3.1 Summary Report:
===================
Category                 |Today                      |Last 7 days        |
-------------------------------------------------------------------------
Port Health              |Out of operating range     |In operating range |
BE Port Health           |No Errors                  |No Errors          |
Fru Health               |In operating range         |In operating range |
Security Violations      |No Errors                  |No Errors          |
Fabric State Changes     |Out of operating range     |In operating range |
Switch Resource          |In operating range         |In operating range |
Traffic Performance      |In operating range         |In operating range |
FCIP Health              |Not applicable             |Not applicable     |
Fabric Performance Impact|In operating range         |In operating range |

3.2 Rules Affecting Health:
===========================
Category(Rule Count)|RptCount|Rule Name           |Execution Time   |Object |Triggered Value(Units)|
-----------------------------------------------------------------------------------------------------
Port Health(5)      |1       |defALL_D_PORTSCRC_1 |05/07/14 08:43:32|D_Port 20|300 Errors          |
                    |4       |defNON_E_F_PORTSLF_0|05/07/14 08:42:56|D_Port 7 |6                   |
                    |        |                    |                 |D_Port 7 |6                   |
                    |        |                    |                 |D_Port 7 |6                   |
                    |        |                    |                 |D_Port 7 |7                   |
```

You can also run the **portdporttest --show** *port_number* command to see details of an individual port. The following example shows the results for port 28.

```
switch:admin> portdporttest --show 28
D-Port Information:
===================
Port: 28
Remote WWNN: 10:00:00:05:1e:e5:e4:00
Remote port: 164
Mode: Manual
Start time: Thu Nov 7 13:43:26 2013
End time: Thu Nov 7 13:53:43 2013
Status: PASSED*
```

Refer to the *Fabric OS Command Reference* for additional information on these commands.

When running the **portdporttest --show** *port_number* command to see details for a 32 Gbps QSPF, the output appears similar to the results for a 16 Gbps QSPF, except the Electric loopback and Optical loopback are skipped.

```
switch:admin> portdporttest --show 48
D-Port Information:
===================
Port:     48
Remote WWNN:    10:00:00:27:f8:f0:26:41
Remote port index:   52
Mode:     Manual
No. of test frames:   1 Million
Test frame size:   1024 Bytes
FEC (enabled/option/active): Yes/No/Yes
CR (enabled/option/active): No/No/No
Start time:    Wed Jun 24 08:57:43 2015
End time:    Wed Jun 24 08:57:50 2015
Status:    PASSED
==============================================================================
Test       Start time Result  EST(HH:MM:SS) Comments
==============================================================================
Electrical loopback -------- SKIPPED     -------- No SFP or chip support
Optical loopback    -------- SKIPPED     -------- No SFP or chip support
Link traffic test   08:57:44 PASSED      -------- ----------
==============================================================================
Roundtrip link latency:  275 nano-seconds
Approximate cable distance: unknown
Buffers required:   1 (for 2112 byte frames at 32Gbps speed)
Egress pwr:  Tx: 0.3  dBm, Rx: -3.6 dBm, Diff: 3.9 dBm(Loss is within tolerable limit)
Ingress pwr: Rx: -4.1 dBm, Tx: 0.3  dBm, Diff: 4.4 dBm(Loss is within tolerable limit)
```

Back-end port monitoring

A back-end port connects a core switching blade to a port or application blade (and vice versa). The primary task of back-end ports is to route packets passing through a switch's ASICs. Switch (and consequently fabric) performance degrades when there are errors in back-end ports. MAPS error notification allows you to take corrective action earlier. In terms of monitoring system functionality, back-end ports can be connected to ports within a fixed-port switch or to other blades within a Backbone chassis, and so their functionality is different from front-end ports, which connect to devices outside of the switch.

In Fabric OS 7.4.0 and later versions, MAPS monitors port counter errors on back-end ports in chassis-based switches, as well as the Brocade 6250 switch . When back-end port rules are triggered,, you can then do SerDes (Serializer/Deserializer) tuning on those ports where errors exceed the desired threshold, which will greatly improve the packet-forwarding performance of these ports.

When a switch is initialized, all back-end ports are automatically brought online and stay online until the slot is powered off or the blade is removed (for chassis-based switches), or if the switch itself goes down (for fixed-port switches). This allows MAPS to continuously monitor the platform for back-end port errors. The monitoring systems are for a given switch or chassis, so all the monitoring systems are monitored only in the default switch.

MAPS monitors the port counter statistics for back-end ports through the group ALL_BE_PORTS, which identifies each port using a combination, such as 3/31. In case of fixed-port switches, the slot number is 0. History data for back-end ports is collected for a period of seven days and is displayed in the "Back-end Port History Data" section of the MAPS dashboard.

MAPS monitors the back-end port errors, and keeps track of the connected port for every back-end port. When a RASLog is generated, the connected port information is added, as shown in the RASLog output. In RASLog, 1/14 is the port where the errors are seen, and the port 5/182 is the connected port. When errors are seen on any back-end port. This additional port information helps in debugging and fixing the issue.

The following example is typical of a RASLog message generated for a back-end port. In this example, the rule sets the threshold at more than 35 CRC errors in a minute (CRC/min>35).

```
2015/06/29-21:40:02, [MAPS-1003], 48, SLOT 6 FID 128, WARNING, dcx_178, BE Port 1/14,
Condition=ALL_BE_PORTS(CRC/5MIN>10), Current Value:[CRC,125 CRCs (Conn. port 5/182)],
RuleName=defALL_BE_PORTSCRC_5M_10, Dashboard Category=BE Port Health.
```

For more information on back-end health monitoring, refer to Back-end Health on page 33 and Back-end port monitoring thresholds on page 148.

Dashboard output of back-end port rule violations

When a back-end port monitoring rule is triggered, the corresponding RASLog rule information appears in the "Rules Affecting Health" section of the dashboard under "BE Port Health".

The following example displays an excerpt from the MAPS dashboard; the items for the back-end port reporting are listed on the line starting with "BE Port Health (1)" and in the section labeled "4.2 Backend port History Data." Be aware that the column headings in the example have been edited slightly to allow the example to display clearly.

```
      (output truncated)
Rules Affecting Health:
================================
Category(Rule Cnt)|Rpt Cnt|Rule Name                | Execution Time    |Object     |Triggered Value
-------------------------------------------------------------------------------------------------------
BE Port Health(1) |1      | defALL_BE_PORTSCRC_5M_10 | 01/21/16 01:30:60 |Port 6/8   | 50 CRCs

4.1 Front end port History Data:
================================
Stats(Units)        Current   01/21/16  01/14/16    --/--/--    --/--/--   --/--/--   --/--/--
                    Port(val) Port(val) Port(val)
-------------------------------------------------------------------------------------------
CRC(CRCs)           1/13(20)   -          -          -          -          -          -
ITW(ITWs)           -          1/13(612)  -          -          -          -          -
LOSS_SYNC(SyncLoss) -          -          -          -          -          -          -
LF                  -          -          -          -          -          -          -
LOSS_SIGNAL(LOS)    -          -          -          -          -          -          -
PE(Errors)          -          -          -          -          -          -          -
STATE_CHG           -          -          -          -          -          -          -
C3TXTO(Timeouts)    -          -          -          -          -          -          -
RX(%)               -          -          -          -          -          -          -
TX(%)               -          -          -          -          -          -          -
UTIL(%)             -          -          -          -          -          -          -
BN_SECS(Seconds)    -          -          -          -          -          -          -

4.2 Backend port History Data:
================================
Stats(Units)        Current   01/21/16   01/14/16   --/--/--   --/--/--   --/--/--   --/--/--
                    Port(val) Port(val)  Port(val)
-------------------------------------------------------------------------------------------
CRC(CRCs)           6/8(50)     -          -          -          -          -          -
```

Port monitoring and pausing

Pausing operations on a port does not affect flow monitoring. Flow monitoring is done at the flow level and the details of the flow passing through a particular port is transparent to MAPS.

Gigabit Ethernet port monitoring

NOTE

Gigabit Ethernet port monitoring can be performed on the following devices:

- 7840 switch
- SX6
- FX8-24 blades

Fabric OS allows you to monitor GE ports in a switch and receive counter errors reported by ASIC drivers as RASLog, SNMP, and e-mail alerts. This reporting helps you identify the nature of FCIP and IP Extension traffic errors at the Level 2 (L2) link layer.

The Ethernet MAC counters are maintained on a 1/10/40 GigE port basis.

MAPS monitors the following errors counters in GE ports using the ALL_EXT_GE_PORTS group:

1. CRC—Frames received with CRC error
2. Carrier—Frames aborted because of carrier sense error, no carrier or loss of carrier
3. Length—Frames received with length error, length type field does not match frame size

The CRC and length error counters track receive errors, and the carrier error counter tracks transmission errors caused by signal loss.

Currently, the rules are created in the default policy for the minute time base, and there are two thresholds. The supported MAPS actions are RASLOG, SNMP, and EMAIL. MAPS send the alerts with a RASLog message and then takes any other actions configured in a rule. The indexing scheme for addressing the members in this group is similar to the addressing scheme of the front-end ports.

The following RASLog message is generated when a GE port rule is triggered due to CRC errors.

```
2016/03/03-02:45:36, [MAPS-1003], 1820, SLOT 7 | FID 1, WARNING, CHASSIS214, (5/ge2),
Condition=ALL_EXT_GE_PORTS(GE_CRC/min>=50), Current Value:[GE_CRC,65 CRCs],
RuleName=test_port_rule_1, Dashboard Category=GE Port category.
```

The following example shows the result of using the **--show -all** option.

```
switch:admin> mapspolicy --show -all

Rule name                        |Condition                                    |Actions          |
---------------------------------------------------------------------------------------------------
defALL_EXT_GE_PORTSCRC_20        |ALL_EXT_GE_PORTS(GE_CRC/MIN > 20)       |RASLOG,SNMP,EMAIL|
defALL_EXT_GE_PORTSLOS_20        |ALL_EXT_GE_PORTS(GE_LOS_OF_SIG/MIN > 20)|                 |
defALL_EXT_GE_PORTSINV_LEN_20    |ALL_EXT_GE_PORTS(GE_INV_LEN/MIN > 20)   |                 |
dflt_conservative_policy:        |                                             |                 |
---------------------------------------------------------------------------------------------------|
defALL_EXT_GE_PORTSCRC_10        |ALL_EXT_GE_PORTS(GE_CRC/MIN > 10)       |RASLOG,SNMP,EMAIL|
defALL_EXT_GE_PORTSLOSS_OF_SIG_10|ALL_EXT_GE_PORTS(GE_LOS_OF_SIG/MIN > 10)|                 |
defALL_EXT_GE_PORTSFRM_LEN_10    |ALL_EXT_GE_PORTS(GE_INV_LEN/MIN > 10)   |                 |
dflt_moderate_policy:                                                                            |
---------------------------------------------------------------------------------------------------|
defALL_EXT_GE_PORTSCRC_5         |ALL_EXT_GE_PORTS(GE_CRC/MIN > 5)        |RASLOG,SNMP,EMAIL|
defALL_EXT_GE_PORTSLOSS_OF_SIG_5 |ALL_EXT_GE_PORTS(GE_LOS_OF_SIG/MIN > 5) |                 |
defALL_EXT_GE_PORTSFRM_LEN_5     |ALL_EXT_GE_PORTS(GE_INV_LEN/MIN > 5)    |                 |
dflt_aggressive_policy           |                                             |                 |
---------------------------------------------------------------------------------------------------
```

GE port monitoring CRC rule creation

Rule creation for Gigabit Ethernet port monitoring is similar to the procedure for creating rules for other ports. You can enable default policies to monitor the counters or create custom rules and policies to monitor the GE ports.

The following example creates a CRC rule for GE port.

```
 switch:admin>mapsrule -create
test_backend_port_rule_1 -group ALL_EXT_GE_PORTS -monitor CRC
-timebase min -op ge -value 35 -action raslog
```

When the CRC error for any GE port is greater than or equal to 35 during one minute, the rule is triggered and a RASLOG message is generated. GE port error counters are important in debugging problems. Rules are created to monitor the port errors when there is FCIP traffic between two switches connected across an IP WAN network.

Monitoring Flow Vision Flows with MAPS

- Monitoring Flow Vision Flow Monitor data with MAPS....107
- Monitoring traffic performance....109
- Monitoring learned flows....110
- I/O latency monitoring....110

Monitoring Flow Vision Flow Monitor data with MAPS

The Monitoring and Alerting Policy Suite (MAPS) can monitor flows created using Flow Vision, which lets you create and track flows. Flow Monitor, when enabled on a flow, can capture various metrics for the flow, such as the number of frames received or transmitted and the number of bytes received or transmitted. MAPS can monitor the metrics captured by Flow Monitor on various flows, evaluate various conditions, and generate RASLogs, SNMP alerts, or e-mails. In order for MAPS to be able to monitor a flow, the Flow Monitor feature of Flow Vision must be enabled on the flow. This capability provides the flexibility to monitor each flow with its specific thresholds.

NOTE

Flows on which the Flow Vision Flow Generator or Flow Mirror features are enabled cannot be monitored using MAPS.

For details on flows and Flow Vision, refer to the Flow Monitor section of the *Flow Vision Administrator's Guide*.

To monitor flows using MAPS, you must perform two processes: (1) importing the flows and (2) adding monitoring flows after importing. Perform these processes as follows:

1. Create the flow in Flow Vision using the **flow --create** command.
2. Import the flow into MAPS using the **mapsconfig --import** command.
3. Enter **logicalgroup --show** to confirm that the flow was correctly imported into MAPS. The imported flow name indicates the groups that can be monitored.
4. Define a MAPS rule using the **mapsrule --create** command (for the supported timebases).

   Refer to MAPS rules overview on page 52 for information on creating and using rules.
5. Enter **mapspolicy --enablepolicy** *policy_name* to activate the policy.

The following example illustrates the flow-monitoring steps. The first command line creates the flow (called myflow_22 for this example), the second command line imports it, and the third command line displays the members of the logical groups. The fourth command line creates a rule for the group and the fifth command line enables the flow with the new rule active.

```
          switch246:FID128:admin> flow --create myflow_22 -feature monitor -egrport 21 -srcdev 0x010200 -
dstdev 0x011500

switch246:FID128:admin> mapsconfig --import myflow_22

switch246:FID128:admin> logicalgroup --show
-------------------------------------------------------------------
Group Name             |Predefined|Type |Member Count|Members
-------------------------------------------------------------------
ALL_PORTS              |Yes       |Port |8           |3/4,3/6-15
NON_E_F_PORTS          |Yes       |Port |2           |3/4,3/6
ALL_E_PORTS            |Yes       |Port |0           |
ALL_F_PORTS            |Yes       |Port |5           |8/0-1,8/4
ALL_OTHER_F_PORTS      |Yes       |Port |1           |8/0
ALL_HOST_PORTS         |Yes       |Port |1           |8/8
ALL_TARGET_PORTS       |Yes       |Port |0           |
ALL_QUARANTINED_PORTS  |Yes       |Port |2           |8/0,8/4
ALL_2K_QSFP            |Yes       |Sfp  |4           |8/28-31
ALL_100M_16GSWL_QSFP   |Yes       |Sfp  |0           |
myflow_22              |No        |Port |3           |Monitored Flow

switch246:FID128:admin> mapsrule --create myRule_22 -group myflow_22
                          -monitor TX_FCNT -timebase hour -op g -value 22
                          -action RASLOG -policy myPolicy_22

switch246:FID128:admin> mapspolicy --enable policy myPolicy_22
```

Importing flows

1. Create the flow in Flow Vision using the **flow --create** command.
2. Import the flow into MAPS using the **mapsconfig --import** command.
3. Enter **logicalgroup --show** to confirm that the flow was correctly imported into MAPS. The imported flow name indicates the groups that can be monitored.
4. Define a MAPS rule using the **mapsrule --create** command (for the supported timebases).

   Refer to MAPS rules overview on page 52 for information on creating and using rules.
5. Enter **mapspolicy --enablepolicy** *policy_name* to activate the policy.

The following example illustrates the steps to import flows.. The first command line creates and activates the flow; the second command line imports it.

```
switch246:FID128:admin> flow --create myflow_22 -feature monitor -egrport 21 -srcdev 0x010200 -dstdev
0x011500

switch246:FID128:admin> mapsconfig --import myflow_22
```

Adding monitoring flows after importing

To add monitoring flows after they have been imported, you must verify that the flows were imported, define a MAPS rule and add it to a MAPS policy, and then enable the MAPS policy. Perform the following steps:

1. Enter **logicalgroup --show** to confirm that the flow was correctly imported into MAPS.
2. Define a MAPS rule using the **mapsrule --create** command (for the supported timebases) and add it to a policy.

   Refer to MAPS rules overview on page 52 for information on creating and using rules.
3. Enter **mapspolicy --enablepolicy** *policy_name* to activate the policy.

The following example illustrates the steps to add monitoring flows after importing. The first command line displays the members of the logical groups, including the imported flow, "myflow22", for this example. The second command line creates a rule for the group, and the third command line enables the flow with the new rule active.

```
switch246:FID128:admin> logicalgroup --show
-----------------------------------------------------------------
Group Name             |Predefined|Type  |Member Count|Members
-----------------------------------------------------------------
ALL_PORTS              |Yes       |Port  |8           |3/4,3/6-15
NON_E_F_PORTS          |Yes       |Port  |2           |3/4,3/6
ALL_E_PORTS            |Yes       |Port  |0           |
ALL_F_PORTS            |Yes       |Port  |5           |8/0-1,8/4
ALL_OTHER_F_PORTS      |Yes       |Port  |1           |8/0
ALL_HOST_PORTS         |Yes       |Port  |1           |8/8
ALL_TARGET_PORTS       |Yes       |Port  |0           |
ALL_QUARANTINED_PORTS  |Yes       |Port  |2           |8/0,8/4
ALL_2K_QSFP            |Yes       |Sfp   |4           |8/28-31
ALL_100M_16GSWL_QSFP   |Yes       |Sfp   |0           |
myflow_22              |No        |Port  |3           |Monitored Flow

switch246:FID128:admin> mapsrule --create myRule_22 -group myflow22 -monitor TX_FCNT -timebase hour -op g -
value 22 -action RASLOG -policy myPolicy

switch246:FID128:admin> mapspolicy --enable policy myPolicy22
```

Monitoring traffic performance

The following examples illustrate how to use MAPS to monitor traffic performance.

Monitoring end-to-end performance

In the following example, MAPS is configured to monitor the throughput of a flow between two specific devices through port 5. To achieve this, you define a flow using the **-feature monitor** for a particular Source ID, Destination ID, and port using the Flow Vision **flow** command. Then, you import the flow into MAPS and create rules to monitor the throughput for the flow.

```
switch246:admin> flow --create E2E_flow -feature monitor -ingrport 5 -scrdev 0x010200 -dstdev 0x020300

switch246:admin> mapsconfig --import E2E_flow

switch246:admin> mapsrule --create E2E_rule -monitor TX_THPUT -group E2E_flow -timebase min -op g -value 10
-action rasLog -policy flowpolicy

switch246:admin> mapspolicy --enable flowpolicy
```

NOTE

The group name needs to match the imported flow name. In this case “E2E_flow”.

Monitoring frames for a specified set of criteria

In the following example, MAPS uses the flow “abtsflow” to watch for frames in a flow going through port 128 that contain SCSI ABORT sequence markers.

```
switch246:admin> flow --create abtsflow -feature mon -ingrport 128 -frametype abts

switch246:admin> mapsconfig --import abtsflow
```

You can then define rules for this flow (group), and then re-enable the policy so they take effect. THe following example creates only one rule, “abts_rule”.

```
switch246:admin> mapsrule --create abts_rule -monitor txfcnt -group abtsflow -timebase min -op ge -value 10
-action raslog -policy flowpolicy
switch246:admin> mapspolicy --enable flowpolicy
```

NOTE

Any new rule you create will not take effect until you enable the policy associated with it.

Monitoring learned flows

Flow Vision allows you use a wild character when creating flows, so that you do not have to specify a SID or DID. In such cases, Flow Vision “learns” all the flows that match the input criteria, and if Flow Monitor feature is enabled on these flows, metrics are captured for each of the learned flows. This ability allows you to capture information about multiple flows or gather information when you do not know a specific flow that is of interest.

If a learned flow is imported into MAPS and rules created for the flow, MAPS will evaluate the configured conditions for each of the learned flows.

The following examples illustrate some of the ways you might want to monitor learned flows.

Excessive throughput notification

To be notified of all the Source ID-Destination ID device pairs for which the RX throughput is greater than a threshold, you would import a learning flow with both the Source ID and Destination ID specified as "*" and define a rule to provide the notification, as shown in the following example.

```
switch246:FID128:admin> flow --create thruputflow -feature monitor -ingrp 123 -srcdev "*" -dstdev "*"

switch246:FID128:admin> mapsconfig --import thruputflow

switch246:FID128:admin> mapsrule --create thruputflow_thput_10 -group thruputflow -timebase hour -m
RX_THRUPUT -op ge -v 10 -a RASLOG,EMAIL
```

I/O latency monitoring

For Gen 6 platforms, MAPS monitors device-level I/O performance by monitoring I/O latency statistics for all the flows provided by the IO Insight capability. You can create new flows and import them to MAPS for monitoring.

MAPS monitors the following matrices:

- **Pending I/O:** Indicates how many I/O requests are pending
- **Completion time:** Indicates total request completion time
- **First Read or First Write time:** Indicates how quickly the target responds to the command

The Gen 6 IO Insight metrics are aonly available with a Flow Vision flow, and they can only be monitored by MAPS when the flow is imported into MAPS.

TABLE 32 I/O latency matrix

Matrix	Size	Monitor
Pending I/O	Less than 8K	RD_PENDING_IO_LT_8K
		WR_PENDING_IO_LT_8K
	8K but less than 64K	RD_PENDING_IO_8_64K
		WR_PENDING_IO_8_64K
	64K but less than 512K	RD_PENDING_IO_64_512K
		WR_PENDING_IO_64_512K
	Greater than or equal to 512K	RD_PENDING_IO_GE_512K
		WR_PENDING_IO_GE_512K
Completion time	Less than 8K	RD_STATUS_TIME_LT_8K
		WR_STATUS_TIME_LT_8K
	8K but less 64K	RD_STATUS_TIME_8_64K
		WR_STATUS_TIME_8_64K
	64K but less than 512K	RD_STATUS_TIME_64_512K
		WR_STATUS_TIME_64_512K
	Greater than or equal to 512K	RD_STATUS_TIME_GE_512K
		WR_STATUS_TIME_GE_512K
First Read or First Write time	Less than 8K	RD_1stDATA_TIME_LT_8K
		WR_1stDATA_TIME_LT_8K
	8K but less than 64K	RD_1stDATA_TIME_8_64K
		WR_1stDATA_TIME_8_64K
	64K but less than 512K	RD_1stDATA_TIME_64K_512K
		WR_1stDATA_TIME_64K_512K
	Greater than or equal to 512K	RD_1stDATA_TIME_GE_512K
		RD_1stDATA_TIME_GE_512K

Monitoring is done for either a minute, an hour, or a day.

- **Minute monitoring**: the difference between two I/O samples is a minute. For example, X1 and X2 are I/O numbers at the beginning and end of a minute. MAPS monitors for (X2-X1)/60 units to send an alert.
- **Hour monitoring**: the difference between two I/O samples is an hour. For example, X1 and X2 are I/O numbers at the beginning and end of an hour. MAPS monitors for (X2-X1)/ (60*60) units to send an alert.
- **Day monitoring**: the difference between two I/O samples is a period of a day. For example, X1 and X2 are I/O numbers at the beginning and end of the day. MAPS monitors (X2-X1)/ (24*60*60) units to send an alert.

Monitoring I/O latency

MAPS monitors the I/O latency statistics for all the flows to monitor the traffic performance. MAPS does not provide any default rules to monitor IO Insight statistics. You need to create the rules for monitoring.

Perform the following steps to monitor I/O latency:

1. Create the flow.

```
flow --create ios_host_flow -srcdev 041900 -dstdev 041b00 -ingrport 25 -fea mo
```

2. Import the flow using the **mapsconfig** command.

```
mapsconfig --import ios_host_flow
```

After importing the flow, MAPS automatically monitors the flow, if the rule is already been created.

3. Create the rule.

```
mapsrule –create ios_mon_rule –group ios_host_flow –monitor RD_1stDATA_TIME_8_64K  –timebase min –op
gt –value 123 –action raslog
```

Fabric performance impact monitoring using MAPS

- MAPS latency monitoring....113
- MAPS and Bottleneck Detection....119
- Slow Drain Device quarantining....121

MAPS allows you to monitor fabrics for performance impacts, including timeouts, latency, and throughput.

There are many distinct elements and layers in a fabric (applications, servers, switches, targets, LUNs, and so on) and consequently multiple places that could possibly be the cause of fabric performance impacts (bottlenecks). Because each application's behavior is unique, the impact of a bottleneck on one individual application might be different from its impact on another application. Each MAPS event needs to be viewed in conjunction with other server or application events to determine the actual root cause of the problem.

The Brocade blades, chassis, and fixed-port switches are also continuously monitored for thermal safety. For more information, refer to "System temperature monitoring" in the *Fabric OS Administrator's Guide*.

NOTE

In Fabric OS 8.0.1, fabric performance impact monitoring no longer requires a Fabric Vision license; it is enabled by default.

MAPS latency monitoring

MAPS latency detection is based on data retrieved from the port on the switch (just one element in the fabric), which is used to determine the potential impact to other flows using the fabric. MAPS monitors the fabric impact state of individual F_Ports (but not F_Port trunks) on both individual switches and switches operating in Access Gateway mode. On an Access Gateway set up with F_Port trunks, fabric performance is monitored only on those F_Ports actually present on the Access Gateway.

MAPS monitors the current latency on F_Ports over different time windows to determine the impact of latency on the fabric. If it determines the latencies on these ports are severe enough to significantly impact fabric performance, the state of that port is changed to IO_PERF_IMPACT or IO_FRAME_LOSS, depending on the severity , and the state change is reported to the MAPS dashboard. When the latencies drop to normal levels, the port state is changed to IO_LATENCY_CLEAR. The IO_PERF_IMPACT value is calculated using buffer credit zero or transient queue latency counters, while IO_FRAME_LOSS is calculated using transient queue latency only.

The following example shows first the MAPS dashboard, displaying the IO_PERF_IMPACT report and then the IO_LATENCY_CLEAR report. The dashboard has been edited to show only Section 3. The back-slash character (\) in the following examples indicates a break inserted because the output is too long to display here as a single line.

```
switch:admin> mapsdb --show
      (output truncated)
3.1 Summary Report:
===================
Category                 |Today                     |Last 7 days               |
-------------------------------------------------------------------------------
Port Health              |Out of operating range    |No Errors                 |
BE Port Health           |No Errors                 |No Errors                 |
Fru Health               |In operating range        |In operating range        |
Security Violations      |No Errors                 |No Errors                 |
Fabric State Changes     |In operating range        |No Errors                 |
Switch Resource          |In operating range        |In operating range        |
Traffic Performance      |In operating range        |In operating range        |
FCIP Health              |In operating range        |No Errors                 |
Fabric Performance Impact|Out of operating range    |In operating range        |

3.2 Rules Affecting Health:
===========================
Category(Rule Count)        |Repeat Count|Rule Name                     |\
-------------------------------------------------------------------------\
Fabric Performance impact(1)|1           |defALL_F_PORTS_IO_PERF_IMPACT|\

\|Execution Time   |Object     |Triggered Value(Units)|
\-----------------------------------------------------
\|03/19/16 22:48:01|F_Port 8/8 |IO_PERF_IMPACT        |
(output truncated)
```

After the latency is cleared on the F_Port 8/8, the MAPS dashboard report changes to the following.

```
switch:admin> mapsdb --show
     (output truncated)
3.1 Summary Report:
===================
Category                 |Today                     |Last 7 days               |
-------------------------------------------------------------------------------
Port Health              |Out of operating range    |No Errors                 |
BE Port Health           |No Errors                 |No Errors                 |
Fru Health               |In operating range        |In operating range        |
Security Violations      |No Errors                 |No Errors                 |
Fabric State Changes     |In operating range        |No Errors                 |
Switch Resource          |In operating range        |In operating range        |
Traffic Performance      |In operating range        |In operating range        |
FCIP Health              |In operating range        |No Errors                 |
Fabric Performance Impact|In operating range           |In operating range        |

3.2 Rules Affecting Health:
===========================
Category(Rule Count)        |Repeat Count|Rule Name                     |\
-------------------------------------------------------------------------\
Fabric Performance impact(1)|1           |defALL_F_PORTS_IO_LATENCY_CLEAR|\

\|Execution Time   |Object     |Triggered Value(Units)|
\-----------------------------------------------------
\|08/19/14 23:48:01|F_Port 8/8 |IO_LATENCY_CLEAR      |
(output truncated)
```

Frame timeout latency monitoring

MAPS monitors for Class 3 frame timeout errors (C3TXTO) on individual ports and when a timeout is detected on a port, MAPS reports them by setting the port state to IO_FRAME_LOSS and posting a RASLog message containing the number of frames that have timed out. This state is also reported on the MAPS dashboard.

The following example displays a typical RASlog entry for this condition.

```
2016/01/30-20:15:59, [MAPS-1001], 2/2, SLOT 5 | FID 1, CRITICAL, DCX_1, F-Port 1/19,
Condition=ALL_F_PORTS(DEV_LATENCY_IMPACT==IO_FRAME_LOSS),
Current Value:[DEV_LATENCY_IMPACT,IO_FRAME_LOSS, 1150 C3TXO Timeouts],
RuleName=C3TXTO_RULE, Dashboard Category=Fabric Performance Impact.
```

Two types of frame loss events are reported:

- Frame timeouts
- Latency causing severe delays but without frame timeouts.

Frame loss at an F_Port is explicitly tracked at the F_Port. However, frame loss at an E_Port might be caused by high R_RDY delays on F_Ports (which could be affected by various conditions, such as edge hold time, number of E_Ports, and other conditions). The delays might not be significant enough to cause timeouts on the F_Ports, but they could cause backups and timeouts at the E_Ports. These delays have an impact as serious as a frame timeout, so both conditions are logged as an IO_FRAME_LOSS state. The RASlog message contains the specific condition that triggered the IO_FRAME_LOSS rule.

Transient queue latency counter monitoring

MAPS monitors port transient queue latency (TXQ) for every port to identify ports that might be causing congestion in the SAN network so that you can take corrective action before the congestion affects other parts of the fabric.

Congestion occurs when the traffic being carried on a port exceeds its capacity to pass that traffic along efficiently. Such congestion increases the latency (the time lag between when a frame is sent and when it is received). Typical sources of congestion are links, hosts, or attached storage devices that are responding more slowly than expected. Early congestion detection is critical to maintaining a fabric's performance, because increased latency on a switch or port can propagate through a switch to the network as a whole. The cumulative effect of latency on many individual devices on a port can degrade the performance of the entire port. While the latency for each device might not be a problem, the presence of too many flows with significant latency passing through a port could become a problem.

Brocade Adaptive Networking uses virtual circuits (VCs) to facilitate fabric-wide resource monitoring. The VC architecture provides a flexible way to apply Quality of Service (QoS) monitoring for applications in virtual server environments. If the latency for a VC queue is high, the performance of the traffic flows assigned to that queue will be degraded. While latency is calculated individually for each VC in a port at a rate of once per second, transient latency is monitored only at the VC level, not at the port level. The latency of the VC having the greatest latency time is what is used to determine when an action is triggered. When the latency value crosses the threshold specified in a rule, the configured actions are taken for the virtual circuit for the given port. Possible actions include RASLog, SNMP, or e-mail notifications, as well as port toggling and slow drain device quarantining. Refer to Port toggling support on page 120 and Slow Drain Device quarantining on page 121 for specific information on these features. Default rules for these actions are included in all three default policies.

In previous versions, MAPS provided notifications only when it detected an impact due to latency (I/O performance impact or frame losses) on an F_Port, and the latency calculation was based on the inter-frame latency time. In this version, the TXQ latency calculation is based on the actual latency time for a frame, that is, the amount of time it takes for a frame to move out of the switch after it arrives at the time when the sample is taken, and is applicable to all ports.

For determining transient queue latency, MAPS has two predefined threshold states: IO_PERF_IMPACT and IO_FRAME_LOSS. The IO_PERF_IMPACT state is set for a port when latency is between the pre-defined low threshold and high threshold values, the IO_FRAME_LOSS state is set for a port when latency is greater than the pre-defined high threshold value.

The following example displays typical RASLogs created when IO_FRAME_LOSS and IO_PERF_IMPACT states are set.

```
 2016/01/19-21:18:00, [MAPS-1001], 979, FID 128, CRITICAL, SWAT_MAPS_TOM, E-Port 9,
Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_FRAME_LOSS),
Current Value:[DEV_LATENCY_IMPACT,IO_FRAME_LOSS, 165 ms Frame Delay],
RuleName=FRAME_LOSS_RULE, Dashboard Category=Fabric Performance Impact.

2016/01/19-21:18:50, [MAPS-1001], 979, FID 128, CRITICAL, SWAT_MAPS_TOM, E-Port 9,
Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_PERF_IMPACT),
Current Value:[DEV_LATENCY_IMPACT,IO_FRAME_LOSS, 15 ms Frame Delay],
RuleName=FRAME_LOSS_RULE, Dashboard Category=Fabric Performance Impact.
```

The following example of the **mapsdb --show** command shows that port 8/8 is having a latency problem. In this example, the output has been trimmed to focus on the explicit rule, and the backslash character (\) indicates a break inserted because the output is too long to display here as a single line.

```
switch:admin> mapsdb --show
     (output truncated)
3.1 Summary Report:
===================
Category                 |Today                      |Last 7 days                |
-------------------------------------------------------------------------------
Port Health              |Out of operating range     |No Errors                  |
BE Port Health           |No Errors                  |No Errors                  |
Fru Health               |In operating range         |In operating range         |
Security Violations      |No Errors                  |No Errors                  |
Fabric State Changes     |In operating range         |No Errors                  |
Switch Resource          |In operating range         |In operating range         |
Traffic Performance      |In operating range         |In operating range         |
FCIP Health              |Out of operating range     |No Errors                  |
Fabric Performance Impact|Out of operating range     |In operating range         |

3.2 Rules Affecting Health:
===========================

Category(Rule Count)        |Repeat Count|Rule Name                         |\
-------------------------------------------------------------------------------\
Fabric Performance impact(1)|1           |defALL_ALL_PORTS_IO_LATENCY_CLEAR|\

\ Execution Time   |Object   |Triggered Value(Units)|
\ ----------------------------------
\ 03/19/16 21:12:01|Port 8/8 |IO_LATENCY_CLEAR      |
```

MAPS only provides the port number of the switch as part of the TXQ latency alert; you must use the **portstatsshow** command to determine exactly which virtual circuits in the port are causing the problem. The following example uses the **portstatsshow** command. The ports shown for tim_latency_vc are the problem ports.

```
switch:admin> portstatsshow 1
     (output truncated)
stat_mc_tx             0           Multicast frames transmitted
tim_rdy_pri            0           Time R_RDY high priority
tim_txcrd_z            0           Time TX Credit Zero (2.5Us ticks)
tim_txcrd_z_vc  0- 3:  0           0           0           0
tim_txcrd_z_vc  4- 7:  0           0           0           0
tim_txcrd_z_vc  8-11:  0           0           0           0
tim_txcrd_z_vc 12-15:  0           0           0           0
tim_latency_vc  0- 3:  1           1           1           1
tim_latency_vc  4- 7:  1           1           1           1
tim_latency_vc  8-11:  1           1           1           1
tim_latency_vc 12-15:  1           1           1           1
fec_cor_detected       0           Count of blocks that were corrected by FEC
(output truncated)
```

Buffer credit zero counter monitoring

Buffer credit zero counter increments are indirect indications of latency; they indicate when frames were not transmitted through a port due to a delay in receiving R_RDY frames.

MAPS monitors this latency using a sliding window algorithm applied over a preset time period. This allows MAPS to monitor the frame delay over multiple window sizes with a different threshold for each time window. When a violation occurs, the latency is reported as IO_PERF_IMPACT in the RASLog message; the message includes both the bandwidth loss amount and the corresponding time window. The message specifies the actual increment of the counter as a percentage.

The following example displays a typical RASLog entry for this condition. In this example, the bandwidth loss is 85% and the time window is 1 second.

```
2016/01/30-21:10:00, [MAPS-1003], 489, SLOT 5 | FID 128, WARNING, SWAT_MAPS_TOM F-Port 7/28,
Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_PERF_IMPACT),
Current Value:[DEV_LATENCY_IMPACT,IO_PERF_IMPACT, 85.0% in 1 secs],
RuleName=PERF_IMPACT_RULE, Dashboard Category=Fabric Performance Impact.
```

NOTE

Buffer credit zero counters are monitored only on F_Ports.

Thresholds are monitored as follows:

- 70% of CRED_ZERO counter increment in 1 second.
- 50% of CRED_ZERO counter increment in 5 seconds.
- 30% of CRED_ZERO counter increment in 10 seconds.

Latency state clearing

When a frame timeout is detected, back pressure caused by either the buffer CRED_ZERO counter value or the queue latency condition is cleared. If there is a rule in the active policy to monitor the IO_LATENCY_CLEAR state, a RASLog message is posted. The following example is a sample of this message.

```
 2016/01/30-21:11:00, [MAPS-1004], 268895, SLOT 4 | FID 128, CRITICAL, SWAT_MAPS_TOM, Port 8/8,
Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_LATENCY_CLEAR),
Current Value:[DEV_LATENCY_IMPACT,IO_LATENCY_CLEAR],RuleName=defALL_ALL_PORTS_IO_LATENCY_CLEAR,
Dashboard Category=Fabric Performance Impact.
```

NOTE

The CRED_ZERO counters are monitored only on F_Ports.

Zoned device ratio monitoring

MAPS allows monitoring of the zoned device ratio per port.

Starting with Fabric OS 8.0.1, devices can be zoned to other devices to allow communications. When MAPS finds that a port has more than the expected number of devices zoned-in, then it alerts the administrator. Zone configuration can cause back pressure in the following situations:

1. If a device is zoned with a disproportionate number of devices.
2. If a port is allowed to communicate with a disproportionate number of ports.

MAPS uses the following rules to support this monitoring.

```
From MAPS FS 03-30-16
Rule name                   |Condition                       |Actions          |Policy      |
-------------------------------------------------------------------------------------------------
defALL_LOCAL_PIDSIT_FLOW_8  |ALL_LOCAL_PIDS(IT_FLOW/NONE>8)  |RASLOG,SNMP,EMAIL|Aggressive  |
-------------------------------------------------------------------------------------------------|
defALL_LOCAL_PIDSIT_FLOW_16 |ALL_LOCAL_PIDS(IT_FLOW/NONE>16)|RASLOG,SNMP,EMAIL|Moderate    |
-------------------------------------------------------------------------------------------------|
defALL_LOCAL_PIDSIT_FLOW_32 |ALL_LOCAL_PIDS(IT_FLOW/NONE>32)|RASLOG,SNMP,EMAIL|Conservative|
-------------------------------------------------------------------------------------------------|
```

Example of dashboard output for zoned device ratio monitoring

```
> mapsdb --show

1 Dashboard Information:
========================

Ports with highest Zoned device ratio:             0x20000,0x200200
```

Examples of using the logical group supporting zoned device ratio monitoring

A logical group, ALL_LOCAL_PIDS, has been added to help you manage all the PIDs. The following are examples of using the logical group:

```
> logicalgroup --show all_local_pids details
-------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type              |Member Count |Members
-------------------------------------------------------------------------------------------------------
ALL_LOCAL_PIDS                  |Yes        |Pid               |6480         |All Pids monitored

B)
> logicalgroup --show | grep ALL_LOCAL_PIDS
ALL_LOCAL_PIDS                  |Yes        |Pid               |6480         |All Pids monitored
```

NOTE

The logical group, ALL_LOCAL_PIDS, is not supported in FICON environments, because they use flat zoning.

Considerations for zoned device ratio monitoring

The following information should be considered for zoned device ratio monitoring:

- MAPS also shows the top five PIDs in the dashboard output.
- Quiet time is not supported with zoned device ratio monitoring; therefore, any rules that specify quiet time will fail.
- Monitoring of zoned device ratios starts an hour after a system reboot, provided the fabric has been formed.

NOTE

Duplicate alerts for the same PID might occur if the system comes up a few minutes before midnight. However, after the system has been up for more than a day, then the rule is executed and alerts are generated once per day.

MAPS and Bottleneck Detection

Starting with Fabric OS 8.0.0, bottleneck detection functionality has been replaced by Fabric Performance Impact (FPI) monitoring; the legacy bottleneck monitoring feature is obsolete.

The MAPS dashboard displays the stuck virtual channel (VC) on any port. It also identifies the ports on which bottlenecks are seen, and then it sorts them based on the number of seconds that they exceeded the bottleneck threshold. This identifies the most strongly affected ports, no matter what the cause.

The bottleneck information appears in the "Rules Affecting Health" section as part of the Port Health category. The "History Data" section displays entries that have cred_zero counters that are not zero. If the cred_zero counter increases for a port but no bottleneck time is recorded, this indicates a potential transient bottleneck on the port.

In the following example, the last three lines list bottlenecks, with the final bottleneck caused by a timeout rather than a numeric value. Be aware that the column headings in the example have been edited slightly so as to allow the example to display clearly.

```
3. Rules Affecting Health:
==========================
Category(RuleCnt)|RptCnt|Rule Name                      |Execution Time  |Object   |Triggered Value
                                                                                      (Units)             |
-------------------------------------------------------------------------------------------------------------
Port Health(12)  |1     |defALL_OTHER_F_PORTSLR_10     |03/21/16 0:30:06|D_Port 23|11                   |
                  1     |defALL_OTHER_F_PORTSLR_5      |03/21/16 0:29:54|D_Port 23|7                    |
                  1     |defALL_OTHER_F_PORTSC3TXTO_3  |03/21/16 0:29:36|D_Port 23|57                   |
                  1     |defALL_OTHER_F_PORTSC3TXTO_10|03/21/16 0:29:36|D_Port 23|57                   |
                  6     |Bottleneck_stuckvc            |03/21/16 0:30:24|D_Port 23|STUCKVC              |
                              |
     (output truncated)
```

When a latency rule is triggered, the instance is listed as part of the Traffic Performance category. In the both the "Front-end port History Data" section and the "Back-end port History Data" sections, the five ports with the longest total backpressure times since the previous midnight are shown, as shown in the following example. Be aware that the headings in the example have been edited slightly so as to allow the example to display clearly.

```
3. Rules Affecting Health:
==================================
Category(RuleCnt)     |RptCnt|Rule Name                       |Execution Time  |Object   |Trig Val(Units)|
Fabric Perf Impact(5)|2     |defALL_PORTS_IO_PERF_IMPACT     |03/21/16 0:30:6 |F_Port 13|IO_PERF_IMPACT
                                                               03/21/16 10:30:6|F_Port 22|IO_PERF_IMPACT
                       3     |defALL_PORTS_IO_FRAME_LOSS      |03/21/16 0:30:6 |F_Port 3 |IO_FRAME_LOSS
                                                               03/21/16 10:30:6|F_Port 2 |IO_FRAME_LOSS
                                                               03/21/02 10:30:6|F_Port 4 |IO_FRAME_LOSS
4.1 Front-end port History Data:
================================
Stats(Units)         Current   03/21/16   03/14/16   --/--/--   --/--/--   --/--/--   --/--/--
                     Port(val) Port(val)  Port(val)
-------------------------------------------------------------------------------------------
CRC(CRCs)            13(20)     -          -          -          -          -          -
ITW(ITWs)            -          13(612)    -          -          -          -          -
LOSS_SYNC(SyncLoss) -          -          -          -          -          -          -
LF                   -          -          -          -          -          -          -
LOSS_SIGNAL(LOS)     12(4)      12(4)      13(5)      -          -          -          -
                     -          13(4)      12(4)      -          -          -          -
                     -          14(4)      14(4)      -          -          -          -
PE(Errors)           -          -          -          -          -          -          -
STATE_CHG            12(5)      12(5)      12(9)      -          -          -          -
                     -          13(5)      13(9)      -          -          -          -
                     -          14(5)      14(9)      -          -          -          -
LR                   -          13(6)      12(10)     -          -          -          -
                     -          12(4)      13(10)     -          -          -          -
                     -          14(4)      14(10)     -          -          -          -
C3TXTO(Timeouts)     -          -          -          -          -          -          -
RX(%)                -          -          -          -          -          -          -
TX(%)                -          -          -          -          -          -          -
UTIL(%)              -          -          -          -          -          -          -
```

```
BN_SECS(Seconds)     -          -          -          -          -          -          -

4.2 Back-end port History Data:
================================
Stats(Units)         Current    03/21/16   03/14/16   --/--/--   --/--/--   --/--/--   --/--/--
                     Port(val)  Port(val)  Port(val)
-------------------------------------------------------------------------------------------
CRC(CRCs)            2/1/0(15)      -          -          -          -          -          -
LOSS_SYNC(SyncLoss)  2/1/0(1)   3/3/1(2)    3/3/1(2)          -          -          -          -
```

NOTE

The MAPS dashboard will continue to log events whether RASLogs are set to on or off in the bottleneck configuration.

Refer to Bottleneck detection with the MAPS dashboard for additional details. Refer to the "Bottleneck Detection" chapter in the *Fabric OS Administrator's Guide* for specific command details and bottleneck monitoring parameters.

Port toggling support

MAPS supports port toggling, which allows Fabric OS to recover a port that has been bottlenecked by a target device. While there are many reasons why the target device could be bottlenecked, one of the most common is a temporary glitch in an adapter or its software.

Port toggling in MAPS temporarily disables a port and then re-enables it, allowing the port to reset and recover from the issue. If the port does not recover, Fabric OS suspends the port, forcing the port traffic to switch over to a different path if one is available.

To enable recovering ports using port toggling, MAPS assumes that there is a redundant path to the target device. It does not check to see if there is one, nor can it check to see if traffic to or from the target device has been switched over to a redundant path. MAPS also assumes that while the port is being toggled, the operational state of the port will not be changed by any other mechanism, such as an administrator disabling or moving the port, or a port fencing operation.

NOTE

Port toggling cannot be used in conjunction with automatic VC quarantining, as this might result in unpredictable behavior.

Port toggling is enabled within MAPS by including the toggle action within the rule and specifying a value from 2 through 3600 seconds as the length of time that the port is to be disabled.

The following example defines a rule that toggles a port offline for 180 seconds (3 minutes) when the number of CRC errors in a minute on the port is greater than 0.

```
switch:admin> mapsrule --config toggle_rule -group DB_PORTS -monitor DEV_LATENCY_IMPACT -timebase none -op
eq -value IO_PERF_IMPACT -action TOGGLE -tt 180
```

When a port has been toggled by a MAPS rule, TOGGLE appears as a notification action in the output of the **mapsconfig** and **mapsdb** commands. The following example displays a sample of the **mapsdb--show** command that illustrates this result.

```
switch:admin> mapsdb --show
1 Dashboard Information:
========================
DB start time:                  Fri Mar 11 18:38:12 2016
Active policy:                  test_xy1
Configured Notifications:       RASLOG,FENCE
Fenced Ports :                  None
Decommissioned Ports :          None
Fenced circuits :               38/0,38/1,38/2,38/3,38/4,38/5,38/6,38/7
Quarantined Ports :             None

       (output truncated)
```

When MAPS toggles a port, the **switchshow** command lists the reason for the port being disabled as "Transient".

The following example displays a sample output for **switchshow** when a port has been toggled. In this example, Port 65 is listed as "Disabled (Transient)".

```
switch:admin> switchshow
444  8 28   ------   cu   8G    No_Sync   FC
445  8 29   ------   cu   8G    No_Sync   FC
446  8 30   ------   cu   8G    No_Sync   FC
447  8 31   ------   cu   8G    No_Sync   FC
 64  9  0   014000   id   N2    Online    FC  F-Port  10:00:00:00:00:01:00:01
 65  9  1   014100   id   N4    In_Sync   FC  Disabled (Transient)
 66  9  2   014200   --   N8    No_Module FC
 67  9  3   014300   --   N8    No_Module FC
 68  9  4   014400   --   N8    No_Module FC
```

Slow Drain Device quarantining

In a fabric, many flows share the same link or virtual circuit (VC). However, the credits used to send traffic or packets across the link are common to all the flows using the same link. Because of this, a slow-draining device may slow down the return of credits and have a negative effect on the healthy flows through the link.

The following figure illustrates how this functions. In it, the inability of the slow-draining device (1) to clear frames quickly enough is causing backpressure not just to the edge device (2), but also to the core (3) and the appliance (4) that is trying to reach the slow-draining device.

FIGURE 1 Slow-drain data flow

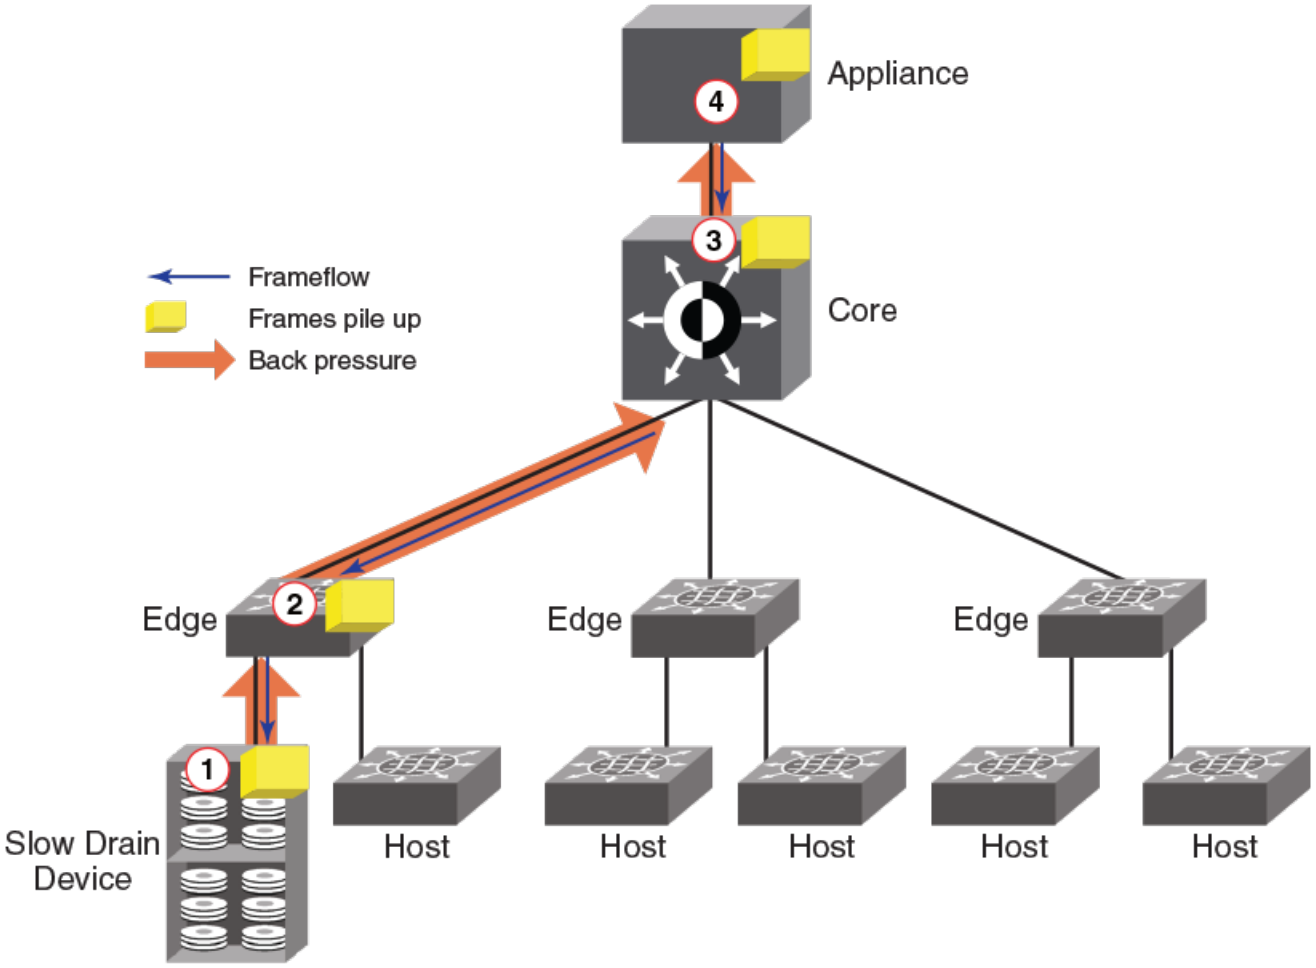

To remedy this, Brocade created Slow Drain Device Quarantine (SDDQ), a feature which —in conjunction with Quality of Service (QoS) monitoring— allows MAPS to identify a slow-draining device and quarantine it by automatically moving all traffic destined to the F_Port that is connected to the slow-draining device to a low-priority VC so that the traffic in the original VC does not experience backpressure.

MAPS accomplishes the isolation of the slow-draining device by moving the slow-draining port to the MAPS ALL_QUARANTINED_PORTS group, and then notifying the Name Server to use a low-priority virtual circuit for all flows directed to the slow-draining port. Once the traffic targeted to the port is isolated, MAPS moves the port to the quarantine group. After this is done, MAPS continues to monitor the device latency on the port. If MAPS clears the latency impact state by changing the status to IO_LATENCY_CLEAR, it creates a RASLog or other notice for that change, but the port is kept in the quarantine group. If any quarantined ports go offline or become disabled, the ports remain in the quarantined group. When the quarantined ports come online, MAPS automatically notifies the Name Server that the port is still quarantined.

To restore the traffic targeted to a quarantined port to normal QoS priority, the **sddquarantine --clear** command is the only way to move the port out of the quarantine group. Refer to Clearing quarantined ports on page 127 for more details on this procedure.

Slow Drain Device Quarantine licensing

For Slow Drain Device Quarantining (SDDQ) to take effect, the Fabric Vision license must be installed on the switch where the slow draining device is detected, as well as on the switch where the quarantine action is to occur. Intermediate switches do not need the Fabric Vision license for this feature to work, but they must have QoS enabled on all ISLs.

> NOTE
> The Quality of Service (QoS) feature license, also known as the Adaptive Networking (AN) feature, does not need to be installed in order to use the QoS feature. The QoS feature can be configured without the licenses in Fabric OS 7.2.0 and later.

If the Fabric Vision license is not installed on a switch which has slow-draining devices (either remote or local) attached, and no quarantine action has been applied previously to those devices, simply enabling the Fabric Vision license on that switch will not immediately trigger the quarantining action. However, the quarantine action will be applied to all zoned devices coming online after the Fabric Vision license has been enabled. In order to have the same behavior across all Fabric Vision-licensed switches, Brocade recommends that you toggle the switch, devices, or ports in the zone involving the slow-draining device after you enable the Fabric Vision license.

When MAPS is installed and the license is active, SDDQ is available to all switches in the fabric that can receive the SDDQ registered state change notification (RSCN). It does not matter if the quarantine action is explicitly enabled on that switch or not. However, those switches that do not have the SDDQ action enabled will not be able to flag or isolate any slow-draining devices.

Notes on Slow Drain Device Quarantining

The following items should be kept in mind when using Slow Drain Device Quarantining (SDDQ):

- SDDQ is disabled by default on installation.
- SDDQ is supported:
  - For N_Port ID Virtualization (NPIV) devices connected to F_Ports, but not for ports connected to Access Gateway devices.
  - On all Fabric OS platforms that are running Fabric OS 7.4.0, including FICON environments.
- SDDQ is not supported:
  - On switches running Fabric OS versions earlier than 7.4.0.
  - On switches running in Access Gateway mode.
  - On F_Port trunks.
  - In Fibre Channel Routing (FCR) backbone fabrics, but can be supported on devices within an FCR edge fabric. However, the edge fabric will not apply this quarantine feature to the flows for any FCR-imported devices.
- The SDDQ action is supported by the default rules defALL_IO_PERF_IMPACT and defALL_IO_FRAME_LOSS.
- While MAPS does allow you to enable and disable SDDQ for individual logical switches, Brocade recommends that you enable SDDQ for the entire fabric. The reason for this is that it is difficult to predict when a device may become slow-draining.
- You cannot add or delete members from the ALL_QUARANTINED_PORTS group. No default rules are defined for the group, but you can create custom rules to monitor the ports in this group.
- When a device is marked as being slow-draining, only the flows destined to it are shifted to the low-priority Virtual Circuit (VC). Flows in the reverse direction are not affected.
- QoS zoning rules may also route flows to low-priority virtual channels, which may further slow these flows. In addition, if a device is already in a QoS zone and the device has been identified as slow-draining, then any flows to that device will be assigned a low priority independent of the QoS zone priority until you explicitly restore the device to its original VC. Alternatively, if a QoS zone is newly configured with the slow-draining device and pushed across the fabric, the flow will not

shift to the new priority because it is still marked as a slow-draining flow. It will be kept in the low-priority VC until you remove it from that VC.

- If device latency is due to Class 3 timeouts (C3TXTO) and the active C3TXTO rule has port fencing as an action, then the port fencing (disabling) may be performed first and quarantining will not occur.
- If there are switches in the fabric that do not have QoS feature support, then end-to-end slow drain flow isolation may not be possible.
- The maximum number of devices that can be isolated per unit (chassis or fixed-port switch) is 32. The default value is 8. This is controlled by the chassis-wide "Chassis SDDQ limit" user-configurable key, which is set using the **configurechassis** command. Refer to the *Fabric OS Command Reference* for more information on this command.
- To prevent SDDQ from consuming an excessive amount of system resources in FICON environments and large-scale fabrics, the SDDQ action is blocked in the following scenarios to prevent spurious IOCTLs (input/output control actions) due to quarantine action on hundreds of source ports:
  - If the total number of ports zoned to the slow drain port is set to 32.
  - If the defzone is labeled as "all access", and there is no user-defined zoning configuration.

  NOTE
  The Port Toggling action should not be used in conjunction with SDDQ, as this may result in unpredictable behavior.

Enabling Slow Drain Device Quarantining

Slow Drain Device Quarantining (SDDQ) is enabled by having a valid Fabric Vision license, and is activated by including the SDDQ action in the configured MAPS action list.

To enable SDDQ, complete the following steps.

1. Connect to the device and log in using an account with admin permissions.
2. Use the **mapsconfig --actions** command, and include SDDQ as one of the actions.

The following example enables SDDQ as an available action.

```
switch:admin> mapsconfig --actions SNMP,RASLOG,EMAIL,SDDQ
```

Disabling Slow Drain Device Quarantining

To disable the Slow Drain Device Quarantine feature, exclude SDDQ from the configured list of MAPS actions, as shown in the following procedure.

1. Connect to the device and log in using an account with admin permissions.
2. Use the **mapsconfig --actions** command, and do not include SDDQ as one of the actions.

The following example removes Slow Drain Device Quarantining from the list of available actions.

```
switch:admin> mapsconfig --actions SNMP,RASLOG,EMAIL
```

Confirming the slow-draining status of a device

You can use the **nsshow**, **nscamshow**, and **nodefind** commands to verify that a device is slow-draining.

The following example shows the output of the **nsshow** command where there is a slow-draining device. The identifying line has been called out in this example. If there not a slow-draining device, the line would not appear.

```
switch:admin> nsshow
{
Type Pid    COS     PortName                NodeName                 TTL(sec)
N    010000;    2,3;20:00:00:05:1e:92:e8:00;20:00:00:05:1e:92:e8:00; na
    Fabric Port Name: 20:00:00:05:1e:92:e8:00
    Permanent Port Name: 20:00:00:05:1e:92:e8:00
    Port Index: 0
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
    Port Properties: SIM Port
N    014a00;    2,3;30:0a:00:05:1e:84:b5:c3;10:00:00:05:1e:84:b5:c3; na
    FC4s: FCP
    NodeSymb: [42] "dsim:fdmi_host:sw_5300edsim[172.26.26.188]"
    Fabric Port Name: 20:4a:00:05:1e:92:e8:00
    Permanent Port Name: 30:0a:00:05:1e:84:b5:c3
    Port Index: 74
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
N    015000;      3;10:00:00:00:00:01:00:01;10:00:00:00:00:00:01:01; na
    Fabric Port Name: 20:50:00:05:1e:92:e8:00
    Permanent Port Name: 10:00:00:00:00:01:00:01
    Port Index: 80
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
    Slow Drain Device: Yes     <== Slow-draining device identified
N    015100;      3;10:00:00:00:00:02:00:01;10:00:00:00:00:00:02:01; na
    Fabric Port Name: 20:51:00:05:1e:92:e8:00
    Permanent Port Name: 10:00:00:00:00:02:00:01
    Port Index: 81
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
The Local Name Server has 4 entries }
```

The following example shows the output of the **nscamshow** command where there is a slow-draining device. The identifying line has been called out in this example. If there was not a slow-draining device, the line would not appear.

```
switch:admin> nscamshow
nscam show for remote switches:
Switch entry for 2
  state  rev    owner  cap_available
  known   v740  0xfffc01 1
  Device list: count 7
    Type Pid    COS     PortName                NodeName
    N    020a00;      3;10:00:00:00:00:04:00:01;10:00:00:00:00:00:04:01;
        Fabric Port Name: 20:0a:00:05:1e:84:98:c7
        Permanent Port Name: 10:00:00:00:00:04:00:01
        Port Index: 10
        Share Area: No
        Device Shared in Other AD: No
        Redirect: No
        Partial: No
    N    020b00;      3;10:00:00:00:00:0a:00:01;10:00:00:00:00:00:0a:01;
        Fabric Port Name: 20:0b:00:05:1e:84:98:c7
        Permanent Port Name: 10:00:00:00:00:0a:00:01
        Port Index: 11
        Share Area: No
        Device Shared in Other AD: No
        Redirect: No
        Partial: No
        Slow Drain Device: Yes     <== Slow-draining device identified
    N    021000;      3;10:00:00:00:00:05:00:01;10:00:00:00:00:00:05:01;
        Fabric Port Name: 20:10:00:05:1e:84:98:c7
        Permanent Port Name: 10:00:00:00:00:05:00:01
        Port Index: 16
        Share Area: No
        Device Shared in Other AD: No
        Redirect: No
        Partial: No
```

The following example shows the output of the **nodefind** command where there is a slow-draining device. The identifying line has been called out in this example. If there was not a slow-draining device, the line would not appear.

```
switch:admin> nodefind 015000
Local:
 Type Pid COS PortName NodeName SCR
 N 015000; 3;10:00:00:00:00:01:00:01;10:00:00:00:00:00:01:01; 0x00000003
    Fabric Port Name: 20:50:00:05:1e:92:e8:00
    Permanent Port Name: 10:00:00:00:00:01:00:01
    Device type: Physical Unknown(initiator/target)
    Port Index: 80
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
    Slow Drain Device: Yes     <== Slow-draining device identified
    Aliases:
```

Displaying quarantined ports

MAPS allows you to display a list of ports which have been quarantined in the ALL_QUARANTINED_GROUP.

To display a list of ports which have been quarantined by MAPS, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **sddquarantine --show** to view a list of the quarantined ports.

The following example shows the offline quarantined local ports and the online quarantined device information across the fabric.

```
switch:admin> sddquarantine --show
-----------------------------------------------------------
Ports marked as Slow Drain Quarantined in the Local Switch:   2/1, 1/3
-----------------------------------------------------------
Ports marked as Slow Drain Quarantined but not enforced:   1/3
-----------------------------------------------------------

Online Quarantined Devices across the fabric
-----------------------------------------------------------
 Port Index |   PID   |            PWWN
-----------------------------------------------------------
      17    | 051100 | 30:10:00:05:33:ac:c6:13
      17    | 051101 | 30:10:01:05:33:ac:c6:13
-----------------------------------------------------------
```

Clearing quarantined ports

MAPS allows you to remove ports which have been quarantined from the quarantine group (ALL_QUARANTINED_GROUP).

To remove ports which have been quarantined by MAPS from the quarantine group, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **sddquarantine --show** to view a list of the quarantined ports.
3. Enter **sddquarantine --clear** followed by either:
   - The slot/port ID that is to be removed from the quarantine group
   - The keyword **all** to clear all quarantined ports

   If the port is offline, the port is removed from the ALL_QUARANTINED_GROUP. There is no confirmation of this action.

   A port is not allowed to be cleared from quarantine while the latency or frame loss condition that caused it to be quarantined persists, as the port will not be flagged again until the condition is cleared and then retriggered. However, you can override this by using the **-force** keyword as part of the **sddquarantine** command.

   NOTE
   The “Ports marked as Slow Drain Quarantined but not enforced” line only appears in the output when devices could not be quarantined because the chassis SDDQ limit of 32 quarantined devices has been reached.

The first part of the following example uses the **sddquarantine --show** command to display the offline quarantined local ports and the online quarantined device information across the fabric. The second part shows using the **sddquarantine --clear** command to remove the ports from quarantine. Notice that port 3 was not able to be cleared using the **all** option as it was still in either the IO_FRAME_LOSS or the IO_PERF_IMPACT condition, but this port was able to be cleared using the **-force** option. These two options can also be used together.

```
switch:admin> sddquarantine --show
-----------------------------------------------------------
Ports marked as Slow Drain Quarantined in the Local Switch:   2/1, 1/3
-----------------------------------------------------------
Ports marked as Slow Drain Quarantined but not enforced:   1/3
-----------------------------------------------------------

Online Quarantined Devices across the fabric
-----------------------------------------------------------
 Port Index |   PID   |            PWWN
-----------------------------------------------------------
      17    | 051100 | 30:10:00:05:33:ac:c6:13
      17    | 051101 | 30:10:01:05:33:ac:c6:13
-----------------------------------------------------------

switch:admin> sddquarantine --clear 2/1, 1/3
Initiated clearing port from quarantined state

switch:admin> sddquarantine --clear 1/3
Clear failed since port is still in latency state

switch:admin> sddquarantine --clear all
The clear action was not initiated for the following port(s). Try with individual ports
                3
Initiated quarantine action on other ports

switch:admin> sddquarantine --clear 1/3 -force
Initiated clearing port from quarantined state
```

Slow Drain Device quarantining and FICON

Brocade does not recommend that you enable the Slow Drain Device Quarantine (SDDQ) feature in a FICON environment because a FICON environment typically has a single zone for all devices. However, if you are using 1-to-1 zoning, SDDQ may help with slow-draining device issues.

When SDDQ is enabled, all traffic to a slow-draining device is moved to the lowest-priority virtual circuit. In a single-zone environment this impacts all traffic in the zone.

In a FICON environment, the maximum number of devices that can be isolated per unit (chassis or fixed-port switch) is 32, and the default number of devices is 10. The first time a port is detected as a slow drain, the Slow Drain Device quarantine will be applied to it. The maximum number of ports that can be zoned with a slow-draining port is 32. If the 32-port zone limit is exceeded, the quarantine action is not taken, and any new zoned device or (33rd) port coming online after this limit is reached will not be quarantined, even if it is slow-draining. However, MAPS will add the problem port to the quarantine list. This means that once the number of zoned ports is back below the limit, the next time the listed port comes online the listed port may be quarantined.

If there are quarantined slow-draining ports in a fabric and their zone changes, these ports will remain quarantined even if this violates the limit on the number of ports that can be zoned together. Once the limit is reached, any reduction in the number of ports in the zone will not affect the quarantine decision until the next time MAPS detects the slow drain condition. Be aware that if the slow drain conditions persist, you may not get another alert, as MAPS has already raised an alert. If a slow-draining device port avoids quarantine due to a limit on the number of ports in the zone, any zoned device connecting to that port that comes online after the limit check will not be quarantined.

The **sddquarantine --show** command output displays the list of successfully quarantined ports as well as those that were identified as problematic but were not quarantined, so that you can take any necessary measures to recover these ports. For example, you might choose to disable the ports, or (after correcting what is causing the slow drain) restore the ports. The restoration action for quarantined ports does not have any constraints and is not affected by any zone restrictions.

NOTE

To avoid excessive consumption of system resources, Slow Drain Device quarantining based on IOCTL errors is not permitted.

Other MAPS monitoring capabilities

- Scalability limit monitoring....131
- MAPS Service Availability Module....136
- MAPS monitoring for Extension platforms....138
- IPEXT monitoring....140
- E-mail delivery monitoring....142
- Fan air-flow direction monitoring....143
- Updating monitoring policies for devices with four PSUs....145

Scalability limit monitoring

MAPS monitors changes of fabric-level monitoring systems. These systems all have scalability limits which MAPS can monitor and send alerts using RASLog entries, SNMP messages, or e-mail. The monitoring results are captured in the MAPS dashboard under the "Fabric State Changes" category. Although there are default rules that monitor these values, MAPS allows you to define new rules with different thresholds and actions.

MAPS can monitor the following scalability limits:

- The number of logged-in device connections in a pure Layer 2 fabric.
- The size of the zone configuration resource that is used.
- The number of Fiber Channel Router configurations.
- The number of Logical SAN (LSAN) device connections (this includes both edge fabric and Backbone fabric device connections).

> **NOTE**
> MAPS monitors the device count per FCR switch.

When a rule is triggered, the corresponding RASLogs appear in the summary section of the dashboard. The following example shows two rules (LSAN_DEVCNT_PER and L2_DEVCNT_PER) have been triggered. The column headings in the example have been edited slightly so as to allow the example to display clearly.

```
3.1 Summary Report:
===================
Category                  |Today                      |Last 7 days                |
--------------------------------------------------------------------------------
Port Health               |No Errors                  |No Errors                  |
BE Port Health            |No Errors                  |No Errors                  |
Fru Health                |In operating range         |In operating range         |
Security Violations       |No Errors                  |No Errors                  |
Fabric State Changes      |No Errors                  |No Errors                  |
Switch Resource           |In operating range         |In operating range         |
Traffic Performance       |In operating range         |In operating range         |
FCIP Health               |Not applicable             |Not applicable             |
Fabric Performance Impact|In operating range         |In operating range         |

3.2 Rules Affecting Health:
===========================
Category(Rule Count)    |RptCnt|Rule Name      |Execution Time  |Object    |Triggered Value (Units)|
-------------------------------------------------------------------------------------------------
Fabric State Changes(2)|1     |LSAN_DEVCNT_PER|03/21/16 00:30:6|D_Port 23|12 %                   |
                       |1     |L2_DEVCNT_PER  |03/21/15 01:04:6|D_Port 23|12 %                   |
```

For more detailed information on scalability limits, refer to *Brocade SAN Scalability Guidelines: Brocade Fabric OS 8.X*.

Layer 2 fabric device connection monitoring

A pure Layer 2 fabric is a collection of Fibre Channel switches and devices and switches that doesn't participate in a metaSAN. In such a fabric, rules for device counts are calculated as a percentage of the total number of devices. For example, a Layer 2 fabric with 5500 devices logged in is using 92 percent of the maximum limit of 6000 devices for a Layer 2 fabric. So if user have configured a rule to trigger an alert at 90 percent or greater, then MAPS triggers the action configured for that rule and sends the data to the dashboard.

LSAN device connection monitoring in a metaSAN

The collection of all devices, switches, edge and Backbone fabrics, LSANs, and routers that make up a physically connected but logically partitioned storage network is called a metaSAN. Using MAPS, the total number of LSAN device connections (including the total number of devices from all edge fabrics) in a metaSAN can be monitored for a scalability limit.

> NOTE
> MAPS rules for monitoring imported LSAN device connections in a metaSAN can be configured only on switches that are a part of the Backbone fabric.

Device counts in this framework are calculated as a percentage of the total number of LSAN devices in a metaSAN (including imported devices from all edge fabric). For example: if a fabric has four switches in the Backbone fabric and four switches each in four edge fabrics, the total number of LSAN devices in this metaSAN (including imported devices from all edge fabrics) is 1200. Given a maximum of 10000 devices, this is 12 percent. If you have configured a rule to trigger at 10 percent or greater, then MAPS triggers the action configured for the rule, but only on those switches that are part of the Backbone fabric, and caches the data in the dashboard.

Backbone fabric Fibre Channel router count monitoring

In a Backbone fabric, there can be maximum number of 12 Fibre Channel routers (FCRs). MAPS rules can be configured to monitor the number of Fibre Channel routers in the Backbone fabric as an absolute value. If the number of Fibre Channel routers reaches the configured threshold, MAPS triggers the action configured for the rule and caches the data in the dashboard. Refer to Default rules for scalability limit monitoring on page 135 for these values.

The following example shows a typical RASLog entry for exceeding the threshold for the number of Fibre Channel routers in the Backbone fabric:

```
2014/05/27-17:02:00, [MAPS-1003], 14816, SLOT 4 | FID 20, WARNING, switch_20, Switch,
Condition=SWITCH(BB_FCR_CNT>12), Current Value:[BB_FCR_CNT,13], RuleName= defSWITCHBB_FCR_CNT_12, Dashboard
Category=Fabric State Changes.
```

The following example shows a typical MAPS dashboard entry for exceeding the threshold for the number of Fibre Channel routers in the Backbone fabric. Be aware that the column headings in section 3.2 of the example have been edited slightly so as to allow the example to display clearly.

```
3.1 Summary Report:
===================
Category                 |Today                    |Last 7 days         |
-------------------------------------------------------------------------
Port Health              |No Errors                |No Errors           |
BE Port Health           |No Errors                |No Errors           |
Fru Health               |In operating range       |In operating range  |
Security Violations      |No Errors                |No Errors           |
Fabric State Changes     |Out of operating range   |No Errors           |
Switch Resource          |In operating range       |In operating range  |
Traffic Performance      |In operating range       |In operating range  |
FCIP Health              |No Errors                |No Errors           |
Fabric Performance Impact|In operating range       |In operating range  |

3.2 Rules Affecting Health:
===========================
Category(RuleCount) |RptCount|Rule Name             |Execution Time   |Object|Triggered Value(Units)|
-----------------------------------------------------------------------------------------------------
Fabric State Changes|1       |defSWITCHBB_FCR_CNT_12|05/27/14 17:02:00|Switch|13                    |
(1)                 |        |                      |                 |      |                      |
```

Zone configuration size monitoring

In Fabric OS 7.3.0 and later, MAPS can monitor zone configuration size. Based on the platform, a switch supports either a maximum zone configuration size of 1 MB or 2 MB. The monitoring value is calculated as a percentage of the zone configuration space used. If the configuration size reaches the configured threshold limit, MAPS triggers the action configured for the rule and caches the data in the dashboard. Refer to Default rules for scalability limit monitoring on page 135 for these limit values.

NOTE

MAPS zone configuration size monitoring is only for the default switch, as the total memory size is for the chassis as a whole. The maximum available zone configuration limit is determined at the chassis level and shared by all logical switches.

Monitoring NPIV logins to F_Ports

MAPS can monitor the number of N_Port ID Virtualization (NPIV) logins to individual F_Ports and generates RASLog, SNMP, or e-mail alerts if the login threshold for a port is reached.

NPIV is a Fibre Channel mechanism that allows host virtual machines to obtain a virtual world wide node (WWN) name. This allows multiple virtual machines to share the same physical Fibre Channel host bus adapter (HBA). When a switch gets a new connection request as part of the FLOGI, a unique N_Port ID is assigned to the device. This means that for each device there is a one-to-one mapping between the virtual world wide node name and the N_Port ID.

Monitoring NPIV logins to F_Ports is important because Access Gateway leverages NPIV to present open system Fibre Channel host connections as logical devices to SAN fabrics. This reduces switch-to-switch interoperability problems by allowing servers to seamlessly connect to SAN fabrics. However, as there is a limit to the number of logins that can be made to a switch, it is important to monitor this number and be able to alert administrators when the limit is approached.

NOTE

NPIV monitoring is not High Availability (HA)-capable. As a consequence, if there is a reboot or an HA failover, existing NPIV logins are not preserved, and new ones are assigned on a first-come, first-served basis.

MAPS supports monitoring all F_Ports in a switch for the number of NPIV logins as part of scalability limit monitoring. The value monitored is calculated as a percentage of the number of devices logged in relative to the maximum number of logins configured for that port. When the number of devices logged in to the switch reaches the rule threshold, MAPS posts a RASLog, SNMP, or e-mail message, allowing you to block any further logins. This threshold monitoring helps keep the switch from being overloaded with connection requests. For each port, MAPS allows you to configure the maximum number of logins value that is used in the calculation.

When a rule is triggered, the triggered rule name appears in the dashboard summary section as a Port Health item. The following example shows the relevant portion of the **mapsdb --show** command output for a switch configured with a maximum NPIV login limit of 100 for all ports. Be aware that the column headings in the example have been edited slightly so as to allow the example to display clearly.

```
     (output truncated)
Category                  Today               Last 7 Calendar Days
================================================================================
Port Health          :  Out of range             In range
BE Port Health       :  In range                 In range
FRU Health           :  In range                 In range
FCIP Health          :  In range                 In range
Security Violations  :  No Errors                   In range
Fabric Health        :  In range                    In range
Switch Resources     :  In range                    In range
Traffic Performance  :  In range                    In range

Rules Affecting Health:
================================
Category(RuleCount)|RptCnt|Rule Name                    |Execution Time    |Object     |Triggered Value|
                                                                                 (Units)         |
-------------------------------------------------------------------------------------------------------|
Port Health(2)     |1     |defALL_F_PORTSDEV_NPIV_LOGIN |01/14/15 12:59:58 |F-Port 7/11| 65.00 %       |
                   |1     |defALL_OTHER_F_PORTSLOSS_SYNC|01/14/15 12:52:34 |F-Port 7/1 | 1 SyncLoss    |
     (output truncated)
```

MAPS includes the following default rules for monitoring NPIV logins to F_Ports in the default policies. The counter monitored is the integer value in percentage and the timebase is "NONE."

TABLE 33 Default rules for monitoring NPIV logins to F_Ports

Policy	Rule Name	Condition	Actions
dflt_aggressive_policy	defALL_NPIV_LOGIN_PER_60	ALL_F_PORTS(DEV_NPIV_LOGINS/NONE > 60)	SNMP, RASLOG, EMAIL
dflt_moderate_policy	defALL_NPIV_LOGIN_PER_75	ALL_F_PORTS(DEV_NPIV_LOGINS/NONE > 75)	SNMP, RASLOG, EMAIL
dflt_conservative_policy	defALL_NPIV_LOGIN_PER_90	ALL_F_PORTS(DEV_NPIV_LOGINS/NONE > 90)	SNMP, RASLOG, EMAIL

Scalability limit monitoring assumptions and dependencies

The following assumptions and dependencies should be kept in mind when considering scalability limit monitoring.

- All the scalability limits are soft limits, not hard limits; the monitored value can be greater than 100 percent.
- The Backbone fabric can also have Layer 2 switches; these switches are not considered as part of any of the scalability limit metrics.
- The number of device connections in an edge fabric or Backbone fabric also have scalability limits themselves, and these cannot be monitored using MAPS.

- Scalability limit monitoring (using L2_DEVCNT_PER) occurs only at midnight. Therefor, if a switch is moved from being a part of the Layer 2 fabric to being a part of the edge fabric, the device count metrics (how many devices in the fabric) will not change until the next midnight.
- The "LSAN-imported device" metric is only monitored in switches that are a part of a Backbone fabric.
- Scalability limits that are determined internally by a device cannot be monitored by MAPS.

Default rules for scalability limit monitoring

The following table lists the scalability monitoring default rules in each of the default policies, and shows the actions and condition for each rule.

TABLE 34 Scalability monitoring default rules

Policy name	Rule name	Rule condition	Rule action
dflt_conservative_policy	defSWITCHL2_DEVCNT_PER_90	L2_DEVCNT_PER greater than 90	RASLOG, SNMP, EMAIL
	defSWITCHLSAN_DEVCNT_PER_90	LSAN_DEVCNT_PER greater than 90	
	defSWITCHZONE_CFGSZ_PER_90	ZONE_CFGSZ_PER greater than 90	
	defSWITCHBB_FCR_CNT_12	BB_FCR_CNT greater than 12	
dflt_moderate_policy	defSWITCHL2_DEVCNT_PER_75	L2_DEVCNT_PER greater than 75	RASLOG, SNMP, EMAIL
	defSWITCHLSAN_DEVCNT_PER_75	LSAN_DEVCNT_PER greater than 75	
	defSWITCHZONE_CFGSZ_PER_80	ZONE_CFGSZ_PER greater than 80	
	defSWITCHBB_FCR_CNT_12	BB_FCR_CNT greater than 12	
dflt_aggressive_policy	defSWITCHL2_DEVCNT_PER_60	L2_DEVCNT_PER greater than 60	RASLOG, SNMP, EMAIL
	defSWITCHLSAN_DEVCNT_PER_60	LSAN_DEVCNT_PER greater than 60	
	defSWITCHZONE_CFGSZ_PER_70	ZONE_CFGSZ_PER greater than 70	
	defSWITCHBB_FCR_CNT_12	BB_FCR_CNT greater than 12	

Examples of scalability limit rules

The following examples show the patterns for creating device counts, Fibre Channel router counts, and zone configuration usage rules for MAPS.

Rule for device counts in a Layer 2 fabric

In the following example, when the total device count in all switches that are part of the Layer 2 fabric rises above 90 percent of the total permissible count in the fabric, MAPS reports the threshold violation using a RASLog message.

```
switch123:FID128:admin> mapsrule --create L2_Dev_Count -group SWITCH -monitor L2_DEVCNT_PER -timebase none -
op ge -value 90 -action RASLOG -policy scalability_policy
```

Rule for LSAN device counts

In the following example, when the total device count in all switches that are part of the metaSAN (edge plus Backbone) fabric rises above 90 percent of the total permissible count in the fabric, MAPS reports the threshold violation using a RASLog message on that platform.

```
switch123:FID128:admin> mapsrule --create LSan_Dev_Count -group SWITCH -monitor LSAN_DEVCNT_PER -timebase
none -op ge -value 90 -action RASLOG -policy scalability_policy
```

Rule for Fibre Channel router count in Backbone fabric

In the following example, when the maximum limit of 12 Fibre Channel routers in the Backbone fabric is reached, MAPS reports the threshold violation using a RASLog message.

```
switch123:admin> mapsrule --create FCRCount -group SWITCH -monitor BB_FCR_CNT -timebase none -op ge -value
12 -action RASLOG -policy scalability_policy
```

Rule for zone configuration size

In the following example, when the zone configuration size limit reaches 90 percent of the total size, MAPS reports the threshold violation using a RASLog message.

```
switch123:admin> mapsrule --create ZoneConfigSize -group SWITCH -monitor ZONE_CFGSZ_PER -timebase none -op
ge -value 90 -action RASLOG -policy scalability_policy
```

For ZONE_CFGSZ_PER policy, the default time base is "none" and the system performs a daily check. The policy does not support other time base.

MAPS Service Availability Module

The MAPS Service Availability Module (MAPSSAM) reports display CPU, RAM, and flash memory usage and the port status for every physical and virtual Fibre Channel port on the switch.

There are three options for the **mapsSam** command:

- **--show**: Displays the MAPSSAM report. Additional option parameters are shown in the table below.
- **--clear**: Clears the SAM report.
- **--help**: Displays how to use the command.

Only **mapsSam --show** has additional option parameters. These parameters are listed in the following table and illustrated in the following examples.

TABLE 35 MAPSSAM --show command option parameters

Option	Details
--show (default option)	For each physical and virtual Fibre Channel port on a switch, this displays the total up, down, and offline time (as percentages), and the number of times the port has been down. This enables you to see if a particular port is failing more often than the others. NOTE The report does not distinguish why a port is recorded as down, it only reports how long the port has been down.
--show cpu	Displays the CPU usage as a percentage.
--show memory	Displays the general RAM memory usage as a percentage, along with total, used, and free memory values.
--show flash	Displays the flash memory usage as a percentage.

The following examples demonstrate using the various **mapsSam --show** option parameters.

Using only "--show"

When you enter simply **mapsSam --show**, the report lists the following information for each port:

- Port Number
- Port type
  - DIS (disabled port)
  - DIA (D_Port)
  - DP (persistently disabled port)
  - E (E_Port)
  - F (F_Port)
  - G (G_Port)
  - T (Trunk port)
  - TF (F_Port trunk)
  - U (U_Port)

    NOTE
    The MAPSSAM report does not include the health status of gigabyte Ethernet (GbE) ports.

- Total up time — Percentage of time the port was up.
- Total down time — Percentage of time the port was faulty.
- Down occurrence — Number of times the port was faulty.
- Total Offline time — Percentage of time the port was offline.
- Number of ports

All percentages are based on the total time the switch was up or down since the switch was rebooted, MAPS was activated, or the **mapsSam --clear** command was last run.

The following example shows typical output for **mapsSam --show**.

```
switch:admin> mapssam --show
                      Total        Total        Down        Total
Port      Type        Up Time      Down Time    Occurrence  Offline Time
                      (Percent)    (Percent)    (Times)     (Percent)
=============================================================================
0           F          100.00       0.00        0             0.00
1           F          100.00       0.00        0             0.00
2           U            0.00       0.00        0           100.00
3           F          100.00       0.00        0             0.00
4           DIS          0.00       0.00        0           100.00
5           DIS          0.00       0.00        0           100.00
6           DIS          0.00       0.00        0           100.00
7           DIS          0.00       0.00        0           100.00
Number of ports: 8
```

Using "--show cpu"

The following example shows the output for **mapsSam --show cpu**.

```
switch:admin> mapssam --show cpu memory
Showing Cpu Usage:
   CPU Usage   : 3.0%
```

Using "--show memory"

The following example shows the output for **mapsSam --show memory**.

```
switch:admin> mapssam --show memory
Showing Memory Usage:
   Memory Usage   : 22.0%
   Used Memory    : 225301k
   Free Memory    : 798795k
   Total Memory   : 1024096k
```

Using "--show flash"

The following example shows the output for **mapsSam --show flash**.

```
switch:admin> mapssam --show flash
Showing Flash Usage:
   Flash Usage   : 59%
```

MAPS monitoring for Extension platforms

FCIP Quality of Service (QoS) monitoring uses policies based on a combination of data characteristics and delivery requirements to appropriately prioritize data traffic. For example, while ordinary data traffic is tolerant of delays and dropped packets, real-time voice and video data are not. MAPS QoS policies provide a framework for accommodating these differences in traffic packets as they pass through a network, and can help you investigate issues such as packet loss, excessive bandwidth utilization, and similar issues.

MAPS can monitor the following on all FCIP-supported platforms:

- Tunnel
- Tunnel QoS
- Circuit
- Circuit QoS

QoS monitoring using MAPS uses the predefined QoS priorities of High, Medium, Low, and F-class. You can configure the values used by these priorities at both the FCIP tunnel and circuit level. The attributes monitored by MAPS for QoS at circuit level are throughput and packet loss. Throughput is the percentage of QoS circuit or tunnel utilization in a configured time period (hour, minute, or day); packet loss is the percentage of the total number of packets that have had to be retransmitted. MAPS monitors the state changes and throughput of each individual tunnel using these QoS priorities. QoS monitoring is not High Availability (HA)-capable.

For tunnel-level monitoring, MAPS can monitor the predefined groups ALL_TUNNELS, ALL_TUNNEL_HIGH_QOS, ALL_TUNNEL_MED_QOS, ALL_TUNNEL_LOW_QOS, and ALL_TUNNEL_F_QOS. These groups correspond to the FCIP tunnels.

For circuit-level monitoring, MAPS can monitor the predefined groups ALL_CIRCUIT_HIGH_QOS, ALL_CIRCUIT_MED_QOS, ALL_CIRCUIT_LOW_QOS, ALL_CIRCUIT_F_QOS, as well as the ALL_CIRCUITS group for round-trip time (RTT) and connection variance (Jitter) in addition to the CIR_STATE, CIR_UTIL, and CIR_PKTLOSS parameters. Monitoring the RTT and Jitter values helps to alert you to possible network disruptions and congestion in the network.

Brocade FCIP monitoring parameters and groups

The available statistics are broken out in the following table, and rules corresponding to these statistics are in the default policies.

TABLE 36 Use of Brocade FCIP monitoring groups as metrics

Parameter	Groups where the parameter is used as a metric
State change (STATE_CHG)	ALL_TUNNELS
Percent utilization (UTIL)	ALL_CIRCUIT_HIGH_QOS, ALL_CIRCUIT_MED_QOS, ALL_CIRCUIT_LOW_QOS, ALL_CIRCUIT_F_QOS, ALL_TUNNELS, ALL_TUNNEL_HIGH_QOS, ALL_TUNNEL_MED_QOS, ALL_TUNNEL_LOW_QOS, and ALL_TUNNEL_F_QOS
Percentage of packets lost in transmission (PKTLOSS)	ALL_CIRCUIT_HIGH_QOS, ALL_CIRCUIT_MED_QOS, ALL_CIRCUIT_LOW_QOS, ALL_CIRCUIT_F_QOS, ALL_TUNNEL_HIGH_QOS, ALL_TUNNEL_MED_QOS, ALL_TUNNEL_LOW_QOS , ALL_TUNNEL_F_QOS
Round-trip time in milliseconds (RTT)	ALL_CIRCUITS
Variance in RTT in milliseconds (Jitter)	ALL_CIRCUITS
CIR_STATE, CIR_UTIL, and CIR_PKTLOSS	Refer to FCIP Health on page 36 for descriptions of these parameters.

TABLE 37 FCIP monitoring parameters

Platform	Supported parameters
All FCIP platforms:	• CIR_UTIL • CIR_STATE • CIR_PKTLOSS • RTT • JITTER • PKTLOSS (Circuit QoS, Tunnel QoS) • UTIL (Tunnel, Circuit QoS, Tunnel QoS) • STATE_CHG (Tunnel)

Quality of Service monitoring example

The following MAPS rule states that when the packet loss percentage for ALL_CIRCUIT_HIGH_QOS group members becomes greater than or equal to 0.5 in a given minute, a RASLog entry will be posted.

```
switch:admin> mapsrule --create urule -group ALL_CIRCUIT_HIGH_QOS -monitor PKTLOSS -t min -op ge -value .5 -
action raslog
```

On triggering the rules, the corresponding RASLogs will appear under the summary section of the dashboard. In the following example, there is one RASLog, triggered by the rule "low_tunnel_mon". This rule has the format `-group ALL_TUNNEL_LOW_QOS -monitor PKTLOSS -timebase HOUR -op ge -value 30 -action raslogs`. The "Conditions contributing to health" column headings have been edited so as to allow the example to display clearly.

```
3.1 Summary Report:
===================
Category                  |Today                      |Last 7 days          |
--------------------------------------------------------------------------------
Port Health               |In operating range         |In operating range   |
BE Port Health            |No Errors                  |No Errors            |
Fru Health                |In operating range         |In operating range   |
Security Violations       |No Errors                  |In operating range   |
Fabric State Changes      |No Errors                  |No Errors            |
Switch Resource           |In operating range         |In operating range   |
Traffic Performance       |In operating range         |In operating range   |
FCIP Health               |In operating range         |In operating range   |
Fabric Performance Impact|In operating range         |In operating range   |

Conditions contributing to health:
==================================
Category(RuleCnt)|RptCnt|Rule Name                                  |Execution Time |Trigger Val(Units)|
-------------------------------------------------------------------------------------------------
FCIP Health (1)  |1     |defALL_CIRCUIT_HIGH_QOS_UTIL_75      |8/11/14 06:19:6|Circuit Qos 23/0  |
FCIP Health (1)  |1     |defALL_CIRCUIT_HIGH_QOS_PKTLOSS_PER_1|8/11/14 07:02:5|Circuit Qos 23/0  |
```

IPEXT monitoring

MAPS provide monitoring of IP Quality of Service (QoS) priorities and IP traffic over a tunnel or circuit along with FCIP monitoring. This feature is supported on Brocade Fabric OS IP extension platforms.

IPEXT provides Layer 3 (IP) extension for IP storage replication. QoS refers to policies for handling differences in data traffic based on data characteristics and delivery requirements. For example, ordinary data traffic is tolerant of delays and dropped packets, but real-time voice and video data are not tolerant. QoS policies provide a framework for accommodating the differences in data as it passes through a network. QoS for IP extension is provided through internal IP QoS priorities. Following are the advantages of IPEXT monitoring:

- Monitoring IP QoS parameters helps administrator to handle packet loss and high bandwidth utilization issues
- Monitoring overall IP traffic over a tunnel or circuit helps administrator effectively allocate bandwidth or resources configured on a IPEXT enabled VE tunnel

The following table explains the mapping between monitors and corresponding groups.

TABLE 38 Brocade IPEXT monitoring parameters and groups

Monitor	Logical groups
UTIL	ALL_TUNNEL_IP_HIGH_QOS
PKTLOSS	ALL_TUNNEL_IP_MED_QOS
	ALL_TUNNEL_IP_LOW_QOS
	ALL_CIRCUIT_IP_HIGH_QOS

TABLE 38 Brocade IPEXT monitoring parameters and groups (continued)

Monitor	Logical groups
	ALL_CIRCUIT_IP_MED_QOS ALL_CIRCUIT_IP_LOW_QOS
IP_UTIL	ALL_TUNNELS ALL_CIRCUITS
Percentage of packets lost in transmission (PKTLOSS)	ALL_CIRCUITS
Round-trip time in milliseconds (RTT)	ALL_CIRCUITS
Variance in RTT in milliseconds (Jitter)	ALL_CIRCUITS

IPEXT rule creation

The following example shows when circuit Qos utilization for ALL_CIRCUIT_IP_HIGH_QOS group members is greater than or equal to 5 in a given minute.

```
switch:admin>mapsrule --create urule -group
ALL_CIRCUIT_IP_HIGH_QOS -monitor UTIL -timebase
min -op ge -value 5 -action raslog
```

The following example shows when packet loss for ALL_TUNNEL_IP_HIGH_QOS group members is greater than or equal to 0.5 in a given minute.

```
switch:admin>mapsrule --create urule -group
ALL_TUNNEL_IP_HIGH_QOS -monitor PKTLOSS -timebase
min -op ge -value .5 -action raslog
```

Tunnel IP Qos monitors dashboard output sample

```
Category                |Today                    |Last 7 days              |
------------------------------------------------------------------------------
Port Health             |In operating range       |No Errors                |
BE Port Health          |No Errors                |No Errors                |
GE Port Health          |In operating range       |No Errors                |
Fru Health              |In operating range       |In operating range       |
Security Violations     |No Errors                |No Errors                |
Fabric State Changes    |No Errors                |No Errors                |
Switch Resource         |In operating range       |In operating range       |
Traffic Performance     |In operating range       |In operating range       |
FCIP Health             |Out of operating range   |No Errors                |
Fabric Performance Impact|In operating range      |In operating range       |
------------------------------------------------------------------------------

3.2 Rules Affecting Health:
===========================

Category    |Repeat|Rule Name              |Execution Time  |Object        |Triggered|
(Rule Count)|Count |                       |                |              |Value    |
-------------------------------------------------------------------------------------
FCIP Health |2     |ip_cir_high_qos_pktloss|08/10/1508:58:01|Circuit/Qos 24| 0 %     |
(8)         |      |                       |                |Circuit/Qos 24| 0 %     |
            |2     |ip_tnl_high_qos_pktloss|08/10/1508:58:01|Tunnel/Qos 24 | 0 %     |
            |      |                       |                |Tunnel/Qos 24 | 0 %     |
            |2     |ip_cir_high_qos_util   |08/10/1508:58:01|Circuit/Qos 24| 46.96 % |
            |      |                       |                |Circuit/Qos 24| 47.45 % |
            |2     |ip_tnl_high_qos_util   |08/10/1508:58:01|Tunnel/Qos 24 | 46.96 % |
            |      |                       |                |Tunnel/Qos 24 | 47.45 % |
-------------------------------------------------------------------------------------
```

Circuit IP monitors dashboard output sample

```
Category     |Repeat |Rule Name    |Execution Time    |Object         |Triggered      |
(Rule Count)|Count   |             |                  |               |Value(Units)   |
-------------------------------------------------------------------------------------
FCIP Health |8       |ipcktjitter |08/10/15 09:11:01|Port/Cir 24/0 |0 %            |
(34)         |        |             |                  |Port/Cir 24/0 |0 %            |
             |        |             |                  |Port/Cir 24/0 |0 %            |
             |        |             |                  |Port/Cir 24/0 |0 %            |
             |        |             |                  |Port/Cir 24/0 |0 %            |
             |8       |ipcktrtt     |08/10/15 09:11:01|Port/Cir 24/0 |1 Milliseconds|
             |        |             |                  |Port/Cir 24/0 |1 Milliseconds|
             |        |             |                  |Port/Cir 24/0 |1 Milliseconds|
             |        |             |                  |Port/Cir 24/0 |1 Milliseconds|
             |        |             |                  |Port/Cir 24/0 |1 Milliseconds|
             |9       |ipcktutil    |08/10/15 09:11:30|Port/Cir 24/0 |47.44 %        |
             |        |             |                  |Port/Cir 24/0 |47.43 %        |
             |        |             |                  |Port/Cir 24/0 |46.97 %        |
             |        |             |                  |Port/Cir 24/0 |47.43 %        |
             |        |             |                  |Port/Cir 24/0 |47.18 %        |
             |9       |ipcktpkt     |08/10/15 09:11:30|Port/Cir 24/0 |0 %            |
             |        |             |                  |Port/Cir 24/0 |0 %            |
             |        |             |                  |Port/Cir 24/0 |0 %            |
             |        |             |                  |Port/Cir 24/0 |0 %            |
             |        |             |                  |Port/Cir 24/0 |0 %            |
-------------------------------------------------------------------------------------
```

Tunnel IP monitors dashboard output sample

```
Category       |Repeat |Rule Name    |Execution Time    |Object         |Triggered    |
(Rule Count)   |Count  |             |                  |               |Value(Units)|
-------------------------------------------------------------------------------------
FCIP Health(1)|1       |iptnlutil    |08/10/15 09:15:36|Tunnel 24      |47.16 %      |
-------------------------------------------------------------------------------------
```

E-mail delivery monitoring

MAPS allows monitoring of e-mails that are sent out of the system, and it sends a notification if an e-mail is not delivered.

There is only one e-mail failure RASLog generated within an hour for every e-mail address configured. The RASlog contains the e-mail address of the failed destination. There is another RASLog which is sent every hour which indicates the total number of failures in the one-hour period. Sending only one RASLog avoids generating too many RASLog failure messages in the event of e-mail server problem or network issues.

An e-mail can be sent from a system, but it might not be successfully delivered to the destination e-mail server. In MAPS, e-mails are sent using the **sendmail** command in Linux. The initial steps to debug the problem are to check for the following:

- whether the network is running
- whether the routing tables are configured correctly
- whether the DNS is configured correctly

Even when the network connectivity is working properly, the **sendmail** command can have the following failure conditions:

- There are no routes in the system to use to send an e-mail. For example, there is a network error and TCP/IP service is not available.
- A fatal configuration problem occurs while reading the configuration file. Failure of a delivery agent to function correctly can lead to this kind of failure.
- Problems can result from various operating system errors.

- The **sendmail** process cannot be forked by the **cron** job script due to a memory problem, and therefore, the command cannot be successfully executed.
- Other internal software errors.

Examples of RASLog messages

The following is an example of the RASLog which is generated when the **sendmail** command fails to send an e-mail from the switch.

```
2016/02/01-19:49:00, [MAPS-1206], 1468, SLOT 6 CHASSIS, INFO, dcx_178, A MAPS notification sent
from the switch to abc@brocade.com could not be delivered to the mail server.
```

The following is an example of the RASLog which is generated every hour to indicate the number of failures seen in one hour. This RASLog is generated only if there is more than one failure for any failed address during the one-hour period.

```
2016/02/02-20:49:00, [MAPS-1206], 1468, SLOT 6 CHASSIS, INFO, dcx_178, There were 12 or more
notifications that could not be delivered in the last one hour.
```

Fan air-flow direction monitoring

Fabric OS allows you to monitor air-flow direction of the fans. If two fans are running in opposite directions, then the switch is marked as marginal.

MAPS monitors the system and if a switch has fans that are running in opposite directions (mixed mode), then it changes the state of the switch to "marginal" and sends an alert. You can check the state of the switch in the dashboard. To support this feature, the following rule has been added as part of all moderate, conservative, aggressive, and base policies.

```
Rule name              |Condition                                 |Actions               |
-------------------------------------------------------------------------------------------
defALL_FAN_AIR_FLOW    |                                          |                      |
_MISMATCH              |CHASSIS(FAN_AIRFLOW_MISMATCH/NONE==TRUE)  |SW_MARGINAL,SNMP,EMAIL|
-------------------------------------------------------------------------------------------|
```

Dashboard output example for fan air-flow direction monitoring

The following is an example of dashboard output when the fan air-flow direction rule has been triggered:

```
From MAPS FS 3-30-16
mapspolicy -show test_1
Policy Name: test_1

Rule Name             |Condition                                  |Actions                 |
-----------------------------------------------------------------------------------------
test_rule_air_flow_60 | chassis(FAN_AIRFLOW_MISMATCH/none==TRUE) |raslog,sw_marginal      |

mapsdb --show

1 Dashboard Information:
========================

DB start time:                  Tue Jan 12 17:39:09 2016
Active policy:                  test_1
Configured Notifications:       RASLOG,SW_CRITICAL,SW_MARGINAL
Fenced Ports :                  None
Decommissioned Ports :          None
Fenced circuits :               None
Quarantined Ports :             None

2 Switch Health Report:
========================

Current Switch Policy Status: MARGINAL
Contributing Factors:
---------------------
*FAN_AIRFLOW_MISMATCH (MARGINAL).

3.1 Summary Report:
===================

Category                 |Today                      |Last 7 days               |
-------------------------------------------------------------------------------
Port Health              |No Errors                  |No Errors                 |
BE Port Health           |No Errors                  |No Errors                 |
GE Port Health           |No Errors                  |No Errors                 |
Fru Health               |In operating range         |In operating range        |
Security Violations      |No Errors                  |No Errors                 |
Fabric State Changes     |No Errors                  |No Errors                 |
Switch Resource          |Out of operating range     |In operating range        |
Traffic Performance      |In operating range         |In operating range        |
FCIP Health              |No Errors                  |No Errors                 |
Fabric Performance Impact|In operating range         |In operating range        |


3.2 Rules Affecting Health:
===========================

Category(Rule Count)|RepeatCount|Rule Name             |Execution Time   |Object    |Triggered Value(Units)|
-----------------------------------------------------------------------------------------------------------
Switch Resource (2) |2          |test_rule_air_flow_60 |01/12/16 17:42:54|Chassis   |TRUE                  |
```

RASLog example for fan air-flow direction monitoring

The following is an example of a RASLog message sent when MAPS detects the condition when there is a mismatch in the air flow direction on switch.

```
Example comes from FS 3-30-16
2016/01/12-17:39:30, [MAPS-1003], 1507, FID 128, WARNING, sw128_top, Chassis,
Condition=CHASSIS(FAN_AIRFLOW_MISMATCH==TRUE),
Current Value:[ FAN_AIRFLOW_MISMATCH,TRUE], RuleName=test_rule_air_flow_1,
Dashboard Category=Switch Resource.

2016/01/12-17:39:30, [MAPS-1021], 1508, FID 128, WARNING, sw128_top,
RuleName=test_rule_air_flow_1, Condition=CHASSIS(FAN_AIRFLOW_MISMATCH==TRUE),
Obj:Chassis [ FAN_AIRFLOW_MISMATCH,TRUE] has contributed to switch status MARGINAL.

2016/01/12-17:39:30, [MAPS-1020], 1509, FID 128, WARNING, sw128_top,
Switch wide status has changed from HEALTHY to MARGINAL.
```

Updating monitoring policies for devices with four PSUs

If you have a chassis that supports more than two power supply units (PSUs), when you increase the number of PSUs and want to enable the Call Home feature to be activated for all the PSUs, you must change the active MAPS power supply "switchstatus" policy settings. This capability requires that you have a Brocade Direct Support maintenance contract.

NOTE

This procedure applies only to the following switches:

- Brocade DCX 8510-8
- Brocade X6-8 Director

To change the active MAPS power supply "switchstatus" policy settings, complete the following steps.

1. Display the MAPS policy rules using one of the following commands:
   - **mapspolicy --show -all**
   - **mapsrule --show fw_CHASSISBAD_PWRCrit_3**

   The following examples show the results of each of these commands when the policy is set for two PSUs.

   ```
   switch:admin> mapspolicy --show -all
   fw_CHASSISBAD_PWRCrit_3 SW_CRITICAL
   CHASSIS(BAD_PWR/none>=3)
   .
   .
   .
   Active Policy is 'fw_active_policy'.

   switch:admin> mapsrule --show fw_CHASSISBAD_PWRCrit_3
   Rule Data:
    ----------
   RuleName: fw_CHASSISBAD_PWRCrit_3
   Condition: CHASSIS(BAD_PWR/none>=3)
   Actions: SW_CRITICAL
   Associated Policies: fw_active_policy
   ```
2. Verify that the chassis has the condition BAD_PWR/none>=3 (**CHASSIS(BAD_PWR/none>=3)**).
3. Record the Active Policy name. This name is required to complete the following step.
4. Use the **mapsrule --config fw_CHASSISBAD_PWRCrit_3 -policy** *policy_name from Step 3* **-monitor BAD_PWR -group CHASSIS -timebase none -op ge -value 1 -action SW_CRITICAL** command to change the CHASSIS(BAD_PWR/none>=3)

setting to CHASSIS(BAD_PWR/none>=1) . When you change a policy, you must enter all the values for the policy, even if you are changing only one value.
In this example, the **-policy** name is **fw_active_policy** which you noted earlier in step 3.

```
switch:admin> mapsrule --config fw_CHASSISBAD_PWRCrit_3 -policy fw_active_policy -monitor BAD_PWR -
group CHASSIS -timebase none -op ge -value 1 -action SW_CRITICAL
Associated Policies: fw_active_policy
```

5. Enter **mapsrule --show fw_CHASSISBAD_PWRCrit_3** to verify the value is set to 1.

```
switch:admin> mapsrule --show fw_CHASSISBAD_PWRCrit_3
Rule Data:
 ----------
RuleName: fw_CHASSISBAD_PWRCrit_3
Condition: CHASSIS(BAD_PWR/none>=1)   <- Note the changed value.
Actions: SW_CRITICAL
Associated Policies: fw_active_policy
```

If you have a Brocade Direct Support maintenance contract, all four PSUs can now use the Call Home event notification capability to automatically send an e-mail or dial into a support center to report system problems.

MAPS Threshold Values

- Viewing monitoring thresholds....147
- Back-end port monitoring thresholds....148
- Fabric state change monitoring thresholds....148
- Extension monitoring thresholds....149
- FRU state monitoring thresholds....149
- Port Health monitoring thresholds....150
- Resource monitoring thresholds....153
- Security monitoring thresholds....153
- SFP monitoring thresholds....154
- Fabric Performance Impact thresholds....155
- Switch status policy monitoring thresholds....156
- Traffic Performance thresholds....158

Viewing monitoring thresholds

You can use the CLI to view the thresholds for a policy, or for a group within a policy.

To view monitoring thresholds, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **mapspolicy --show** *policy_name*. To see only the thresholds for a specific group in a policy, use **--show** *policy_name* | **grep** *group_name*.

   The following example shows all the thresholds for the ALL_D_PORTS group in the policy named "dflt_conservative_policy".

```
switch:admin> mapspolicy --show dflt_conservative_policy | grep ALL_D_PORTS
defALL_D_PORTSCRC_3             RASLOG,SNMP,EMAIL   ALL_D_PORTS(CRC/MIN>3)
defALL_D_PORTSPE_3              RASLOG,SNMP,EMAIL   ALL_D_PORTS(PE/MIN>3)
defALL_D_PORTSITW_3             RASLOG,SNMP,EMAIL   ALL_D_PORTS(ITW/MIN>3)
defALL_D_PORTSLF_3              RASLOG,SNMP,EMAIL   ALL_D_PORTS(LF/MIN>3)
defALL_D_PORTSLOSS_SYNC_3       RASLOG,SNMP,EMAIL   ALL_D_PORTS(LOSS_SYNC/MIN>3)
defALL_D_PORTSCRC_H90           RASLOG,SNMP,EMAIL   ALL_D_PORTS(CRC/HOUR>90)
defALL_D_PORTSPE_H90            RASLOG,SNMP,EMAIL   ALL_D_PORTS(PE/HOUR>90)
defALL_D_PORTSITW_H90           RASLOG,SNMP,EMAIL   ALL_D_PORTS(ITW/HOUR>90)
defALL_D_PORTSLF_H90            RASLOG,SNMP,EMAIL   ALL_D_PORTS(LF/HOUR>90)
defALL_D_PORTSLOSS_SYNC_H90     RASLOG,SNMP,EMAIL   ALL_D_PORTS(LOSS_SYNC/HOUR>90)
defALL_D_PORTSCRC_D1500         RASLOG,SNMP,EMAIL   ALL_D_PORTS(CRC/DAY>1500)
defALL_D_PORTSPE_D1500          RASLOG,SNMP,EMAIL   ALL_D_PORTS(PE/DAY>1500)
defALL_D_PORTSITW_D1500         RASLOG,SNMP,EMAIL   ALL_D_PORTS(ITW/DAY>1500)
defALL_D_PORTSLF_D1500          RASLOG,SNMP,EMAIL   ALL_D_PORTS(LF/DAY>1500)
defALL_D_PORTSLOSS_SYNC_D1500   RASLOG,SNMP,EMAIL   ALL_D_PORTS(LOSS_SYNC/DAY>1500)
```

   The first column is the name of the statistic being monitored. The second is the actions for that statistic that will be triggered if the threshold is passed. The third column lists the group being monitored, followed by the metric, followed by the threshold. This means that "defALL_D_PORTSCRC_3 RASLOG,SNMP,EMAIL ALL_D_PORTS(CRC/MIN>3)" :

   - Is named "defALL_D_PORTSCRC_3"
   - Has the actions RASLog, SNMP, and e-mail
   - Applies to all D_Ports
   - Measures CRC errors per minute. The threshold to trigger the listed actions is "more than three errors in a minute".

Back-end port monitoring thresholds

All Back-end port monitors support the Minute, Hour, Day, and Week timebases.

The following table lists the errors MAPS monitors for on back-end ports, the trigger thresholds, and the default actions to be taken when the threshold is crossed.

TABLE 39 Back-end port monitoring default thresholds and actions

Errors	Thresholds	Actions
CRC	• 10 per 5 minutes • 100 per day	RASLOG, SNMP, EMAIL, FMS
Link Reset	• 10 per 5 minutes • 100 per day	RASLOG, SNMP, EMAIL, FMS
ITW	• 10 per 5 minutes • 100 per day	RASLOG, SNMP, EMAIL, FMS
BAD_OS	• 10 per 5 minutes • 100 per day	RASLOG, SNMP, EMAIL, FMS
Frame too long	• 10 per 5 minutes • 100 per day	RASLOG, SNMP, EMAIL, FMS
Frame truncated	• 10 per 5 minutes • 100 per day	RASLOG, SNMP, EMAIL, FMS

Fabric state change monitoring thresholds

All the fabric state change monitors support the Minute, Hour, and Day timebases. They do not support the “None” timebase.

The following table lists the default monitoring thresholds for fabric state change criteria used by MAPS. All thresholds are measured per minute and actions are triggered when the reported value is greater than the threshold value.

TABLE 40 Default fabric state change monitoring thresholds and actions

	Fabric state change monitoring threshold values by policy			
Monitoring statistic	Aggressive	Moderate	Conservative	Actions
Domain ID change	1	1	1	RASLOG, SNMP, EMAIL
Fabric logins	4	6	8	RASLOG, SNMP, EMAIL
Fabric reconfigurations	1	2	4	RASLOG, SNMP, EMAIL
E_Ports down	1	2	4	RASLOG, SNMP, EMAIL
Segmentation changes	1	2	4	RASLOG, SNMP, EMAIL
Zone changes	2	5	10	RASLOG, SNMP, EMAIL
L2 Device Count	60	75	90	RASLOG, SNMP, EMAIL
LSAN Device Count	60	75	90	RASLOG, SNMP, EMAIL
Zone Configuration size	70	80	90	RASLOG, SNMP, EMAIL
FCR Count	12	12	12	RASLOG, SNMP, EMAIL

Extension monitoring thresholds

These FCIP monitors support Minute, Hour, Day, and Week timebases: Tunnel state change, Tunnel throughput, Tunnel QoS throughput, Tunnel QoS Packet loss, FCIP Circuit State Changes, FCIP Circuit Utilization, FCIP Packet loss, FCIP Circuit Round Trip Time, and FCIP connection variance.

The following tables list the default monitoring thresholds for Fiber Channel over IP (FCIP) criteria used by MAPS. All actions are triggered when the reported value is greater than the threshold value.

TABLE 41 Default FCIP monitoring thresholds and actions

Monitoring statistic	Units	FCIP monitoring threshold values by policy			Actions
		Aggressive	Moderate	Conservative	
Circuit state change (CIR_STATE)	Changes per minute	0	3	5	RASLOG, SNMP, EMAIL, FENCE
Circuit utilization percentage (CIR_UTIL)	Percentage per hour	50	75	90	RASLOG, SNMP, EMAIL
Circuit packet loss percentage (CIR_PKTLOSS)	Percentage per minute	0.01	0.05	0.1	RASLOG, SNMP, EMAIL
Circuit round-trip times (RTT)	Total delay in milliseconds	250	250	250	RASLOG, SNMP, EMAIL
Circuit jitter (JITTER)	Percentage of delay, calculated from the difference of two successive minutes. The total delay in milliseconds is averaged per minute for each minute. Values less than 5 ms in the converted percentage are ignored.	5	15	20	RASLOG, SNMP, EMAIL
Tunnel (STATE_CHG)	Changes per minute	0	1	3	RASLOG, SNMP, EMAIL
Tunnel or circuit QoS (UTIL)	Percentage per hour	60	75	90	RASLOG, SNMP, EMAIL
QoS Packet loss percentage (PKTLOSS)	Percentage per minute	0.01	0.05	0.1	RASLOG, SNMP, EMAIL

FRU state monitoring thresholds

For all FRU monitoring statistics, the default MAPS thresholds are part of blade and WWN rules for Brocade DCX and Brocade DCX+ systems. All threshold conditions are absolute, and actions are triggered when the statistic value either does or does not match the value (depending on how the rule is written). FRU monitoring statistics do not use any timebases.

TABLE 42 FRU monitoring statistics, states, and actions

Monitored Statistic	Supported States	Actions
Power Supply (PS_STATE)	ON, OUT, FAULTY	RASLOG, SNMP, EMAIL
Fan (FAN_STATE)	ON, OUT, FAULTY	RASLOG, SNMP, EMAIL
Slot (BLADE_STATE)	ON, OFF, OUT, FAULTY	RASLOG, SNMP, EMAIL
SFP (SFP_STATE)	IN, OUT, FAULTY	RASLOG, SNMP, EMAIL
WWN (WWN)	ON, OUT, FAULTY	RASLOG, SNMP, EMAIL

Port Health monitoring thresholds

All Port Health monitoring thresholds used by MAPS are triggered when they exceed the listed value. For thresholds that have both an upper value and a lower value, the threshold is triggered when it exceeds the upper value or drops below the lower value. All thresholds other than RXP, TXP, and Utilization percentage are measured per minute. The RXP, TXP, and Utilization percentage thresholds are measured per hour.

The following Port Health monitors support Minute, Hour, and Day timebases: CRC Errors, Invalid Transmit Words, Loss of sync, Link Failure, Loss of Signal, Protocol Errors, Link Reset, C3 Time outs, and State change. The SFP Current, SFP Receive Power, SFP Transmit Power, SFP Voltage, SFP Temperature, and SFP Power On Hours monitors support only the "None" timebase.

D_Port default Port Health monitoring thresholds

The following tables list the default D_Port Port Health monitoring threshold values and actions, broken out by policy.

TABLE 43 Aggressive policy default D_Port Port Health monitoring threshold values and actions

Monitoring statistic	Unit	Threshold	Actions
CRC Errors (defALL_D_PORTSCRC_1)	Min	1	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_1)	Min	1	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_1)	Min	1	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_1)	Min	1	EMAIL, SNMP, RASLOG
CRC Errors (defALL_D_PORTSCRC_H30)	Hour	30	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_H30)	Hour	30	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_H30)	Hour	30	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_H30)	Hour	30	EMAIL, SNMP, RASLOG
CRC Errors (defALL_D_PORTSCRC_D500)	Day	500	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_D500)	Day	500	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_D500)	Day	500	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_D500)	Day	500	EMAIL, SNMP, RASLOG

TABLE 44 Moderate policy default D_Port Port Health monitoring threshold values and actions

Monitoring statistic	Unit	Threshold	Actions
CRC Errors (defALL_D_PORTSCRC_2)	Min	2	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_2)	Min	2	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_2)	Min	2	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_2)	Min	2	EMAIL, SNMP, RASLOG
CRC Errors (defALL_D_PORTSCRC_H60)	Hour	60	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_H60)	Hour	60	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_H60)	Hour	60	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_H60)	Hour	60	EMAIL, SNMP, RASLOG
CRC Errors (defALL_D_PORTSCRC_D1000)	Day	1000	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_D1000)	Day	1000	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_D1000)	Day	1000	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_D1000)	Day	1000	EMAIL, SNMP, RASLOG

TABLE 45 Conservative policy default D_Port Port Health monitoring threshold values and actions

Monitoring statistic	Unit	Threshold	Actions
CRC Errors (defALL_D_PORTSCRC_3)	Min	3	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_3)	Min	3	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_3)	Min	3	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_3)	Min	3	EMAIL, SNMP, RASLOG
CRC Errors (defALL_D_PORTSCRC_H90)	Hour	90	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_H90)	Hour	90	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_H90)	Hour	90	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_H90)	Hour	90	EMAIL, SNMP, RASLOG
CRC Errors (defALL_D_PORTSCRC_D1500)	Day	1500	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_D1500)	Day	1500	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_D1500)	Day	1500	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_D1500)	Day	1500	EMAIL, SNMP, RASLOG

E_Port default Port Health monitoring thresholds

The following table lists the default E_Port Port Health monitoring threshold values and actions.

TABLE 46 Default E_Port Port Health monitoring threshold values and actions

Monitoring statistic	E_Port monitoring threshold values by policy			Actions
	Aggressive	Moderate	Conservative	
C3 Time out (C3TX_TO)	5	10	20	EMAIL, SNMP, RASLOG
CRC Errors (CRC)	Low: 0 High: 2	Low: 10 High: 20	Low: 21 High: 40	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Invalid Transmit Words (ITW)	Low: 15 High: 20	Low: 21 High: 40	Low: 41 High: 80	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Link Reset (LR)	Low: 2 High: 4	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
State Change (STATE_CHG)	Low: 2 High: 4	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Loss of signal (LOSS_SIGNAL)	0	3	5	EMAIL, SNMP, RASLOG
Link Failure (LF)	0	3	5	EMAIL, SNMP, RASLOG
Sync Loss (LOSS_SYNC)	0	3	5	EMAIL, SNMP, RASLOG

F_Port default Port Health monitoring thresholds

The following table lists the default Host F_Port Port Health monitoring threshold values and actions.

TABLE 47 Default Host F_Port Port Health monitoring threshold values and actions

Monitoring statistic	Host F_Port monitoring threshold values by policy			Actions
	Aggressive	Moderate	Conservative	
C3 Time out (C3TX_TO)	Low: 2 High: 4	Low: 3 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLOG, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS

TABLE 47 Default Host F_Port Port Health monitoring threshold values and actions (continued)

Monitoring statistic	Host F_Port monitoring threshold values by policy			Actions
	Aggressive	Moderate	Conservative	
CRC Errors (CRC)	Low: 0 High: 2	Low: 10 High: 20	Low: 21 High: 40	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Invalid Transmit Words (ITW)	Low: 15 High: 20	Low: 21 High: 40	Low: 41 High: 80	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Link Reset (LR)	Low: 2 High: 4	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
State Change (STATE_CHG)	Low: 2 High: 4	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Loss of signal (LOSS_SIGNAL)	0	3	5	EMAIL, SNMP, RASLOG
Link Failure (LF)	0	3	5	EMAIL, SNMP, RASLOG
Sync Loss (LOSS_SYNC)	0	3	5	EMAIL, SNMP, RASLOG

The following table lists the default Target F_Port Port Health monitoring threshold values and actions.

TABLE 48 Default Target F_Port Port Health monitoring threshold values and actions

Monitoring statistic	Target F_Port monitoring threshold values by policy			Actions
	Aggressive	Moderate	Conservative	
C3 Time out (C3TX_TO)	Low: 0 High: 2	Low: 3 High: 5	Low: 6 High: 10	Low: EMAIL, SNMP, RASLOG, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS
CRC Errors (CRC)	Low: 0 High: 2	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Invalid Transmit Words (ITW)	Low: 5 High: 10	Low: 11 High: 20	Low: 21 High: 40	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Link Reset (LR)	Low: 0 High: 2	Low: 3 High: 5	Low: 6 High: 10	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
State Change (STATE_CHG)	Low: 0 High: 2	Low: 3 High: 7	Low: 8 High: 15	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Loss of signal (LOSS_SIGNAL)	0	3	5	EMAIL, SNMP, RASLOG
Link Failure (LF)	0	3	5	EMAIL, SNMP, RASLOG
Sync Loss (LOSS_SYNC)	0	3	5	EMAIL, SNMP, RASLOG

NOTE

If an F_Port cannot be identified as either a host or a target, the thresholds for it are the same as those for Host F_Ports.

Non-F_Port default Port Health monitoring thresholds

The following table lists the default non-F_Port Port Health monitoring threshold values and actions.

TABLE 49 Default non-F_Port Port Health monitoring threshold values and actions

Monitoring statistic	Non-F_Port monitoring threshold values by policy			Actions
	Aggressive	Moderate	Conservative	
C3 Time out (C3TX_TO)	N/A	N/A	N/A	N/A
CRC Errors (CRC)	Low: 0 High: 2	Low: 10 High: 20	Low: 21 High: 40	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Invalid Transmit Words (ITW)	Low: 15 High: 20	Low: 21 High: 40	Low: 41 High: 80	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Link Reset (LR)	Low: 2 High: 4	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
State Change (STATE_CHG)	Low: 2 High: 4	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Loss of signal (LOSS_SIGNAL)	0	3	5	EMAIL, SNMP, RASLOG
Link Failure (LF)	0	3	5	EMAIL, SNMP, RASLOG
Sync Loss (LOSS_SYNC)	0	3	5	EMAIL, SNMP, RASLOG

Resource monitoring thresholds

The only timebase the Resource monitors support is “None”.

The following table lists the default monitoring threshold values and associated actions for the switch resource criteria monitored by MAPS. All thresholds are measured per minute and are triggered when they are greater than the shown value.

TABLE 50 Default resource monitoring thresholds and actions

Monitoring statistic	Resource monitoring threshold values by policy			Actions
	Aggressive	Moderate	Conservative	
Flash memory percentage used (FLASH_USAGE)	90	90	90	RASLOG, SNMP, EMAIL
CPU percentage used (CPU)	80	80	80	RASLOG, SNMP, EMAIL
Memory percentage used (MEMORY_USAGE)	75	75	75	RASLOG, SNMP, EMAIL
Ethernet management port state (ETH_MGMT_PORT_STATE)	Up/Down	Up/Down	Up/Down	RASLOG, SNMP, EMAIL
Temperature Sensor (TEMP)	OUT_OF_RANGE	OUT_OF_RANGE	OUT_OF_RANGE	RASLOG, SNMP, EMAIL

Security monitoring thresholds

All the Security Health monitors support the Minute, Hour, and Day timebases. They do not support the “None” timebase.

The following table lists the default monitoring thresholds for security criteria used by MAPS. Unless noted otherwise, all thresholds are measured per minute and actions are triggered when the reported value is greater than the threshold value.

TABLE 51 Default security monitoring thresholds and actions

Monitoring statistic	Security monitoring threshold values by policy			Actions
	Aggressive	Moderate	Conservative	
DCC violations	0	2	4	RASLOG, SNMP, EMAIL
HTTP violation	0	2	4	RASLOG, SNMP, EMAIL
Illegal command	0	2	4	RASLOG, SNMP, EMAIL
Incompatible security DB	0	2	4	RASLOG, SNMP, EMAIL
Login violations	0	2	4	RASLOG, SNMP, EMAIL
Invalid certifications	0	2	4	RASLOG, SNMP, EMAIL
No-FCS	0	2	4	RASLOG, SNMP, EMAIL
SCC violations	0	2	4	RASLOG, SNMP, EMAIL
SLAP failures	0	2	4	RASLOG, SNMP, EMAIL
Telnet violations	0	2	4	RASLOG, SNMP, EMAIL
TS out of sync	1 per hour 2 per day	2 per hour 4 per day	4 per hour 10 per day	RASLOG, SNMP, EMAIL

SFP monitoring thresholds

SFP monitoring statistics do not use any timebases.

All SFP monitoring thresholds used by MAPS are triggered when the reported value exceeds the threshold value. For thresholds with both an upper value and a lower value, actions are triggered when the reported value exceeds the upper threshold value or drops below the lower threshold value.

10 Gbps, 16 Gbps, and 32 Gbps SFP monitoring threshold defaults

The following table lists the default thresholds for 10 Gbps, 16 Gbps, and 32 Gbps SFPs.

TABLE 52 Default SFP monitoring thresholds and actions for 10 Gbps, 16 Gbps, and 32 Gbps SFPs

Monitoring statistic	SFP monitoring thresholds for all policies							Actions
	10 Gbps SWL	10 Gbps LWL	16 Gbps SWL	16 Gbps LWL	25Km 16 Gbps LWL	32 Gbps SWL	32 Gbps LWL	
Current (CURRENT) (mA)	10	95	12	70	90	12	60	SFP_MARGINAL, RASLOG, SNMP, EMAIL
Receive Power (RXP) (µW)	1999	2230	1259	1995	2338	1259	1995	SFP_MARGINAL, RASLOG, SNMP, EMAIL
Temperature (SFP_TEMP) (°C)	-5 to 90	-5 to 90	-5 to 85	-5 to 90	-5 to 75	-5 to 85	-5 to 75	SFP_MARGINAL, RASLOG, SNMP, EMAIL
Transmit Power (TXP) (µW)	1999	2230	1259	1995	4466	1259	1584	SFP_MARGINAL, RASLOG, SNMP, EMAIL
Voltage (VOLTAGE) (mV)	3000 to 3600	2970 to 3600	3000 to 3600	3000 to 3600	2850 to 3750	3000 to 3600	3000 to 3600	SFP_MARGINAL, RASLOG, SNMP, EMAIL

Quad SFPs and all other SFP monitoring threshold defaults

The following table lists the default threshold and actions for Quad SFPs (QSFPs) and all other SFPs.

TABLE 53 Default SFP monitoring thresholds and actions for QSFPs and all other SFPs

Monitoring statistic	QSFP and other SFP monitoring thresholds for all policies					Actions
	100M 16Gbps SWL QSFP	2K QSFP	32 Gbps SWL QSFP	QSFP	All Other SFPs	
Current (CURRENT) (mA)	10	39	10	10	50	RASLOG, SNMP, EMAIL
Receive Power (RXP) (µW)	2187	2000	3400	2180	5000	RASLOG, SNMP, EMAIL
Temperature (TEMP) (°C)	-5 to 85	-15 to 85	-5 to 75	-5 to 85	-13 to 85	RASLOG, SNMP, EMAIL
Transmit Power (TXP) (µW)	—	—	—	—	5000	RASLOG, SNMP, EMAIL
Voltage (VOLTAGE) (mV)	2970 to 3630	2900 to 3600	2970 to 3630	2940 to 3600	2960 to 3630	RASLOG, SNMP, EMAIL

Fabric Performance Impact thresholds

The following FPI monitors support the Minute, Hour, and Day timebases: Receive Bandwidth usage percentage, Transmit Bandwidth usage percentage, and Trunk Utilization percentage. The Fabric Performance Impact monitor (DEV_LATENCY_IMPACT) supports only the "None" timebase.

The following table lists the default latency threshold values for FPI monitoring. They are binary, in that the threshold value is either present or it is not. When the latency returns within the threshold values, another message is issued, IO_LATENCY_CLEAR.

TABLE 54 Default Fabric Performance Impact latency monitoring threshold values and actions

Monitoring statistic	Value (Y/N) for all policies	Actions
Fabric Performance Impact (DEV_LATENCY_IMPACT)	IO_FRAME_LOSS	RASLOG, SNMP, EMAIL, SDDQ, TOGGLE
	IO_PERF_IMPACT	RASLOG, SNMP, EMAIL, SDDQ, TOGGLE
	IO_LATENCY_CLEAR	RASLOG, SNMP, EMAIL

The following table lists the default Receive Bandwidth usage percentage, Transmit Bandwidth usage percentage, and Trunk Utilization percentage value monitoring thresholds for E_Ports, Host F_Ports, Target F_Ports, and non-F_Ports.

TABLE 55 Default Fabric Performance Impact RX, TX, and UTIL monitoring threshold values and actions

Monitoring statistic	FPI monitoring threshold values by policy			Actions
	Aggressive	Moderate	Conservative	
Receive Bandwidth usage percentage (RX)	60	75	90	EMAIL, SNMP, RASLOG
Transmit Bandwidth usage percentage (TX)	60	75	90	EMAIL, SNMP, RASLOG
Trunk Utilization percentage (UTIL)	60	75	90	EMAIL, SNMP, RASLOG

Switch status policy monitoring thresholds

The only timebase the Switch status monitors support is the "None" timebase. The following tables list the default switch status policy monitoring thresholds used by MAPS. All threshold conditions are absolute and actions are triggered when the reported value is greater than or equal to the threshold value. For thresholds with both an upper value and a lower value, an action is triggered when the reported value exceeds the upper threshold value or drops below the lower threshold value. The following tables list the default Switch Status monitoring threshold values and actions, broken out by policy.

The following table lists the default Switch Status monitoring thresholds and actions for the default aggressive policy.

TABLE 56 Aggressive policy default Switch Status monitoring thresholds and actions

Monitoring statistic	Switch status threshold values (Marginal/ Critical)	Actions
Absent or faulty power supply (BAD_PWR)	DCX, DCX+: -/3 All other supported platforms: 1/2	SW_CRITICAL, SNMP, EMAIL
Temperature sensors outside range (BAD_TEMP)	1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
Absent or faulty fans (BAD_FAN)	1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
Flash usage (FLASH_USAGE)	90	RASLOG, SNMP, EMAIL
Percentage of marginal ports (MARG_PORTS)	-/5	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Percentage of error ports (ERR_PORTS)	-/5	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Percentage of faulty ports (FAULTY_PORTS)	-/5	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Faulty blades (FAULTY_BLADE)	DCX, DCX+: 1/-	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: No Action
Faulty WWN (WWN_DOWN)	DCX, DCX+: -/1	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Core blade monitoring (DOWN_CORE)	DCX, DCX+: 1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
HA Sync (HA_SYNC)	DCX, DCX+,: sync=0	SW_MARGINAL, SNMP, EMAIL
Missing SFP (MISSING_SFP)	??	??
Mismatched fan airflows (FAN_AIRFLOW_MISMATCH)	??	??

The following table lists the default Switch Status monitoring thresholds and actions for the default moderate policy.

TABLE 57 Moderate policy default Switch Status monitoring thresholds and actions

Monitoring statistic	Switch status threshold values (Marginal/ Critical)	Actions
Absent or faulty power supply (BAD_PWR)	DCX, DCX+: -/3 All other supported platforms: 1/2	SW_CRITICAL, SNMP, EMAIL
Temperature sensors outside range (BAD_TEMP)	1/2	Marginal: SW_MARGINAL, SNMP, EMAIL

TABLE 57 Moderate policy default Switch Status monitoring thresholds and actions (continued)

Monitoring statistic	Switch status threshold values (Marginal/ Critical)	Actions
		Critical: SW_CRITICAL, SNMP, EMAIL
Absent or faulty fans (BAD_FAN)	1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
Flash usage (FLASH_USAGE)	90	RASLOG, SNMP, EMAIL
Percentage of marginal ports (MARG_PORTS)	6/10	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
Percentage of error ports (ERR_PORTS)	6/10	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
Percentage of faulty ports (FAULTY_PORTS)	6/10	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
Faulty blades (FAULTY_BLADE)	DCX, DCX+: 1/-	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: No Action
Faulty WWN (WWN_DOWN)	DCX, DCX+: -/1	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Core blade monitoring (DOWN_CORE)	DCX, DCX+: 1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
HA Sync (HA_SYNC)	DCX, DCX+, : sync=0	SW_MARGINAL, SNMP, EMAIL

The following table lists the default Switch Status monitoring thresholds and actions for the default conservative policy.

TABLE 58 Conservative policy default Switch Status monitoring thresholds and actions

Monitoring statistic	Switch status threshold values (Marginal/ Critical)	Actions
Absent or faulty power supply (BAD_PWR)	DCX, DCX+: -/3 All other supported platforms: 1/2	SW_CRITICAL, SNMP, EMAIL
Temperature sensors outside range (BAD_TEMP)	1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
Absent or faulty fans (BAD_FAN)	1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
Flash usage (FLASH_USAGE)	90	RASLOG, SNMP, EMAIL
Percentage of marginal ports (MARG_PORTS)	11/25	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Percentage of error ports (ERR_PORTS)	11/25	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Percentage of faulty ports (FAULTY_PORTS)	11/25	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Faulty blades (FAULTY_BLADE)	DCX, DCX+: 1/-	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: No Action
Faulty WWN (WWN_DOWN)	DCX, DCX+: -/1	Marginal: No Action

TABLE 58 Conservative policy default Switch Status monitoring thresholds and actions (continued)

Monitoring statistic	Switch status threshold values (Marginal/ Critical)	Actions
		Critical: SW_CRITICAL, SNMP, EMAIL
Core blade monitoring (DOWN_CORE)	DCX, DCX+: 1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
HA Sync (HA_SYNC)	DCX, DCX+, : sync=0	SW_MARGINAL, SNMP, EMAIL

Traffic Performance thresholds

The following Traffic Performance monitors support the Minute, Hour, and Day timebases: Receive throughput, Transmit Frame Count, Receive Frame Count, Transmit throughput, IO Read Command Count, IO Write Command Count, IO Read Data, IO Write Data. The "Throughput Degradation" monitor supports Minute, Hour, and Day timebases as well as the "None" timebase.