Brocade® Fabric OS® MAPS User Guide, 8.2.x

User Guide
14 October 2022

Table of Contents

Introduction

About This Document

This document provides an overview of the Monitoring and Alerting Policy Suite (MAPS), its elements and categories, and how you can set up and operate MAPS. This document also explains MAPS groups, conditions, rules, policies, and the monitoring dashboard.

Supported Hardware and Software

The following hardware platforms are supported by Brocade® Fabric OS® 8.2.x.

Brocade Gen 5 (16Gb/s) Fixed-Port Switches

- Brocade 6505 Switch
- Brocade 6510 Switch
- Brocade 6520 Switch
- Brocade M6505 Blade Server SAN I/O Module
- Brocade 6542 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 (16Gb/s) Directors

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 (32Gb/s) Fixed-Port Switches

- Brocade G610 Switch
- Brocade G620 Switch
- Brocade G630 Switch
- Brocade 7810 Extension Switch

Brocade Gen 6 (32Gb/s) Directors

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

MAPS Rules and Groups Altered in This Version

For Fabric OS 8.2.1 and 8.2.1a, changes have been made to the rules and groups. You must update your MAPS policies to use the correct rules and groups.

NOTE
Obsolete rules are automatically removed from the default and user-defined policies.

The following table lists the rules that are newly added for different platforms.

Table 1: Rules Added for Different Platforms

Rule Name	Condition	Actions
defALL_F_PORTS DEVICE_LOGIN_DISTRIBUTION _BALANCED	ALL_F_PORTS (DEVICE_LOGIN_DISTRIBUTION == BALANCED) RASLOG	RASLog
defALL_F_PORTS DEVICE_LOGIN_DISTRIBUTION _IMBALANCED	ALL_F_PORTS (DEVICE_LOGIN_DISTRIBUTION == IMBALANCED)	RASLog, RE_BALANCE
defALL_F_PORTS DEVICE_LOGIN_DISTRIBUTION _RE_BALANCE_FAILED	ALL_F_PORTS (DEVICE_LOGIN_DISTRIBUTION == RE_BALANCE_FAILED	RASLog, ACTIONS
defALL_25Km_32GELWL_SFPCURRENT_70	ALL_25Km_32GELWL_SFP(CURRENT>=70)	SFP_MARGINAL, RASLog, SNMP, EMAIL
defALL_25Km_32GELWL_SFPRXP_1995	ALL_25Km_32GELWL_SFP(RXP>=1995)	SFP_MARGINAL, RASLog, SNMP, EMAIL
defALL_25Km_32GELWL_SFPSFP_TEMP_75	ALL_25Km_32GELWL_SFP(SFP_TEMP>=75)	SFP_MARGINAL, RASLog, SNMP, EMAIL
defALL_25Km_32GELWL_SFPSFP_TEMP_n5	ALL_25Km_32GELWL_SFP(SFP_TEMP<=-5)	SFP_MARGINAL, RASLog, SNMP, EMAIL
defALL_25Km_32GELWL_SFPTXP_6309	ALL_25Km_32GELWL_SFP(TXP>=6309)	SFP_MARGINAL, RASLog, SNMP, EMAIL
defALL_25Km_32GELWL_SFPVOLTAGE_3600	ALL_25Km_32GELWL_SFP(VOLTAGE>=3600)	SFP_MARGINAL, RASLog, SNMP, EMAIL
defALL_25Km_32GELWL_SFPVOLTAGE_3000	ALL_25Km_32GELWL_SFP(VOLTAGE<=3000)	SFP_MARGINAL, RASLog, SNMP, EMAIL

Contacting Technical Support for Your Brocade® Product

If you purchased Brocade® product support from a Broadcom® OEM or solution provider, contact your OEM or solution provider for all your product support needs.

- OEM and solution providers are trained and certified by Broadcom to support Brocade products.
- Broadcom provides backline support for issues that cannot be resolved by the OEM or solution provider.
- Brocade Supplemental Support augments your existing OEM support contract, providing direct access to Brocade expertise. For more information on this option, contact Broadcom or your OEM.
- For questions regarding service levels and response times, contact your OEM or solution provider.

If you purchased Brocade product support directly from Broadcom, use one of the following methods to contact the Technical Assistance Center 24x7. For product support information and the latest information on contacting the Technical Assistance Center, go to www.broadcom.com/support/fibre-channel-networking/contact-brocade-support.

Online	Telephone
For nonurgent issues, the preferred method is to log on to the Support portal at support.broadcom.com. (You must initially register to gain access to the Support portal.) Once registered, log on and then select **Brocade Products**. You can now navigate to the following sites: • **Case Management** • **Software Downloads** • **Licensing** • **SAN Reports** • **Brocade Support Link** • **Training & Education**	For Severity 1 (critical) issues, call Brocade Fibre Channel Networking Global Support at one of the phone numbers listed at www.broadcom.com/support/fibre-channel-networking/contact-brocade-support.

Document Feedback

Quality is our first concern. We have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission or if you think that a topic needs further development, we want to hear from you. Send your feedback to documentation.pdl@broadcom.com. Provide the publication title; topic heading; publication number and page number (for PDF documents); URL (for HTML documents); and as much detail as possible.

Monitoring and Alerting Policy Suite Overview

MAPS Overview

The Monitoring and Alerting Policy Suite (MAPS) is a storage area network (SAN) health monitoring and alerting utility with the capability of early fault detection and isolation. MAPS helps the administrators to

- monitor the health and performance of the SAN infrastructure proactively to ensure application uptime and availability by leveraging pre-defined, rule-/policy-based templates
- automate the policy-based real-time monitoring and alert thresholds on a per-port or group basis
- configure each MAPS-enabled switch to monitor itself constantly for potential faults and automatically alert you to problems before they become costly failures
- configure an entire fabric (or multiple fabrics) at one time using common rules and policies, or custom policies for specific ports or switch elements with the simplified fabric-wide threshold configuration, monitoring, and alerting.
- remediate the conditions that require attention and resolve them before they impact operations using the timely alerts.
- fence, toggle, or quarantine ports automatically for the fabric without operations intervention.
- monitor fabric-wide events, ports, FRUs, environmental parameters, security, traffic flows, and performance impacts

MAPS requires an active and valid Fabric Vision license. If you already have a Fabric Watch license and an Advanced Performance Monitoring (AMP) license, you can automatically get the MAPS functionality without a separate license. If you only have one of these licenses, acquire the other license to use all the MAPS features. Refer to the *Brocade Fabric OS Software Licensing Guide* for more information about licensing and how to obtain the necessary license keys.

MAPS gets shipped with predefined rules, groups, and policies. However, you can create custom rules, groups, or policies. MAPS enables cloning of pre-defined rules, groups, and policies to facilitate the management of the monitoring. Predefined rules, groups are system specific and each system has its own rules and groups. All the MAPS configuration is persistent and retain across reboot or HA fail over and can be uploaded and downloaded.

MAPS License Requirements

MAPS requires an active and valid Fabric Vision license.

Alternatively, if you are upgrading and already have a Fabric Watch license and an Advanced Performance Monitoring license, you can automatically get the MAPS functionality without a separate license. Refer to the *Brocade Fabric OS Software Licensing User Guide* for more information about licensing and how to obtain the necessary license keys. If you only have one of these licenses, acquire the other license to use all the MAPS features.

MAPS Activation

Note the following information when you activate MAPS:

- The Fabric Vision license must be activated (enabled) in the Fabric OS software to use the full set of MAPS options.
- Activating MAPS enables it for all logical switches in the chassis, but each logical switch can have its own MAPS configuration.

MAPS Configuration

MAPS automatically generates one configuration for each platform that is known as the *default* configuration. You can add rules, groups, and policies to create a user-created configuration.

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

Automated Generation of Default Configuration

When MAPS generates a configuration for a particular platform, it includes only those rules that are applicable to that platform. For example, Extension Health monitoring rules are not included in the automatic generation of the configuration for platforms that do not support Extension Health. Also, some rules are replaced with others, because they are deprecated for a particular platform. The deprecated rules exist in some policies to support backward compatibility.

User-Created Configuration

A MAPS user-created 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.

Deleting a User-Created MAPS Configuration

Perform the following steps to delete the user-created MAPS configuration:

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 *Brocade Fabric OS Command Reference Manual*.

MAPS Interaction with Other Fabric OS Features

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

The following table outlines how MAPS interacts with specific features in the Fabric OS software:

Table 2: Interactions between Fabric OS Features and MAPS

Feature	MAPS interaction
Admin Domains	Admin Domains are deprecated in the Fabric OS 8.0.1 release. 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 an HA failover or HA reboot; however, MAPS restarts monitoring after an 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:

- The IP_EXTN_FLOW monitor runs in a polling mode and polls the DP stats at every 10 seconds of interval. There is a scenario where if the threshold value is crossed within the polling interval and again comes back to a value below the threshold values and alerts may not be seen.
- On Brocade G620, MAPS does not monitor the ports 44-47 if the in-flight encryption or compression is enabled on these ports.
- Small form-factor pluggable (SFP) transceivers on the 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), the event might not be shown on the dashboard.

See Monitoring Flow Vision Flow Monitor Data with MAPS 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 the Fabric OS 8.2.x release, there are two main scenarios:

- 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. Also, 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. The *Bottleneck Monitor* functionality was enabled by default in the Fabric OS software prior to the 8.0.x release. MAPS *Fabric Performance Impact (FPI)* monitoring is enabled by default in the Fabric OS 8.0.x and later releases.

Upgrading if Basic MAPS Monitoring is in Use without a Fabric Vision License

If MAPS is enabled implicitly (without the Fabric Vision license), any existing user-defined configurations are not initialized, and MAPS is enabled with the base policy, *dflt_base_policy*. This policy contains 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 with a Fabric Vision License on the switch, MAPS continues to monitor the switch afterward without any change in operation. MAPS is enabled with the same active policy that was previously in force on each logical switch and continues to monitor the fabric based on that policy.

MAPS is always enabled. If the Fabric Vision license is installed on the switch at the time the firmware is upgraded, MAPS is enabled after the upgrade with the active policy named in the configuration. If there is no active policy named, MAPS is enabled with the default conservative policy, *dflt_conservative_policy*. You can choose a different MAPS policy from one of the default policies or custom defined policies.

Firmware Downgrade Considerations for MAPS

You cannot downgrade from Fabric OS 8.2.x to an earlier version if the rebalance action is configured at the switch level using `mapsconfig --action` command.

The following are the primary downgrade scenarios for the Fabric OS 8.2.x release for MAPS:

- If any user defined MAPS rule is created to monitor the number of IP Extension Flows, the firmware downgrade to the pre-Fabric OS 8.2.x release displays a warning message to remove such rules before the downgrade.

  ```
  WARNING:MAPS user defined rules for IP_EXTN_FLOW will not be monitored in pre-8.2.0 release. Please delete
    these rules from FID: 128
  ```

  After the downgrade, new rule creation for IP_EXTN_FLOW is not allowed, and the existing predefined rules from 8.x are not monitored. However, you can still be able to view and delete the rules.
- When downgrading to any supported version of Fabric OS software without an active Fabric Vision license, the switch continues to use the same MAPS policy.

When downgrading from the Fabric OS 8.1.x software to a previous version, you can expect the following MAPS-related behaviors:

- Downgrading to previous versions of Fabric OS software fails 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, for example, if either MAPS actions or rules include Fabric Performance Impact monitoring or port decommissioning.
- Downgrading from Fabric OS software 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 is not blocked; however, MAPS displays the following warning message:

  ```
  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 with 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 software prior to 7.3.x 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 is enabled and slow-drain isolations are already performed. Restore the isolated slow-drain flows to their original virtual circuits before proceeding with the downgrade because the commands to restore them are not available in the downgraded firmware version.
- Downgrading is not allowed if the SDDQ action is enabled. You must clear the ports from the quarantined state using the `sddquarantine --clear` *`<slot/port>`* command or `sddquarantine all` command before you downgrade the switch firmware.
- When downgrading from Fabric OS 8.1.x to a version prior to Fabric OS 8.0.x, MAPS prompts with a warning whether you want to proceed if user-defined rules are present that contain the following monitoring systems:
  – SFPs that were added for the Fabric OS 8.0.x or Fabric OS 8.1.x release
  – Fan air-flow direction
  – Port zoned device ratio
  – IO Insight
  – Extension rules
  – Certificate rules
  – GE_Port rules

  Rules containing these monitoring systems are not monitored in the versions prior to Fabric OS 8.0.x version and the rules should be deleted. These monitoring systems include the following:

Table 3: Monitoring Systems that cannot be Included in Rules Prior to Fabric OS 8.0.x Software

ALL_32GSWL_QSFP ALL_32GSWL_SFP ALL_32GLWL_SFP ALL_25Km_16GLWL_SFP ALL_25Km_32GELWL_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 to a version prior to Fabric OS 8.1.x, a warning message is displayed to determine if any rules include user-configured severity values. If you continue downgrading, the user-configured severity values are ignored.
- Downgrading is not allowed if user-defined rules are present that include:
  - Rule-on-rule (RoR) rules that monitor base rules
  - Rules that contain the UNQUAR action
  - Rules that contain the UNINSTALL_VTAP action
  - Rules that contain names with more than 40 characters

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

Features That Do Not Require a Fabric Vision License

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

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

- Switch status policies
- Switch resource changes (FRU state, Flash memory space, Temperature, CPU usage, Memory usage, and Ethernet management port)
- FPI monitoring
- FRU health monitoring (not including SFPs)

For more information, refer to MAPS Commands That do not Need a License.

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 software earlier than this version.

Table 4: MAPS Commands Accessible without a Fabric Vision License

Command	Effect	Description
`logicalgroup --show`	Displays unlicensed feature groups details.	Viewing Group Information
`mapsconfig --show`	Displays the global actions that are currently allowed on a switch.	Enabling or Disabling Rule Actions at a Global Level
`mapsconfig --action`	Specifies which global actions are allowed on a switch.	Enabling or Disabling Rule Actions at a Global Level
`mapsdb --show`	Displays the MAPS dashboard.	MAPS Dashboard Display Options
`mapspolicy --show –summary`	Displays a summary of all the policies on the switch.	Viewing Policy Information
`mapspolicy --show <policy_name>`	Displays the rules for the specified MAPS policy on the switch.	Viewing Policy Information
`mapsrule --show <rule_name>`	Displays the details of the specified MAPS rule on the switch.	Viewing MAPS Rules
`mapsrule --show all`	Displays all the MAPS rules on the switch.	Viewing MAPS Rules
`mapssam --show flash`	Displays the flash memory usage as a percentage.	MAPS Service Availability Module
`mapssam --show cpu`	Displays the CPU usage as a percentage.	MAPS Service Availability Module
`mapssam --show memory`	Displays the general RAM memory usage as a percentage, along with total, used, and free memory values.	MAPS Service Availability Module
`mapshelp`	Displays list of MAPS commands.	

NOTE
Even without a Fabric Vision license, help for all the MAPS commands is displayed.

MAPS Setup and Operation

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 Fabric OS 7.4.0 release. Without a Fabric Vision license, the MAPS functionality is limited to rules included in the base policy, *dflt_base_policy*.

Activating MAPS without a Fabric Vision License

If you have not activated a Fabric Vision license, you can still use a limited set of MAPS functions.

MAPS is automatically enabled when you install Fabric OS 7.4.0 or later versions; however, if you have not installed and activated the Fabric Vision license, you can use only the 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 the Fabric Vision license installed or activated:

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

switch:admin> mapsdb --show

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

DB start time:                  Wed Apr 13 15:29:10 2018
Active policy:                  dflt_base_policy
Configured Notifications:       SW_CRITICAL,SW_MARGINAL,SFP_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, perform the following:
   a) Enter `licenseadd` , followed by the license key provided for your license.
   b) Use `licenseshow` to ensure the license is 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 installed a Fabric Vision license, you should use the conservative, aggressive, or moderate policies. Use the base policy only for basic monitoring similar to using MAPS without a license. See Features Not Requiring a License 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
      You can configure any basic actions such as RASLog, Email, and SNMP. SW_CRITICAL, SW_MARGINAL, and SFP_MARGINAL are enabled by default.

   b) Use `mapsconfig --show` to display and verify the configuration 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:

   - See Predefined Policies for details on the default MAPS policies.
   - See Viewing the MAPS Dashboard for details on the `mapsdb` command output.
   - See MAPS Rule Actions 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 installed a Fabric Vision license, you should use the conservative, aggressive, or moderate policies. Use the base policy only for basic monitoring, similar to using MAPS without a license. See Features Not Requiring a License 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. Modify the rules used by the policy as required.
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, 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 --create` <*severe_rule_name*> to confirm the rule values.
4. Enter `mapsrule --create` <*persistent_rule_name*> to confirm the rule values.

   Both of the following cases 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 can 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.

Perform the following steps to set the active MAPS 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. The default policies are:
   - dflt_conservative_policy
   - dflt_aggressive_policy
   - dflt_moderate_policy
   - dflt_base_policy

   NOTE
   If you installed a Fabric Vision license, you should use the conservative, aggressive, or moderate policies. Use the base policy only for basic monitoring similar to using MAPS without a license. See Feature Monitors Not Requiring a License for details.
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, see Predefined Policies.

Pausing MAPS Monitoring

You can stop monitoring ports, FCIP circuits, or SFPs in MAPS. You may stop monitoring 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 another 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

The Monitoring and Alerting Policy Suite (MAPS) contains various structural elements.

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

Table 5: MAPS Structural Elements

Element	Description
Action	The activity performed by MAPS if a condition defined in a rule evaluates to true. For more information, see MAPS Actions.
Category	A grouping of similar elements that can be monitored (for example, *Security Violations*). For more information, see MAPS Monitoring Categories.
Condition	A 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, see MAPS Conditions.
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, see MAPS Groups Overview.
Dashboard	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. For more information, see MAPS Dashboard Overview.
Monitoring system	A value (measure or statistic) that can be monitored. For more information, see MAPS Monitoring Categories.
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, see MAPS Rules Overview.
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, see MAPS Policies Overview.

MAPS Monitoring Categories

MAPS provides the following categories for monitoring:

- Back-End Port Health
- Gigabit Ethernet Ports
- Fabric Performance Impact
- Fabric State Changes
- Extension Health
- FRU Health
- Port Health
- Security Health
- Switch Resources
- Switch Status Policy
- Traffic Performance

In addition to setting alerts and other actions based on these categories, the MAPS dashboard also displays their status. See MAPS Dashboard Overview for information on using the MAPS dashboard.

Back-End Port Health

The Back-End Port Health category lets you monitor the health of the back-end switch ports for invalid ordered sets, CRC and link reset error rates, frame length (either too long or truncated), and invalid transmission words.

The following table lists the monitored parameters in this category:

Table 6: Back-End Port Health Category Parameters

Parameter Name	Description	Data Type	Range of Data
BAD_OS	The number of invalid ordered sets (OS) outside the frame.	Uint32	0 - 999999999
CRC	The number of cyclic redundancy check (CRC) errors.	Uint32	0 - 999999999
FRM_LONG	The number of frames that were detected that are longer than expected (greater than 2148 bytes).	Uint32	0 - 999999999
FRM_TRUNC	The number of frames that were detected that are truncated and shorter than expected (less than 36 bytes).	Uint32	0 - 999999999
ITW	The number of times an invalid transmission word (ITW) 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.	Uint32	0 - 999999999
LR	The number of times a link reset (LR) error occurs on a back-end port.	Uint32	0 - 999999999

For more information on back-end port monitoring, see Back-End Port Monitoring.

Fabric Performance Impact

Fabric Performance Impact (FPI) monitors congestion related issues on all physical E_Ports and F_Ports at all times. Specifically, FPI monitors abnormal device latency and oversubscription issues using the following monitoring systems and rules.

NOTE
Whenever IO_PERF_IMPACT is used, IO_LATENCY_CLEAR must also be included in the active policy to clear the latency record.

The following table lists the monitored parameters in this category:

Table 7: Fabric Performance Impact Category Parameters

Parameter Name	Description	Data Type	Range of Data
BE_LATENCY_IMPACT	In all fabric edge switches connecting to Brocade Analytics Monitoring Platform with active VTAP flows, MAPS monitors the *tim_txcrd_z* back-end port on a mirrored traffic path.	enum	IO_PERF_IMPACT
DEV_LATENCY_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. When a timeout is seen on a port, the bottleneck state of that port is changed to IO_FRAME_LOSS. Average R_RDY_DELAY is greater than or equal to 80ms.	enum	IO_PERF_IMPACT, IO_FRAME_LOSS
IT_FLOW	Initiator to target flow ratio.	Uint32	0 - 999999999

Parameter Name	Description	Data Type	Range of Data
RX	The percentage of port bandwidth being used by received, incoming (RX) traffic. For example, if the port speed is 10Gb/s and the port receives 5Gb of data in one second, then the percentage of RX utilization is 50% (5Gb*100/(10Gb*1 second)). For a master trunk port, this indicates the RX percentage for the entire trunk.	Float	0 - 100
TX	The percentage of port bandwidth being used by transmitted, outgoing (TX) traffic. For example, if the port speed is 10Gb/s and the port sends 5Gb of data in one second, then the percentage of TX utilization is 50% (5Gb*100/(10Gb*1 second)). For a master trunk port, this indicates the TX percentage for the entire trunk.	Float	0 - 100
DEV_LOGIN_DIST	The port group state. This parameter has the following threshold values: • BALANCED: All ports in the group have close to an equal number of devices connected to them. • IMBALANCED: At least one port in the group does not have any devices connected or has more than one device connected but less than the devices connected to the rest of the ports. • BALANCE_FAILED: MAPS performed the rebalance action but failed to rebalance the group.	enum	BALANCED, IMBALANCED, BALANCE_FAILED
UTIL	The percentage of individual port (or trunk) bandwidth being used at the time of the most recent poll.	Float	0 - 100

For more information on Fabric Performance Impact monitoring, refer to Fabric Performance Monitoring Using MAPS.

Fabric State Changes

The Fabric State Changes category contains areas of potential fabric related or switch related 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 8: Fabric State Change Category Parameters

Parameter name	Description	Data Type	Range of Data
BB_FCR_CNT	Monitors the number (count) of Fibre Channel routers (FCR) configured in a backbone (BB) fabric.	Uint32	0 - 999999999
DID_CHG	Monitors the number of forced domain ID (DID) 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.	Uint32	0 - 999999999
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).	Uint32	0 - 999999999

Parameter name	Description	Data Type	Range of Data
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	Uint32	0 - 999999999
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	Uint32	0 - 999999999
FLOGI	Activates when ports and devices initialize with the fabric (fabric logins).	Uint32	0 - 999999999
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 (Layer 2 device count percentage). The switches in a pure Layer 2 fabric do not participate in the metaSAN.	Float	0 - 100
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 (LSAN device count percentage). This percentage includes devices imported from any attached edge fabrics.	Float	0 - 100
ZONE_CFGSZ_PER	Monitors the *used zone configuration* size relative to the maximum zone configuration size on the switch.	Float	0 - 100
ZONE_CHG	Tracks the number of zone changes. As zoning is a security provision, frequent zone changes can indicate a security breach or weakness. Zone change messages occur whenever there is a change in zone configurations.	Uint32	0 - 999999999

Extension Health

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

> **NOTE**
> The FCIP Health category name is changed to Extension Health Category from Fabric OS 8.2.x and later releases. All the monitoring that was done as part of FCIP Health continues to be monitored under the Extension Health category.

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

Table 9: Extension Health Category Parameters

Parameter Name	Description	Data Type	Range of Data
CIR_PKTLOSS	The percentage of the total number of FCIP circuit packets that had to be re-transmitted.	Float	0 - 100
CIR_STATE	The state of the FCIP circuit has changed for one of the following reasons: • The circuit has gone offline. • The circuit has come online. • The circuit is faulty.	Uint32	0 - 999999999
CIR_UTIL	The percentage of FCIP circuit utilization in the configured time period (this can be a minute, or an hour, or a day).	Float	0 - 100
CIR_QOS_UTIL	Circuit QoS utilization percentage	Float	0 - 100
CIR_QOS_PKTLOSS	Circuit QoS packet loss percentage	Float	0 - 100
TUNNEL_IP_UTIL	Tunnel IP utilization percentage	Float	0 - 100
TUNNEL_STATE	Tunnel state changes	Uint32	0 - 999999999
TUNNEL_UTIL	Tunnel utilization	Float	0 - 100
QOS_UTIL	QoS utilization percentage	Float	0 - 100
CIR_IP_UTIL	Circuit IP utilization percentage	Float	0 - 100
CIR_IP_PKTLOSS	Circuit IP packet loss percentage	Float	0 - 100
IP_JITTER	The amount of jitter in an IP circuit. This is a calculated percentage and only applies to the circuit group.	Float	0 - 100
IP_RTT	The IP circuit round-trip latency. This is an absolute value and only applies to the circuit group.	Uint32	0 - 999999999
JITTER	The amount of jitter in an FCIP circuit. This is a calculated percentage and only applies to the circuit group.	Float	0 - 100
PKTLOSS	The percentage of the total number of packets that have had to be re-transmitted in each QoS level. This applies to each FCIP QoS group only.	Float	0 - 100
RTT	The FCIP circuit round-trip latency. This is an absolute value and only applies to the circuit group.	Uint32	0 - 999999999
IP_EXTN_FLOW	The number of LAN IP Extension flows (TCP connections)	Uint32	1 - 511

Table 10: Brocade 7840 Extension Health Category Parameters

Parameter Name	Description	Data Type	Range of Data
STATE_CHG	The count of FCIP tunnel state changes. This applies to the tunnel group only.	Uint32	0 - 999999999
UTIL	The percentage of FCIP utilization. This applies to both the tunnel and the tunnel QoS groups.	Float	0 - 100

See Extension Monitoring Thresholds for the default values.

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:

Table 11: FRU Health Category Parameters

Monitored Parameter	Description	Data Type	Range of Data
PS_STATE	The state of a power supply (PS) has changed.	enum	OUT, ON, FAULTY
FAN_STATE	The state of a fan has changed.	enum	OUT, ON, FAULTY
BLADE_STATE	The state of a blade has changed.	enum	FAULTY, OFF, ON, OUT
SFP_STATE	The state of the SFP transceiver has changed (for both FC or FCoE).	enum	IN, OUT, FAULTY
WWN	The state of a WWN card has changed.	enum	OUT, ON, FAULTY

Gigabit Ethernet Ports

The Gigabit Ethernet (GE) Ports category monitors statistics for GE ports and takes action based on the configured thresholds and actions. You can configure thresholds and apply the configuration to all ports.

The following table describes the monitored parameters in this category. The value in the first column is the parameter name you specify in the `mapsrule -monitor` command.

Table 12: Gigabit Ethernet Ports Category Parameters

Parameter Name	Description	Data Type	Range of Data
GE_CRC	The number of times an invalid cyclic redundancy check (CRC) error occurs on a GE 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.	Uint32	0 - 999999999
GE_LOS_OF_SIG	The number of times that a signal loss occurs in offline GE ports. Signal loss indicates that no data is moving through the port. A loss of signal usually indicates a hardware problem.	Uint32	0 - 999999999

NOTE
The monitoring systems listed in the table above are applicable to the extension platforms.

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.

NOTE
VE_Ports are not monitored.

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 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). This is 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.

Port Health Category Parameters

The following table describes the monitored parameters in this category. In the value in the first column is the parameter name you specify in the `mapsrule -monitor` command.

Table 13: Port Health Category Parameters

Parameter Name	Description	Data Type	Range of Data
C3TXTO C3TXTO is not monitored for resident N_Ports with any default policy in AG mode.	The number of Class 3 discard frames due to timeouts.	Uint32	0 - 999999999
CRC	The number of times an invalid cyclic redundancy check (CRC) 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.	Uint32	0 - 999999999
CURRENT	The amperage supplied to the SFP transceiver in milliamps (mA). Current area events indicate hardware failures.	Uint32	0 - 999999999
DEV_NPIV_LOGINS	The number of NPIV logins to the device. See Monitoring NPIV Logins to F_Ports for details.	Uint32	0 - 100

Parameter Name	Description	Data Type	Range of Data
ITW	The number of times an invalid transmission word (ITW) error occurs on a port. A word did not transmit successfully, resulting in encoding errors. Invalid word messages usually indicate a hardware problem. 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 (8Gb/s or lower) speeds, and PCS violations occur only at high (10Gb/s or higher) speeds.	Uint32	0 - 999999999
LF	The number of times a link failure (LF) 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.	Uint32	0 - 999999999
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.	Uint32	0 - 999999999
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.x and later versions, the loss of sync monitoring happens only for online ports.	Uint32	0 - 999999999
LR	The ports on which the number of link resets (LRs) exceeds the specified threshold value.	Uint32	0 - 999999999

Parameter Name	Description	Data Type	Range of Data
PE	The number of times a protocol error (PE) occurs on a port. Occasionally, PEs occur due to software glitches. Persistent errors generally occur due to hardware problems.	Uint32	0 - 999999999
PWR_HRS	The number of hours that an SFP transceiver has been powered up.	Uint32	0 - 999999999
RXP	The power of the incoming laser (receive power) 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.	Uint32	0 - 999999999
SFP_TEMP	The temperature of the SFP transceiver in degrees Celsius. A high temperature indicates that the SFP transceiver could be in danger of damage.	Int32	-40 to 100
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.	Uint32	0 - 999999999
TXP	The power of the outgoing laser (transmit power) in microwatts (µW). This is used to determine if the SFP transceiver is in good working condition. If the counter often exceeds the threshold, the SFP transceiver is deteriorating.	Uint32	0 - 999999999
VOLTAGE	The voltage supplied to the SFP transceiver in millivolts (mV). If this value exceeds the threshold, the SFP transceiver is deteriorating.	Uint32	0 - 999999999
ENCR_BLK	Encryption block	Uint32	0 - 999999999
ENCR_DISC	Encryption discard	Uint32	0 - 999999999
ENCR_SHRT_FRM	Encryption short frame	Uint32	0 - 999999999

Security Violations

The Security Violations Health 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 Health Category Parameters

Parameter Name	Description	Data Type	Range of Data
DAYS_TO_EXPIRE	The number of days before the certificate expires. This parameter notifies the user when the threshold is reached.	Uint32	0 - 99999
EXPIRED_CERTS	The number of expired certifications detected.	Uint32	0 - 999999999
SEC_AUTH_FAIL	The number of SLAP failures that occurred when packets tried to pass from a non-secure switch to a secure fabric (authorization failed).	Uint32	0 - 999999999
SEC_CERT	The number of invalid certificates detected.	Uint32	0 - 999999999
SEC_CMD	The number of times that commands that are permitted only to the primary Fibre Channel Switch (FCS) were executed on another switch (illegal commands).	Uint32	0 - 999999999
SEC_DCC	The number of times an unauthorized device attempted to log in to a secure fabric (DCC violations).	Uint32	0 - 999999999
SEC_FCS	The number of times an unauthorized device attempted to log in to a secure fabric (FCS violations).	Uint32	0 - 999999999
SEC_HTTP	The number of times a browser access request reached a secure switch from an unauthorized IP address (HTTP violations).	Uint32	0 - 999999999
SEC_IDB	The number of times a secure switch with a different version stamp was detected (incompatible security DB).	Uint32	0 - 999999999
SEC_LV	The number of login violations (LV) that occurred when a secure fabric detected a login failure.	Uint32	0 - 999999999
SEC_SCC	The number of SCC violations that occurred when an unauthorized switch tried to join a secure fabric. The WWN of the unauthorized switch appears in the ERRLOG.	Uint32	0 - 999999999
SEC_TELNET	The number of Telnet violations that occurred when a Telnet connection request reached a secure switch from an unauthorized IP address.	Uint32	0 - 999999999
SEC_TS	The number of Time Server (TS) violations that occurred when an out-of-synchronization error was detected.	Uint32	0 - 999999999

Switch Resources

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

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

- Configure thresholds for the 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 trigger a set of specified MAPS alerts if exceeded. 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 15: Switch Resources Category Parameters

Parameter Name	Description	Data Type	Range of Data
CPU	The percentage of CPU available, calculated by comparing the percentage of CPU consumed with the configured threshold value.	Float	0 - 100
ETH_MGMT_PORT_STATE	The status of the Ethernet management port (Bond0).	enum	UP, DOWN
FLASH_USAGE	The percentage of storage space used, calculated by comparing the percentage of flash space consumed with the configured high threshold value.	Uint32	0 - 100
MEMORY_USAGE	The available memory, calculated by comparing the percentage of memory consumed with the configured threshold value.	Float	0 - 100
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 might occur to the switch.	enum	IN_RANGE, OUT_OF_RANGE
VTAP_IOPS	VTAP I/Os per second	Uint64	0 - 0xffffffffffffffff

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 power supply failures, it is considered to be in a marginal state, or if it has two failures, it is in a critical (down) state.

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.

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

Table 16: Switch Status Policy Category Parameters

Parameter Name	Description	Data Type	Range of Data
BAD_FAN	The number of problematic fans including missing fans and faulty fans.	Uint32	0 - 999999999
BAD_PWR	The number of times a power (PWR) supply threshold detected absent or failed power supplies, and power supplies that were not in the correct slot for redundancy.	Uint32	0 - 999999999
BAD_TEMP	The number of time temperature thresholds detected faulty temperature sensors.	Uint32	0 - 999999999
DOWN_CORE	The number of faulty (down) core blades (applies to modular switches only).	Uint32	0 - 999999999

Parameter Name	Description	Data Type	Range of Data
ERR_PORTS	The percentage of ports with errors, including • Ports that are fenced and decommissioned by MAPS or Brocade Network Advisor • Loopback ports • Segmented ports • Ports that are segmented due to security violations	Float	0 - 100
EXPIRED_CERTS	The details of install but expired switch certificates.	Uint32	0 - 999999999
FAULTY_BLADE	The number of faulty (down) blades (applies to modular switches only).	Uint32	0 - 999999999
FAULTY_PORTS	The percentage of hardware-related port faults.	Float	0 - 100
FLASH_USAGE	The percentage of flash usage by the system, such as faulty SFPs or laser FTL.	Float	0 - 100
HA_SYNC	Indicates whether the system is in or out of sync.	Uint32	0 - 1
MARG_PORTS	The percentage of physical ports, E_Ports, and F_Ports (both optical and copper) that exceed threshold settings. Whenever these thresholds are persistently high, the port does not have sufficient credits to operate.	Float	0 - 100
MISSING_SFP	The percentage of ports that are missing SFP media.	Float	0 - 100
WWN_DOWN	The number of faulty (down) WWN cards (applies to modular switches only).	Uint32	0 - 999999999
FAN_AIRFLOW_MISMATCH	Mismatched air flow of fans.	enum	TRUE

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).

Traffic Performance

The Traffic Performance category groups areas that monitor the metrics of flows that were either (1) created in Flow Vision and for which Flow Monitor is enabled or (2) created using VM Insight to monitor applications supported by VM Insight. 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.

Table 17: Traffic Performance Category Parameters

Parameter Name	Description	Data Type	Range of Data
IO_RD	The number of SCSI I/O read command frames recorded for the flow. (Not supported for VM Insight flows.)	Uint64	0 - 0xffffffffffffffff
IO_RD_BYTES	The number of SCSI I/O bytes read as recorded for the flow. (Not supported for VM Insight flows.)	Uint64	0 - 0xffffffffffffffff
IO_WR	The number of SCSI I/O write command frames recorded for the flow. (Not supported for VM Insight flows.)	Uint64	0 - 0xffffffffffffffff
IO_WR_BYTES	The number of SCSI I/O bytes written as recorded for the flow. (Not supported for VM Insight flows.)	Uint64	0 - 0xffffffffffffffff
RX	The percentage of frames successfully received by the flow destination switch.	Float	0 - 100

Parameter Name	Description	Data Type	Range of Data
RX_FCNT	The number of frames received by the flow destination (received frame count).	Uint64	0 - 0xffffffffffffffff
RX_THPUT	The number of megabytes (MB) received per second by the flow destination (received throughput).	Uint64	0 - 0xffffffffffffffff
TX	The percentage of frames successfully transmitted by the flow source switch.	Float	0 - 100
TX_FCNT	The number of frames transmitted from the flow source (transmitted frame count).	Uint64	0 - 0xffffffffffffffff
TX_THPUT	The number of megabytes (MB) transmitted per second by the flow source (transmitted throughput).	Uint64	0 - 0xffffffffffffffff
RD_STATUS_TIME_LT_8K	Time for read I/O request less than 8K	Uint64	0 - 0xffffffffffffffff
RD_STATUS_TIME_8_64K	Time for read I/O request greater than or equal to 8K but less than 64K	Uint64	0 - 0xffffffffffffffff
RD_STATUS_TIME_64_512K	Time for read I/O request greater than or equal to 64K but less than 512K	Uint64	0 - 0xffffffffffffffff
RD_STATUS_TIME_GE_512K	Time for read I/O request greater than or equal to 512K	Uint64	0 - 0xffffffffffffffff
WR_STATUS_TIME_LT_8K	Time for write I/O request less than 8K	Uint64	0 - 0xffffffffffffffff
WR_STATUS_TIME_8_64K	Time for write I/O request greater than or equal to 8K but less than 64K	Uint64	0 - 0xffffffffffffffff
WR_STATUS_TIME_64_512K	Time for write I/O request greater than or equal to 64K but less than 512K	Uint64	0 - 0xffffffffffffffff
WR_STATUS_TIME_GE_512K	Time for write I/O request greater than or equal to 512K	Uint64	0 - 0xffffffffffffffff
RD_1stDATA_TIME_LT_8K	Time for first data read less than 8K	Uint64	0 - 0xffffffffffffffff
RD_1stDATA_TIME_8_64K	Time for first data read greater than or equal to 8K but less than 64K	Uint64	0 - 0xffffffffffffffff
RD_1stDATA_TIME_64_512K	Time for first data read greater than or equal to 64K but less than 512K	Uint64	0 - 0xffffffffffffffff
RD_1stDATA_TIME_GE_512K	Time for first data read greater than or equal to 512K	Uint64	0 - 0xffffffffffffffff
WR_1stXFER_RDY_LT_8K	First data transfer in the ready state less than 8K	Uint64	0 - 0xffffffffffffffff
WR_1stXFER_RDY_8_64K	First data transfer in ready state greater than or equal to 8K but less than 64K	Uint64	0 - 0xffffffffffffffff
WR_1stXFER_RDY_64_512K	First data transfer in ready state greater than or equal to 64K but less than 512K	Uint64	0 - 0xffffffffffffffff
WR_1stXFER_RDY_GE_512K	First data transfer in ready state greater than or equal to 512K	Uint64	0 - 0xffffffffffffffff
RD_PENDING_IO_LT_8K	Pending read I/O requests less than 8K	Uint64	0 - 0xffffffffffffffff
RD_PENDING_IO_8_64K	Pending read I/O requests greater than or equal to 8K but less than 64K	Uint64	0 - 0xffffffffffffffff
RD_PENDING_IO_64_512K	Pending read I/O requests greater than or equal to 64K but less than 512K	Uint64	0 - 0xffffffffffffffff
RD_PENDING_IO_GE_512K	Pending read I/O requests greater than or equal to 512K	Uint64	0 - 0xffffffffffffffff
WR_PENDING_IO_LT_8K	Pending write I/O requests less than 8K	Uint64	0 - 0xffffffffffffffff
WR_PENDING_IO_8_64K	Pending write I/O requests greater than or equal to 8K but less than 64K	Uint64	0 - 0xffffffffffffffff

Parameter Name	Description	Data Type	Range of Data
WR_PENDING_IO_64_512K	Pending write I/O requests greater than or equal to 64K but less than 512K	Uint64	0 - 0xffffffffffffffff
WR_PENDING_IO_GE_512K	Pending write I/O requests greater than or equal to 512K	Uint64	0 - 0xffffffffffffffff
RD_IO_RATE_LT_8K	Read I/O bytes less than 8K	Uint64	0 - 0xffffffffffffffff
RD_IO_RATE_8_64K	Read I/O bytes greater than or equal to 8K but less than 64K	Uint64	0 - 0xffffffffffffffff
RD_IO_RATE_64_512K	Read I/O bytes greater than or equal to 64K but less than 512K	Uint64	0 - 0xffffffffffffffff
RD_IO_RATE_GE_512K	Read I/O bytes greater than or equal to 512K	Uint64	0 - 0xffffffffffffffff
WR_IO_RATE_LT_8K	Written I/O bytes less than 8K	Uint64	0 - 0xffffffffffffffff
WR_IO_RATE_8_64K	Written I/O bytes greater than or equal to 8K but less than 64K	Uint64	0 - 0xffffffffffffffff
WR_IO_RATE_64_512K	Written I/O bytes greater than or equal to 64K but less than 512K	Uint64	0 - 0xffffffffffffffff
WR_IO_RATE_GE_512K	Written I/O bytes greater than or equal to 512K	Uint64	0 - 0xffffffffffffffff
RD_IOPS_LT_8K	Read I/O count less than 8K	Uint64	0 - 0xffffffffffffffff
RD_IOPS_8_64K	Read I/O count greater than or equal to 8K but less than 64K	Uint64	0 - 0xffffffffffffffff
RD_IOPS_64_512K	Read I/O count greater than or equal to 64K but less than 512K	Uint64	0 - 0xffffffffffffffff
RD_IOPS_GE_512K	Read I/O count greater than or equal to 512K	Uint64	0 - 0xffffffffffffffff
WR_IOPS_LT_8K	Written I/O count less than 8K	Uint64	0 - 0xffffffffffffffff
WR_IOPS_8_64K	Written I/O count greater than or equal to 8K but less than 64K	Uint64	0 - 0xffffffffffffffff
WR_IOPS_64_512K	Written I/O count greater than or equal to 64K but less than 512K	Uint64	0 - 0xffffffffffffffff
WR_IOPS_GE_512K	Written I/O count greater than or equal to 512K	Uint64	0 - 0xffffffffffffffff

MAPS Groups, Conditions, Rules, and Policies

MAPS Groups Overview

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

MAPS provides a set of predefined groups. MAPS also allows you to create a user-defined group and use the same group in rules to simplify 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 has different feature sets, and then import them into MAPS as a group.

Viewing Group Information

MAPS allows you to view the information for all logical groups collectively or 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 the information on 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_25Km_32GELWL_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 the MAPS perspective. The state of 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 is imported and 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 *Brocade Fabric OS Command Reference Manual* 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 18: Predefined MAPS Group

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 also includes all the ports in E_Port, EX_Port trunks, and AE ports.
ALL_F_PORTS	FC Port	All F_Ports in the logical switch. This also includes all the ports in F_Port trunks.
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 16Gb/s long wavelength (LWL) small form-factor pluggable (SFP) transceivers in the logical switch.
ALL_25Km_32GELWL_SFP	SFP	All 25 kilometer-capable 32Gb/s extended long wavelength (ELWL) small form-factor pluggable (SFP) transceivers in the logical switch.
ALL_32GLWL_SFP	SFP	All 32Gb/s long wavelength (LWL) small form-factor pluggable (SFP) transceivers in the logical switch.
ALL_32GSWL_SFP	SFP	All 32Gb/s short wavelength (SWL) small form-factor pluggable (SFP) transceivers in the logical switch.
ALL_32GSWL_QSFP	QSFP	All 4 X 32Gb/s short wavelength (SWL) quad small form-factor pluggable (QSFP) transceivers in the logical switch.
ALL_10GSWL_SFP	SFP	All 10Gb/s Short Wavelength (SWL) SFP transceivers on FC Ports in the logical switch.
ALL_10GLWL_SFP	SFP	All 10Gb/s Long Wavelength (LWL) SFP transceivers on FC Ports in the logical switch.
ALL_16GSWL_SFP	SFP	All 16Gb/s SWL SFP transceivers in the logical switch.
ALL_16GLWL_SFP	SFP	All 16Gb/s LWL SFP transceivers in the logical switch.
ALL_OTHER_SFP	SFP	ALL_OTHER_SFP transceivers do not belong to any other groups except ALL_SFP.
ALL_2K_QSFP	SFP	All 2 kilometer-capable 16Gb/s 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.

Predefined Group Name	Object Type	Description
ALL_WWN	WWN	All WWN cards in the chassis.
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	Control QoS members of a tunnel	All control QoS members of a tunnel.
ALL_TUNNEL_IP_HIGH_QOS	IP QoS members of a tunnel	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.
ALL_DP	Member count	ALL DP members

Table 19: Predefined MAPS Group for FCoE-Supported Fabrics

Predefined Group Name	Object Type	Description
ALL_ETH_PORTS	ETH Port	All VF_Ports and VN_Ports in the logical switch. This also includes all the ports in the VF_Port port-channels.
ALL_FCOE_40G_QSFP	QSFP	All 40GE short wavelength (SWL) quad small form-factor pluggable (QSFP) transceivers in the logical switch.
ALL_FCOE_10G_SFP	SFP	All 10GE short wavelength (SWL) small form-factor pluggable (SFP) transceivers in the logical switch.
ALL_FCOE_25G_SFP	SFP	All 25GE short wavelength (SWL) small form-factor pluggable (SFP) transceivers in the logical switch.
ALL_FCOE_100G_SR4_QSFP	QSFP	All 100GE short wavelength (SWL) quad small form-factor pluggable (QSFP) transceivers in the logical switch.

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 suitable 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.

Note the following points 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*, 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.

NOTE
Ports in the ETH mode are not included as a part of MAPS dynamic user-defined groups and cannot be included as a part of static user-defined groups. Ports in the ETH mode include VN_Ports, VF_Ports, and port-channels. This means that the ports in the ETH mode are not supported with the `logicalgroup --addmember` , `logicalgroup --create` , and the `logical group --update` commands.

Creating a Static User-Defined Group

MAPS allows you to use a static definition to create a group that can be monitored. The membership in that group is explicit and only changes if you redefine the group.

As an example of a static definition, you can define a group called *MY_CRITICAL_PORTS* and specify its members as *2/1-10, 2/15,* and *3/1-20*. In this case, the group has a fixed membership. To add or remove a member from the group, use the `logicalgroup` command and specify whether you want to 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,* and *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 *Brocade Fabric OS Command Reference Manual*.

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).

Perform the following steps to change which ports are in a static user-defined group:

1. Connect to the switch and log in using an account with admin permissions.
2. Use the following commands to add or delete the 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 *Brocade Fabric OS Command Reference Manual*.

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 can specify a port name or an attached device node WWN, and all ports which match the port name or device node WWN are automatically included in this group. When a port meets the criteria, it is automatically added to the group. When it does not 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 20: Group-Definition Operators

Character	Meaning	Explanation
*	Match any set of characters in the position indicated by the asterisk.	Defining the port name as *brcdhost** includes 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?* includes any port name that has exactly one character following *brcdhost*, such as brcdhost1, brcdhostn, and so on. However, brcdhostnew does 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]* includes only the port names brcdhost1, brcdhost2, and brcdhost3.

Character	Meaning	Explanation
!	Match the string and exclude any ports that match. You must include the entire term in single quotation marks (').	Defining the port name as *!brcdhost* includes 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*"
   ```

   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 *Brocade Fabric OS Command Reference Manual*.

Modifying a Dynamic User-Defined Group

MAPS allows you to change the definition pattern used to specify a dynamic user-defined group after you 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 the 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 *Brocade Fabric OS Command Reference Manual*.

Restoring a Group to Its Default Membership

MAPS allows you to restore the membership of any modified MAPS group 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 the GOBLIN_PORTS group, then restores the membership of the group and then 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 is 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 Overview

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, 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 triggers.

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 using `mapsrule --config` or creating a rule using `mapsrule --create` , you can use the *--value* option to specify a threshold value that, when exceeded, triggers specified actions.

For example, you can create a rule that monitors loss-of-signal errors on all host ports and triggers actions when the counter is greater than 14. When the counter reaches 15, the rule triggers the actions. The following is an example of such a rule:

```
switch:admin> mapsrule --create LoS_greater_than_14  -monitor  LOSS_SIGNAL
                                -group  ALL_HOST_PORTS -timebase day
                                -op ge -value 14 -action raslog,email,snmp
```

Timebase

The `-timebase` value specifies the time interval between samples and user-defined threshold values.

You can set the timebase to the following durations:

Table 21: Timebase Durations

Duration	Description	How Often Samples are Compared
`week`	Samples used for comparison are one week apart. (This timebase is valid only for rule-on-rule (RoR) rules.)	Once a week
`day`	Samples used for comparison are one day apart.	Once a day
`hour`	Samples used for comparison are one hour apart.	Once every hour
`minute`	Samples used for comparison are one minute apart.	Once every minute
`second`	Samples used for comparison are one second apart.	Once every second
`none`	A comparison is made between the real-time value and the configured threshold value.	N/A

NOTE
When you are specifying the timebase for a rule that monitors another rule (a rule-on-rule (RoR) rule), and the timebase of the base rule is *None*, then see Supported Timebases for RoR Rules When the Timebase of the Base Rule is NONE to determine which timebases can be defined in the RoR rule.

Supported Timebase

The following tables identify the timebase supported by monitoring systems within each group:

Table 22: Port Health: Monitoring System and Supported Timebase

Parameter Used In Rule Command	Monitoring System	Group	Day	Hour	Minute	None
C3TXTO	C3 transmission time-outs	FC Port	Yes	Yes	Yes	No
CRC	CRC errors	FC Port	Yes	Yes	Yes	No
ENCR_BLK	Encryption block	FC Port	Yes	Yes	Yes	No
ENCR_DISC	Encryption discard	FC Port	No	No	No	Yes
ENCR_SHRT_FRM	Encryption short frame	FC Port	Yes	Yes	Yes	No

Parameter Used In Rule Command	Monitoring System	Group	Day	Hour	Minute	None
ITW	Invalid transmit words	FC Port	Yes	Yes	Yes	No
LF	Link failure	FC Port	Yes	Yes	Yes	No
LOSS_SIGNAL	Loss of signal	FC Port	Yes	Yes	Yes	No
LOSS_SYNC	Loss of sync	FC Port	Yes	Yes	Yes	No
LR	Link reset	FC Port	Yes	Yes	Yes	No
PE	Protocol errors	FC Port	Yes	Yes	Yes	No
STATE_CHG	State changes	FC Port	Yes	Yes	Yes	No
CURRENT	SFP current	SFP	No	No	No	Yes
PWR_HRS	SFP power-on hours	SFP	No	No	No	Yes
RXP	SFP receive power	SFP	No	No	No	Yes
SFP_TEMP	SFP temperature	SFP	No	No	No	Yes
TXP	SFP transmit power	SFP	No	No	No	Yes
VOLTAGE	SFP voltage	SFP	No	No	No	Yes
DEV_NPIV_LOGINS	Device NPIV Logins	FC Port	No	No	No	Yes

Table 23: BE Port Health: Monitoring System and Supported Timebase

Parameter Used In Rule Command	Monitoring System	Group	Day	Hour	Minute	None
BAD_OS	Bad OS error	BE Port	Yes	Yes	Yes	No
CRC	CRC errors	BE Port	Yes	Yes	Yes	No
FRM_LONG	Frame too long	BE Port	Yes	Yes	Yes	No
FRM_TRUNC	Frame truncated	BE Port	Yes	Yes	Yes	No
ITW	Invalid transmit words	BE Port	Yes	Yes	Yes	No
LR	Link reset	BE Port	Yes	Yes	Yes	No

Table 24: GE Port Health: Monitoring System and Supported Timebase

Parameter Used In Rule Command	Monitoring System	Group	Day	Hour	Minute	None
GE_CRC	CRC errors	GE Port	Yes	Yes	Yes	No
GE_LOS_OF_SIG	Loss of signal	GE Port	Yes	Yes	Yes	No

NOTE
The monitoring systems in the table above are applicable to extension platform.

Table 25: Extension Health: Monitoring System and Supported Timebase

Parameter Used In Rule Command	Monitoring System	Group	Day	Hour	Minute	None
IP_RTT	Circuit IP round-trip time percentage	Circuit	Yes	Yes	Yes	No
JITTER	Circuit FC connection variance	Circuit	Yes	Yes	Yes	No
RTT	Circuit FC round-trip time	Circuit	Yes	Yes	Yes	No
CIR_STATE	Circuit state changes	Circuit	Yes	Yes	Yes	No
PKTLOSS	Tunnel QoS packet loss	Tunnel QOS	Yes	Yes	Yes	No
CIR_PKTLOSS	Circuit FC packet loss	Circuit	Yes	Yes	Yes	No
CIR_UTIL	Circuit FC utilization	Circuit	Yes	Yes	Yes	No
IP_EXTN_FLOW	LAN IP extension flow per DP	DP	No	No	No	Yes
TUNNEL_STATE	Tunnel state changes	Tunnel	Yes	Yes	Yes	No
TUNNEL_UTIL	Tunnel utilization	Tunnel	Yes	Yes	Yes	No
QOS_UTIL	QoS utilization percentage	Tunnel QOS	Yes	Yes	Yes	No
CIR_QOS_UTIL	Circuit QoS utilization percentage	Circuit QOS	Yes	Yes	Yes	No
CIR_QOS_PKTLOSS	Circuit QoS packet loss percentage	Circuit QOS	Yes	Yes	Yes	No
TUNNEL_IP_UTIL	Tunnel IP utilization percentage	Tunnel	Yes	Yes	Yes	No
CIR_IP_UTIL	Circuit IP utilization percentage	Circuit	Yes	Yes	Yes	No
IP_JITTER	IP Circuit FC connection variance	Circuit	No	No	No	Yes

Table 26: Fabric Performance Impact: Monitoring System and Supported Timebase

Parameter Used In Rule Command	Monitoring System	Group	Day	Hour	Minute	None
IT_FLOW	IT flow ration	Pid	Yes	Yes	Yes	No
RX	Receive bandwidth usage percentage	Port	Yes	Yes	Yes	No
TX	Transmit bandwidth usage percentage	Port	Yes	Yes	Yes	No
DEV_LATENCY_IMPACT	Device latency impact	Port	Yes	Yes	Yes	Yes
BE_LATENCY_IMPACT	Backend port latency impact	BE Port	No	No	No	Yes
UTIL	Trunk utilization	Port	Yes	Yes	Yes	No
DEV_LOGIN_DIST	Device logins distribution	Port	No	No	No	Yes

Table 27: Traffic Performance: Monitoring System and Supported Timebase

Parameter Used In Rule Command	Monitoring System	Group	Day	Hour	Minute	None
IO_RD	I/O read command count	Flow	Yes	Yes	Yes	No
IO_RD_BYTES	I/O read data	Flow	Yes	Yes	Yes	No
IO_WR	I/O write command count	Flow	Yes	Yes	Yes	No
IO_WR_BYTES	I/O write data	Flow	Yes	Yes	Yes	No

Parameter Used In Rule Command	Monitoring System	Group	Day	Hour	Minute	None
RX_FCNT	Receive frame count	Flow	Yes	Yes	Yes	No
RX_THPUT	Receive throughput	Flow	Yes	Yes	Yes	No
TX_FCNT	Transmit frame count	Flow	Yes	Yes	Yes	No
TX_THPUT	Transmit throughput	Flow	Yes	Yes	Yes	No
RD_STATUS_TIME_LT_8K	Time for read I/O request less than 8K	Flow	Yes	Yes	Yes	No
RD_STATUS_TIME_8_64K	Time for read I/O request greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No
RD_STATUS_TIME_64_512K	Time for read I/O request greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No
RD_STATUS_TIME_GE_512K	Time for read I/O request greater than or equal to 512K	Flow	Yes	Yes	Yes	No
WR_STATUS_TIME_LT_8K	Time for write I/O request less than 8K	Flow	Yes	Yes	Yes	No
WR_STATUS_TIME_8_64K	Time for write I/O request greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No
WR_STATUS_TIME_64_512K	Time for write I/O request greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No
WR_STATUS_TIME_GE_512K	Time for write I/O request greater than or equal to 512K	Flow	Yes	Yes	Yes	No
RD_1stDATA_TIME_LT_8K	Time for first data read less than 8K	Flow	Yes	Yes	Yes	No
RD_1stDATA_TIME_8_64K	Time for first data read greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No
RD_1stDATA_TIME_64_512K	Time for first data read greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No
RD_1stDATA_TIME_GE_512K	Time for first data read greater than or equal to 512K	Flow	Yes	Yes	Yes	No
WR_1stXFER_RDY_LT_8K	First data transfer in the ready state less than 8K	Flow	Yes	Yes	Yes	No
WR_1stXFER_RDY_8_64K	First data transfer in ready state greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No
WR_1stXFER_RDY_64_512K	First data transfer in ready state greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No

Parameter Used In Rule Command	Monitoring System	Group	Day	Hour	Minute	None
WR_1stXFER_RDY_GE_512K	First data transfer in ready state greater than or equal to 512K	Flow	Yes	Yes	Yes	No
RD_PENDING_IO_LT_8K	Pending read I/O requests less than 8K	Flow	Yes	Yes	Yes	No
RD_PENDING_IO_8_64K	Pending read I/O requests greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No
RD_PENDING_IO_64_512K	Pending read I/O requests greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No
RD_PENDING_IO_GE_512K	Pending read I/O requests greater than or equal to 512K	Flow	Yes	Yes	Yes	No
WR_PENDING_IO_LT_8K	Pending write I/O requests less than 8K	Flow	Yes	Yes	Yes	No
WR_PENDING_IO_8_64K	Pending write I/O requests greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No
WR_PENDING_IO_64_512K	Pending write I/O requests greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No
WR_PENDING_IO_GE_512K	Pending write I/O requests greater than or equal to 512K	Flow	Yes	Yes	Yes	No
RD_IO_RATE_LT_8K	Read I/O bytes less than 8K	Flow	Yes	Yes	Yes	No
RD_IO_RATE_8_64K	Read I/O bytes greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No
RD_IO_RATE_64_512K	Read I/O bytes greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No
RD_IO_RATE_GE_512K	Read I/O bytes greater than or equal to 512K	Flow	Yes	Yes	Yes	No
WR_IO_RATE_LT_8K	Written I/O bytes less than 8K	Flow	Yes	Yes	Yes	No
WR_IO_RATE_8_64K	Written I/O bytes greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No
WR_IO_RATE_64_512K	Written I/O bytes greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No
WR_IO_RATE_GE_512K	Written I/O bytes greater than or equal to 512K	Flow	Yes	Yes	Yes	No
RD_IOPS_LT_8K	Read I/O count less than 8K	Flow	Yes	Yes	Yes	No
RD_IOPS_8_64K	Read I/O count greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No

Parameter Used In Rule Command	Monitoring System	Group	Day	Hour	Minute	None
RD_IOPS_64_512K	Read I/O count greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No
RD_IOPS_GE_512K	Read I/O count greater than or equal to 512K	Flow	Yes	Yes	Yes	No
WR_IOPS_LT_8K	Written I/O count less than 8K	Flow	Yes	Yes	Yes	No
WR_IOPS_8_64K	Written I/O count greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No
WR_IOPS_64_512K	Written I/O count greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No
WR_IOPS_GE_512K	Written I/O count greater than or equal to 512K	Flow	Yes	Yes	Yes	No

Table 28: Security Violation: Monitoring System and Supported Timebase

Parameter Used In Rule Command	Monitoring System	Group	Day	Hour	Minute	None
DAYS_TO_EXPIRE	Days to expire	Certificate	No	No	No	Yes
EXPIRED_CERTS	Expired certificates	CHASSIS	No	No	No	Yes
SEC_AUTH_FAIL	SLAP failures	SWITCH	Yes	Yes	Yes	No
SEC_CERT	Invalid certifications	SWITCH	Yes	Yes	Yes	No
SEC_CMD	Illegal command	SWITCH	Yes	Yes	Yes	No
SEC_DCC	DCC	SWITCH	Yes	Yes	Yes	No
SEC_FCS	No FCS	SWITCH	Yes	Yes	Yes	No
SEC_HTTP	HTTP	SWITCH	Yes	Yes	Yes	No
SEC_IDB	Incompatible security database	SWITCH	Yes	Yes	Yes	No
SEC_LV	Login violations	SWITCH	Yes	Yes	Yes	No
SEC_SCC	SCC violations	SWITCH	Yes	Yes	Yes	No
SEC_TELNET	Telnet violations	SWITCH	Yes	Yes	Yes	No
SEC_TS	TS out of sync	SWITCH	Yes	Yes	Yes	No

Table 29: Fabric State Changes: Monitoring System and Supported Timebase

Parameter Used In Rule Command	Monitoring System	Group	Day	Hour	Minute	None
BB_FCR_CNT	FCR count	SWITCH	No	No	No	Yes
DID_CHG	Domain ID change	SWITCH	Yes	Yes	Yes	No
EPORT_DOWN	E_Ports down	SWITCH	Yes	Yes	Yes	No
FAB_CFG	Fabric reconfigurations	SWITCH	Yes	Yes	Yes	No

Parameter Used In Rule Command	Monitoring System	Group	Day	Hour	Minute	None
FAB_SEG	Segmentation	SWITCH	Yes	Yes	Yes	No
FLOGI	FLOGI	SWITCH	Yes	Yes	Yes	No
L2_DEVCNT_PER	L2 device count percentage	SWITCH	Yes	Yes	Yes	No
LSAN_DEVCNT_PER	LSAN device count percentage	SWITCH	No	No	No	Yes
ZONE_CFGSZ_PER	Zone configuration size percentage	SWITCH	No	No	No	Yes
ZONE_CHG	Zone changes	SWITCH	Yes	Yes	Yes	No

Table 30: Switch Resources: Monitoring System and Supported Timebase

Parameter Used In Rule Command	Monitoring System	Group	Day	Hour	Minute	Second	None
TEMP	Temperature sensor	Temperature sensor	No	No	No	No	Yes
CPU	CPU usage	CHASSIS	No	No	No	No	Yes
ETH_MGMT_PORT_STATE	Ethernet management port state	CHASSIS	No	No	No	No	Yes
FLASH_USAGE	Flash usage	CHASSIS	No	No	No	No	Yes
MEMORY_USAGE	Memory usage	CHASSIS	No	No	No	No	Yes
VTAP_IOPS	VTAP I/Os per second	Asic	No	No	No	Yes	No

Table 31: Switch Status Policy: Monitoring System and Supported Timebase

Parameter Used In Rule Command	Monitoring System	Group	Day	Hour	Minute	None
BAD_FAN	Absent or faulty fans	CHASSIS	No	No	No	Yes
BAD_PWR	Absent or faulty power supplies	CHASSIS	No	No	No	Yes
BAD_TEMP	Temperature sensors outside the range	CHASSIS	No	No	No	Yes
DOWN_CORE	Core blade monitoring	CHASSIS	No	No	No	Yes
ERR_PORTS	Percentage of error ports	SWITCH	No	No	No	Yes
FAN_AIRFLOW_MISMATCH	Mismatch between PSU fan airflow directions	CHASSIS	No	No	No	Yes
FAULTY_BLADE	Faulty blades	CHASSIS	No	No	No	Yes
FAULTY_PORTS	Percentage of faulty ports	SWITCH	No	No	No	Yes
FLASH_USAGE	Flash usage	CHASSIS	No	No	No	Yes
HA_SYNC	HA monitoring	CHASSIS	No	No	No	Yes
MARG_PORTS	Percentage of marginal ports	SWITCH	No	No	No	Yes
MISSING_SFP	Percentage of missing SFPs	SWITCH	No	No	No	Yes
WWN_DOWN	WWN	CHASSIS	No	No	No	Yes

Table 32: FRU Health: Monitoring System and Supported Timebase

Parameter Used In Rule Command	Monitoring System	Group	Day	Hour	Minute	None
BLADE_STATE	Blade	Blade	No	No	No	Yes
FAN_STATE	Fan	FAN	No	No	No	Yes
PS_STATE	Power supply	Power supply	No	No	No	Yes
SFP_STATE	SFP	FC PORT and FCoE Port	No	No	No	Yes
WWN	WWN	WWN	No	No	No	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. MAPS rules can exist outside of a MAPS policy, but they are only evaluated by MAPS when the rule is part of an active policy.

Each rule specifies the following items:

- A group of objects to be evaluated. See MAPS Groups Overview for additional information.
- The condition being monitored. Each rule specifies a single condition. A condition includes a timebase and a threshold. See MAPS Conditions for additional information.
- The actions to take if the condition is evaluated to be true. See MAPS Actions for additional information.

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

NOTE
No user operations are allowed on the rules for backend ports. These are monitored by ALL_BE_PORTS.

Rules That Monitor Other Rules (RoR Rules)

You can create rules that monitor the performance of other rules. The monitoring rule is called a rule-on-rule (RoR) rule.

Using RoR rules, you can monitor the execution of a rule. The RoR rule performs different actions based on the performance of the rule it was monitoring.

A RoR rule specifies the following items:

- The base rule that the RoR rule is monitoring.
- A group of objects to be evaluated (optional). See MAPS Groups Overview for additional information.
- The condition being monitored. Each RoR rule specifies a single condition. A condition includes a timebase and a threshold. See MAPS Conditions and Supported Timebases for RoR Rules When the Timebase for the Base Rule is NONE for additional information.
- The actions to take if the condition is evaluated to be true. See MAPS Actions for additional information.

When creating RoR rules, you can use all of the existing MAPS features including:

- Rule management
- Policy management
- Dashboard
- Alerts

See Creating Rules to Monitor Other Rules (RoR) for details.

MAPS Rule Actions

When you create, modify, or clone a rule using the `mapsrule --create` , `mapsrule --config` , or `mapsrule --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
- FICON Alerts
- MAPS SNMP Traps
- Port Fencing
- RASLog Messages
- SFP Marginal
- Slow Drain Device Quarantine
- Switch Critical
- Switch Marginal
- Port Toggling

For each action, you can define a *quiet time* for most rules to reduce the number of alert messages generated. See Quieting a Rule for details.

For rules that are state-bound (such as those for temperature sensing monitoring 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, thenif the CRC threshold is reached, a RASLog message is issued, but the port would not be fenced. To enable global actions, use the `mapsconfig --actions` command. For more details, see Enabling or Disabling Actions at a Global Level. Refer to the *Brocade Fabric OS Command Reference Manual* 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 allow 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 7.2.0 and later releases, 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, 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, decom, and fence notifications are not currently active actions on the switch, and then shows them being added to the list of allowed actions.

    NOTE
    SW_CRITICAL, SW_MARGINAL, and SFP_MARGINAL notifications are enabled by default in a switch, and you cannot disable them.

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

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

```
switch:admin> mapsconfig --show
Configured Notifications:       RASLOG,EMAIL,DECOM,FENCE,SW_CRITICAL,SW_MARGINAL,SFP_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 address and each address must be a complete e-mail address. For example, abc@company.com is a valid e-mail address; abc@company is not. See Sending Alerts Using Email 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

For Fabric OS 8.0.x and later versions, e-mail alerts are enhanced to include information that is more meaningful, including the monitor port name and rephrasing other content to help you understand the error condition or violation in the switch and take action accordingly.

Example of E-Mail Alert Enhancements

The following examples show the enhancements for e-mail alerts from threshold-based rules and state-based rules. The enhanced information is labeled *Subject*, *Group*, and *Current Value* (enhancements are shown in boldface type):

Table 33: 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_PORTSCurrent 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 34: 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_PORTSCurrent 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 *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 the 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 may 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 can be a most efficient notification method, because you can avoid having to log in to each switch individually as you do for error log notifications.

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 *Brocade Fabric OS Command Reference Manual*.

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

- 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 action for rules that monitor changes of state (STATE_CHG). Refer to the Port Health Monitoring Default Thresholds 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. This is because configuring 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, see Port Health Monitoring Default Thresholds.

For F_Ports, port decommissioning will only work if Brocade Network Advisor is actively monitoring the switch. Brocade Network Advisor can decommission F_Ports based on specified criteria (see Port Health Monitoring Default Thresholds.) MAPS notifications are integrated with Brocade Network Advisor, which in turn must coordinate with the switch and the end device to orchestrate the port decommissioning. If Brocade Network Advisor is not configured on a switch, MAPS fences the F_Port.

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

Configuring Port Decommissioning

Port decommissioning is a two-part process. You can configure port decommissioning along with port fencing in the MAPS actions configuration and can 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 Downgrade

- 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 software. That is, a default MAPS rule from the current version of Fabric OS software 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 software.
- A firmware downgrade will be blocked if any user-defined rule is configured using the *decom* action.

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 occurs any time the port fails, while if it is part of a rule, the port is fenced if that rule is triggered. Multiple rules can have port fencing as an action; it occurs 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 is not resolved, the port toggling action suspends the port for a longer period of time, forcing the port traffic to switch over to a different path if available.

> **NOTE**
> The port toggling (PT) action and the SDDQ action are mutually exclusive. When using the `mapsconfig` command, you cannot enable the SDDQ and PT actions at the same time.

See MAPS Support for Port Toggling 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.

> **NOTE**
> The Master port in the trunk group would display as E-port, and for slave ports as T-Port.

Table 35: 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 36: 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 and later releases, 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 *Brocade Fabric OS Message Reference Manual* for a complete listing and explanation of MAPS-related RASLog messages.

Custom RASLog IDs

Starting with Fabric OS 8.2.x release, you can use the custom RASLog IDs (starting from MAPS-2000 through MAPS-2971) so that the third-party applications do not need to parse the RASLog message content to determine the monitoring statistics such as CRC or ITW, and thus automate the response based on the RASLog IDs. Each type of RASLog message supports up to four IDs based on severity levels such as *Critical*, *Error*, *Warning*, and *Info*. Refer to the *Brocade Fabric OS Message Reference Manual* for the complete list of new RASLog message IDs.

For example, the CRC RASLog messages use the following IDs:

- MAPS-2004 for Critical
- MAPS-2005 for Error
- MAPS-2006 for Warning
- MAPS-2007 for Info

For example, the ITW RASLog messages use the following IDs:

- MAPS-2008 for Critical
- MAPS-2009 for Error
- MAPS-2010 for Warning
- MAPS-2011 for Info

To enable the customer RASLog IDs, run the `mapsconfig --raslogmode custom` command. To switch back to the legacy RASLog mode, run the `mapsconfig --raslogmode default` command. When custom RASLog IDs are enabled, if you downgrade to pre-8.2.x version of Fabric OS software, an error is generated.

Severity Configuration

Starting with Fabric OS 8.1.0 release, you can configure the severity of the actions in a rule.

By configuring the severity of actions on rules, you can ensure that certain actions occur that correspond to the way your company filters alerts according to its policies.

The new syntax for the `mapsrule` command includes an option to configure severity when creating, modifying, or cloning a rule. The following commands are affected:

- `mapsrule --create`
- `mapsrule --config`
- `mapsrule --clone`
- `mapsrule --show`

The syntax for the new option is:

```
-severity level | default
```

where the level can be defined as *critical*, *error*, *warning*, *info*, or *default* indicates the alert uses default severities as defined in the *Default Severity* table.

The severity specification affects all alerts and appears in the body of the email alert.

NOTE
You cannot change the severity of an alert that notifies you about the health of a switch.

Configuring the Severity of an Alert

The following example demonstrates using the `-severity` option to create a rule and configure alerts to use severity of *critical*:

```
switch:admin> mapsrule --create newrule1 -group GRP702 -monitor MON876598
```

```
-timebase min -value 1 -action RASLog,email,SNMP
-severity critical
```

SFP Marginal

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

NOTE
This action is valid only in the context of Advanced SFP groups. This action is enabled only on the 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 is 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 is 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 is triggered for Port 03 and Port 19. Notice that the health status of 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 it 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.

See Slow-Drain Device Quarantining for a complete discussion of this action.

Switch Critical

The *switch critical* action sets the state of the affected switch on the MAPS dashboard display to *SW_CRITICAL*. This action does not bring the switch down, but only affects what is displayed on 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 on 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.

Working with MAPS Rules and Actions

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

Viewing MAPS Rule

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

Perform the following steps to view the MAPS rules on a switch:

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 the 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,UVTAP |
defALL_D_PORTSCRC_1            |ALL_D_PORTS(CRC/MIN>1)               |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   UVTAP: UNINSTALL_VTAP
```

Creating a Rule

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

When creating rules, be aware of the following:

- Rule names are case sensitive; *My_Rule* and *my_rule* are not considered to be the same.
- Rule names can be up to 72 characters long.
- A logical switch can have a maximum number of 500 user-defined rules.

To create a rule, perform the following steps:

1. Enter `mapsrule --create <rule_name>` followed by the rule parameters, and optionally the policy you want to assign it to.
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

   - You can only add a rule to an existing policy.
   - If you specify 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> mapsconfig --import io_mon_4

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
```

Creating Rules to Monitor Other Rules (RoR Rules)

You can create rules that monitor the performance of other rules. The monitoring rule is called a rule-on-rule (RoR) rule.

To create a RoR rule to monitor a base rule, perform the following steps:

1. Enter `mapsrule --createRoR <RoR_rule_name> [-group <group_name> -monitor <name_of_rule_to_monitor>` followed by the RoR rule parameters, and optionally the policy you want to assign it to. RoR rule names are not case sensitive; My_Rule and my_rule are considered to be the same.

   NOTE
   The group name is optional. If you do not specify a group name, the group for the base rule is used.

2. Optional: Enter `mapsrule --show <rule_name>` to display the rule.
3. Optional: If you added the RoR rule to an active policy, you must re-enable the policy for the RoR rule to take effect by entering `mapspolicy --enable policy <policy_name>`.

   NOTE
   You can only add a RoR rule to an existing policy. If you specify a group (optional), the group must already exist.

Examples of using a RoR rule to monitor rule performance

A typical rule you might have created would monitor CRC per minute. You could create a RoR rule that monitors how often the base rule is executed. In this example, the base rule triggers if the value of CRC is greater than 12 per minute. The RoR rule is triggered and sends alerts if the base rule is triggered more than five times in an hour.

```
switch:admin> mapspolicy --create test_policy1
switch:admin> mapsgroup --create test_ports_1
switch:admin> mapsrule --create test_CRC_min_1 -group ALL_PORTS -monitor CRC
                          -timebase min -op g -value 12 -action raslog,snmp,email
                          -policy test_policy1

switch:admin> mapsrule --createRoR test_ror_CRC_min_1 -group test_ports_1 -monitor test_CRC_min_1
                          -timebase hour -op g -value 5 -action raslog,snmp,email,fence
                          -policy test_policy1
```

In the following example, the group option is not specified. The RoR rule uses the base rule's group, ALL_PORTS as the scope:

```
switch:admin> mapspolicy --create test_policy2
switch:admin> mapsrule --create test_CRC_min_2 -group ALL_PORTS -monitor CRC
                          -timebase min -op g -value 12 -action raslog,snmp,email
                          -policy test_policy2

switch:admin> mapsrule --createRoR test_ror_CRC_min_2 -monitor test_CRC_min_2
                         -timebase hour -op g -value 5 -action raslog,snmp,email,fence
                         -policy test_policy2
```

Restrictions on Using RoR Rules

When using rule-on-rules (RoR rules) to monitor other rules (base rules), note the following restrictions:

Restrictions on Creating RoR Rules

When you create RoR rules, the following restrictions apply:

- MAPS supports a maximum of 50 RoR rules for each policy.
- You can create a RoR rule only if the base rule is present on the switch.
- You can add a RoR rule to a policy only if the base rule is present and has been added to that policy.
- You cannot specify the FMS action for a RoR rule.
- Chassis monitoring systems are monitored only in the default switch context.
- You can define the timebase of a RoR with these restrictions:
  - The timebase of the RoR rule cannot be *None*.
  - The timebase of the RoR rule must be greater than the timebase of the base rule.
  - If the timebase of the base rule is *None*, the timebase of the RoR rule is defined as shown in Supported Timebases for RoR Rules When the Timebase of the Base Rule is NONE.
- The following monitoring systems are ***not*** supported for RoR rules:
  - Fabric state change monitors (all of them).
  - Fabric Performance Impact monitor: IT_FLOW.
  - Flow monitors (all of them).
  - Security monitors: DAYS_TO_EXPIRE and EXPIRED_CERTS.
  - All switch resource monitors except CPU and MEMORY_USAGE.
  - Switch status policy monitors (all of them).
  - Any monitor for which a rule can be configured with SW_MARGINAL or SW_CRITICAL.
- For chassis monitors (CPU, MEMORY, and DOWN_CORE), you can create RoR rules in only the default switch.
- You cannot create a RoR rule to monitor another RoR rule.

Restrictions that Affect the Modification of Base Rules

If you have created one or more RoR rules to monitor a base rule, then the following restrictions apply when you want to update or configure the base rule:

- If you modify the threshold or operator values of the base rule, the rule's trigger count is reset, and monitoring by the RoR rule is restarted.
- All other parameters of the base rule can be updated without affecting the monitoring by the RoR rule.
- If you modify the timebase value in the base rule more than the timebase in the monitoring RoR rule, then the RoR rule changes to unmonitor (*) state.

Supported Timebase for RoR Rules when the Timebase of the Base Rule is *None*

When defining a RoR rule that monitors a base rule that has a timebase of *None*, the timebase for the RoR rule can be defined as shown in the following table:

Table 37: Supported Timebase for RoR Rules when the Timebase of the Base Rule is *None*

RoR Rules	Timebase
BE_LATENCY_IMPACT	Hour, Day, Week
BLADE_STATE	Min, Hour, Day, Week
CPU	Hour, Day, Week
CURRENT	Hour, Day, Week
DEV_LATENCY_IMPACT	Hour, Day, Week
DEV_LOGIN_DIST	Hour, Day, Week
DEV_NPIV_LOGINS	Day, Week

RoR Rules	Timebase
ENCR_DISC	Hour, Day, Week
FAN_STATE	Min, Hour, Day, Week
IP_EXTN_FLOW	Hour, Day, Week
IP_JITTER	Hour, Day, Week
IP_RTT	Hour, Day, Week
JITTER	Hour, Day, Week
MEMORY_USAGE	Hour, Day, Week
PWR_HRS	Hours, Day, Week
PS_STATE	Min, Hour, Day, Week
RTT	Hour, Day, Week
RXP	Hour, Day, Week
SFP_STATE	Min, Hour, Day, Week
SFP_TEMP	Hour, Day, Week
TXP	Hour, Day, Week
VOLTAGE	Hour, Day, Week
WWN	Min, Hour, Day, Week

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, perform 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
```

NOTE
If you are using rule-on-rule (RoR) rules, you should be aware of restrictions affecting your ability to modify rules. See Restrictions that Affect the Modification of Rules.

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>`* *`<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 *Brocade Fabric OS Command Reference Manual*.

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:
```

Cloning a Group of Rules

Instead of cloning one rule at a time, you can clone an entire group of rules with one command.

When cloning a group of rules, be aware of the following:

- You clone all the rules in an existing default or user-defined group into a new group.
- You can specify that the rules be included in a policy:
  - The policy must already exist.
  - If rules already exist in the policy, the cloned rules will still be added.
- Optionally, you can specify an extension (a *tag*) that will be appended to the name of each cloned rule. If you do not specify a tag, the names of all of the cloned rules will be prefixed with *clone_*.

To clone a group of MAPS rules, perform the following steps:

1. Create or clone a new group, and create or clone a new policy (or use an existing policy).

   Refer to Creating a Static User-Defined Group, Cloning a Group, and Working With MAPS Policies for details.

2. Use the following command to clone the group of rules:

   `mapsrule --cloneByGroup <existingGroupName> -newGroup <newGroupName> --fromPolicy <existingPolicyName> -newPolicy <newPolicyName>.`

3. Optionally, you can append a *tag* to the name of each cloned rule. To add a tag, include the following parameter in the command above: `-tag <tagText>` .

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

Creating a New Group of Cloned Rules

The following is an example of cloning all the rules in a group. It shows the creation of a group and a policy within that group. The example shows adding the rules to an existing policy in a new group and appending a *tag* to the name of all the cloned rules.

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

switch:admin>  mapspolicy --create myPorts_aggressive_policy

switch:admin>  mapsrule --cloneByGroup ALL_E_PORTS -newGroup MY_CRITICAL_PORTS
               -fromPolicy dflt_aggressive_policy  -newPolicy myPorts_aggressive_policy
               -tag _myPorts
```

When you issue these commands, perform the following:

- Create a new group called *MY_CRITICAL_PORTS*.
- Create a policy called *myPorts_aggressive_policy*.
- Clone all the rules in group *ALL_E_PORTS* from the *dflt_aggressive_policy* to *myPorts_aggressive_policy* based on matching criteria based on the new group *MY_CRITICAL_PORTS* and add the cloned rules to new policy *myPorts_aggressive_policy*.
- Append the tag *_myPorts* to all the cloned rules.

NOTE
If you specify a policy with the `-newPolicy` parameter, it must be an existing user-defined policy or the command will fail. The new policy cannot be a default policy.

Quieting a Rule

MAPS supports the concept of *quiet time* as an optional rule parameter for the `mapsrule` command. Including `-qt <seconds>` in a rule keeps MAPS from sending another alert based on the same rule for the specified number of seconds after it sends the initial alert.

NOTE
The quiet time 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 in-box.

MAPS allows you to set quiet time for a rule. The following example shows a rule that includes a 120-second quiet time period. The quiet time element is highlighted in boldface.

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

When a rule for which quiet time has been set is triggered, MAPS performs the configured and enabled actions for the rule the first time. Afterward, if the rule is triggered again within the quiet time period, MAPS does not perform any of the actions until the quiet time expires.

NOTE
The rules get evaluated every one minute. For example, when you configures quiet time value as 620 (10 minutes 20 seconds), that rule is not triggered for the next 10 minutes as the quiet time value does not expire yet while in the 10th minute. The rule is triggered only on the next rule evaluation on the 11th minute where the quiet time value has expired.

MAPS allows you to configure quiet time with unsupported actions, such as port fencing or port decommissioning; however, MAPS ignores the quiet time parameters.

Quiet Time Expiry Action

When quiet time expires, MAPS sends an updated alert and generates an SNMP trap for any rule violation. 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:

```
2016/01/26-01:56:00, [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 Tues Jan 26 01:53:24 2016,
Dashboard Category=Port Health
```

NOTE
The SNMP trap is automatically generated when quiet time expires. You can use the `mibcapability` command to disable the trap if it is not required during setup.

The command `mapsTrapAM` (*AM* represents *Alerting and Monitoring*) 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.

```
mapsTrapAM NOTIFICATION-TYPE
    OBJECTS     {
           mapsConfigRuleName,
                 ….
```

```
  mapsRuleTriggerCount,
  mapsLastRuleExecTime,
  mapsQuietTime
    }
    STATUS current
    DESCRIPTION     "Trap to be sent for MAPS threshold events."
    ::= { mapsTraps 1 }


mapsRuleTriggerCount        OBJECT-TYPE
    SYNTAX          Unsigned32
    DESCRIPTION     "Number of times rule was triggered in quiet time."
    ::= { mapsConfig 15 }

mapsLastRuleExecTime        OBJECT-TYPE
    SYNTAX          DateAndTime
    DESCRIPTION     "Last rule's execution time."
    ::= { mapsConfig 16 }

mapsQuietTime              OBJECT-TYPE
    SYNTAX          Unsigned32
    DESCRIPTION     "Quiet time configured in the rule."
    ::= { mapsConfig 17 }
```

Quiet Time and Timebase

The length of a quiet time period is specified in seconds, with the minimum value determined by the timebase defined in the rule for all timebases other than *NONE*. For example, if the rule has *minute* specified as the timebase, the minimum value for quiet time is 60 seconds (one minute). There is no predefined upper limit.

For all timebases other than *NONE*, the minimum quiet time is determined by the timebase present in the rule. However, if the rule's timebase is *NONE*, then the quiet time value is different for different monitoring `systems` , as shown in the following table:

Table 38: Minimum Quiet Time Values for Monitoring Systems that Support a Timebase of NONE

Monitoring System	Minimum Quiet Time (Seconds)	Maximum Quiet Time (Seconds)
AVG_PENDING_IOS	10	
AVG_ROS	10	
BAD_FAN	60	31536000 (365 days)
BAD_PWR	60	
BAD_TEMP	60	
BB_FCR_CNT	86400	
BE_LATENCY_IMPACT	60	
BLADE_STATE	15	
CPU	120	
CURRENT	720	
DAYS_TO_EXPIRE	86400	

Monitoring System	Minimum Quiet Time (Seconds)	Maximum Quiet Time (Seconds)
DEV_LATENCY_IMPACT	60	
DEV_NPIV_LOGINS	3600	
DOWN_CORE	60	
ENCR_DISC	60	
ERR_PORTS	60	
ETH_MGMT_PORT_STATE	60	
EXPIRED_CERTS	86400	
FAN_AIRFLOW_MISMATCH	6	
FAULTY_BLADE	60	
FAULTY_PORTS	60	
FAN_STATE	15	
FLASH_USAGE	60	
FLOW_COUNT	60	
HA_SYNC	60	
IP_EXTN_FLOW	10	
IP_JITTER	60	
IP_RTT	60	
IT_COUNT	60	
ITL_COUNT	60	
JITTER	60	
L2_DEVCNT_PER	86400	
LSAN_DEVCNT_PER	86400	
MARG_PORTS	60	
MAX_PENDING_IOS	10	
MAX_ROS	10	
MEMORY_USAGE	120	
MISSING_SFP	60	
FAN_STATE	15	
PS_STATE	15	
PWR_HRS	720	
RTT	60	
RXP	720	
SFP_STATE	15	
SFP_TEMP	720	
TEMP	60	
TXP	720	
VOLTAGE	720	

Monitoring System	Minimum Quiet Time (Seconds)	Maximum Quiet Time (Seconds)
WWN_DOWN	60	
WWN	15	
ZONE_CFGSZ_PER	86400	

Clearing Quiet Time Settings

You can clear the quiet time settings in a rule using the `-qtclear` option of the `mapsrule` command. For example, to clear the quiet time settings for the rule used in the first example in this section, enter the following:

```
switch:admin>    mapsrule --config toggle_crc_rule -qtclear
```

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 may require a lot of tedious work if the rule is 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 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 is 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 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. See Configuring E-Mail Server Information 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 Email Alert Testing 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,SFP_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,SFP_MARGINAL
Mail Recipient:                 admin1@mycompany.com,
                                admin2@mycompany.com
Paused members :
PORT :
CIRCUIT :
SFP :
```

Specifying the From E-Mail Address for Alerts

The following example specifies e-mail addresses for e-mail alerts with a From email address. It assumes that you have already correctly configured and validated the e-mail server. If you do not specify a From email address, the default From email address (*`swtich_name@domain.com`*) is used.

```
switch:admin> mapsconfig --emailcfg -address admin1@mycompany.com  -from admin@mycompany.com
switch:admin> mapsconfig --show
Configured Notifications:       RASLOG,EMAIL,FENCE,SW_CRITICAL,SFP_MARGINAL
Mail Recipient:                 admin1@mycompany.com
From Address                    admin@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 39: 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 *Brocade Fabric OS Command Reference Manual*.

Configuring E-Mail Server Information

Fabric OS software 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.

The relay host is a smart relay server which is used to filter e-mail messages coming from the 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, perform the following steps:

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

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.my-company.com"

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

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

Viewing Configured E-Mail Server Information

Fabric OS software 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, perform 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:                     10.01.02.03
   Relay Domain Name:           mail.my-company.com
   ```

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

Deleting E-Mail Server Configuration

Fabric OS software 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, perform 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 are 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 *Brocade Fabric OS Command Reference Manual*.

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 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.

    NOTE
    It is recommended that you avoid using policies that have no rules.
  - 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|
```

Related Links

Predefined Policies on page 81

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

User-Defined Policies on page 82

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.

Creating a Policy on page 85

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

Modifying a Policy on page 86

You can modify existing user-defined policies. For example, you might need to modify a policy if elements in the fabric change or if threshold configurations need to be modified to catch certain error conditions.

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 Feature Monitors Not Requiring a License for details.

This policy contains rules that monitor the unlicensed features which were made available earlier through Fabric Watch. Refer to Unlicenced Feature Monitoring 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:

- User-Defined Policies
- Creating a Policy
- Modifying a Default Policy

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 policy dflt_aggressive_policy. For Open Systems environments and other environments, Brocade recommends that you start with policy dflt_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.

See Working with MAPS Policies 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 used `mapsConfig --fwconvert` . If you do not use this command, 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 is 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.

Changes to Threshold Behaviors After Conversion

MAPS has inherent design differences from Fabric Watch when it handles thresholds. During conversion, the following changes occur in threshold behavior:

- **Above-High** is converted for all the thresholds because they always indicate an error.
- **Above-Low** is converted for port-related classes: PORT, E_PORT, F_PORT, and so on.
- **Below-Low** is converted for SFP parameters, where a Below-Low event for voltage or current indicates an error state.
- **Below-High** is ignored during conversion for all classes.

NOTE
If you still want to configure any unconverted rules, given the above behavioral changes, you can create rules manually using the `--op` option of the `mapsrule` command to add them to the policy.

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.

Perform the following steps to view the MAPS policies on a switch:

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

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

The following restrictions apply to policies:

- Policy names are not case-sensitive; *My_Policy* and *my_policy* are considered to be the same.
- When you create a policy, the policy is automatically saved, but not enabled. The policy is not enabled unless you explicitly enable it.
- A logical switch can have a maximum of 20 user-defined policies.
- A user-defined policy can have a maximum of 350 rules per logical switch.
- The maximum length of a policy name is 31 characters.

Perform the following steps to create policies and add rules to them:

1. To create a new policy, enter `mapspolicy --create <policy_name>`.
2. Create, clone or modify the rules for that policy to configure the required thresholds in the new policy. See Working with MAPS Rules and Actions for details.

   The following example creates a policy and adds a rule to the new policy:

```
switch:admin> mapspolicy --create new_policy1
switch:admin> mapsrule --create CPUmon123 -monitor CPU -group chassis
                        -op ge -value 40 -action raslog -policy new_policy1
```

Cloning a Policy

Cloning a policy allows you to create a policy that is similar to another policy. After you clone a policy, you can modify it to meet your needs.

To clone policies and add rules to them, perform the following steps:

1. To clone an existing policy, enter `mapsPolicy --clone <policy_name> -name <clone_policy_name>`.
2. Create, clone or modify rules to configure the required thresholds in the new policy. See Working with MAPS Rules and Actions for details.

   The following example creates a policy by cloning another policy and 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 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, perform 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

You can modify existing user-defined policies. For example, you might need to modify a policy if elements in the fabric change or if threshold configurations need to be modified to catch certain error conditions.

Perform the following steps to modify a user-defined policy and its associated rules:

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: 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 be 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
```

Applying the Modified User-Defined Active Policy

If you do not know if the modified user-defined active policy is applied or not, check the `mapspolicy --show -summary` command output for the following message, which indicates that the active policy has been modified but not applied.

```
WARNING: The policy configuration for the active policy has changed because of addition/deletion/modification
 of rule. The active policy has to be re-enabled to apply the new changes.
```

You can also look for the "(+)" sign next to the active policy name in the in the dashboard, which indicates that the active policy has been modified but not applied.

```
switch#admin> mapsdb --show all
1 Dashboard Information:
=======================
DB start time:   Sat May 27 11:49:48 2017
Active policy:   test_pol1(+)
<output truncated>
(+): The active policy should be re-enabled to apply the new changes.
```

The active policy is considered to be modified if any of the following changes are made:

- If any of the parameters are changed in the active policy rules.
- If a rule is added to the active policy.

  NOTE
  Though the newly added rule does not get monitored, MAPS does not show a ' * ' in front of the rule name because the policy is not enabled to apply the new set of rules.
- If a rule is removed from the active policy.

  NOTE
  Even the deleted rule gets monitored because the policy is not enabled to apply the new set of rules.
- If any of the user-defined rule is deleted.

For the changes to the active user-defined policy to take effect, run the `mapspolicy --enable <policy_name>` command.

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, perform 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.

The 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_05, 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_05, 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.

MAPS Dashboard

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.

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 extracts 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 2018
Active policy:                     test_xy1
Configured Notifications:          SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,RASLOG,FENCE
Ports fenced by MAPS/BNA :         None
Ports decommissioned by MAPS :     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 health status and lists any factors contributing to that status as defined by the Switch Status Policy rules in the active policy.

The following example shows the Switch Health Report section; revealing that the switch status is CRITICAL.

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

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


-------------Output truncated-----------------
```

See Switch Status Policy 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, 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
- Back End Health
- FRU Health
- Security Violations
- Fabric State Changes
- Switch Resources
- Traffic Performance
- Fabric Performance Impact

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        |
Extension 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 are edited 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_     |02/04/18 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|02/04/18 20:21:01|Switch     |2 Ports     |
Changes (4)   |      |                     |                 |           |            |
              |      |                     |                 |Switch     |2 Ports     |
              |1     |defSWITCHEPORT_DOWN_1|02/04/18 19:17:12|Switch     |2 Ports     |
              |1     |defSWITCHFLOGI_4     |02/04/18 18:43:42|Switch     |7 Logins    |
Extesnion     |2     |defALL_TUNNELSSTATE_ |02/04/18 19:23:48|Tunnel 8/19|1           |
Health(2)     |      |CHG_0                |                 |           |            |
              |      |                     |                 |Tunnel 8/18|1           |
              |      |                     |                 |Tunnel 8/19|1           |
              |      |                     |                 |Tunnel 8/18|1           |
Fabric        |2     |defALL_PORTS_IO_     |02/04/18 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_     |02/04/18 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      |
                        --------------Output truncated----------
```

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.

The following output extract shows a sample History Data section:

```
(output truncated)
4 History Data:
===============
Stats(Units)        Current   07/21/14  07/21/14    --/--/--    --/--/--   --/--/--   --/--/--
                    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)    -          -          -          -          -          -          -

5 History Data for Backend ports:
=================================
Stats(Units)        Current   07/21/14   07/21/14   --/--/--   --/--/--   --/--/--   --/--/--
                    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 Extension health; for example, if no port 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, and 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 details` 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. You can use various options of the `mapsdb` command to display data gathered since midnight, for any one hour since midnight, or for the last seven days on which errors were recorded.

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

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.

Perform the following steps to view the status of the switch as seen by MAPS:

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 2018
Active policy:                  test_xy1
Configured Notifications:       SW_CRITICAL,SW_MARGINAL,SFP_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        |
Extension 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)   -           -           -           -           -           -           -
```

See MAPS Monitoring Categories 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, perform the following steps:

1. Connect to the switch and log on 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 *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. Note that the example column headings are edited to allow it to display clearly.

```
switch:admin> mapsdb --show


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


DB start time:                  Mon Jun 13 20:14:57 2016
Active policy:                  dflt_aggressive_policy
Configured Notifications:       RASLOG,SNMP,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,SDDQ
Fenced Ports :                  None
Decommissioned Ports :          None
Fenced circuits :               None
Quarantined Ports :             3/20,3/45,3/46,4/0,4/19,4/20
```

```
Top Zoned PIDs <pid(it-flows)>: 0x731400(21) 0x734c00(21) 0x734900(4) 0x735400(1) 0x735300(1)

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        |
Extension 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_   |06/13/16 20:14:57 |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 illustrates the violations of the *thruputflow_thput_10* rule. Some of the lines in the output are truncated 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 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 are edited slightly to display clearly.

   ```
   switch:admin> 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,SW_CRITICAL,SW_MARGINAL,SFP_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              |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        |
Extension 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 are 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.

Perform the following steps to view a summarized history of the switch status:

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.

Perform the following steps to view detailed historical data about a switch:

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 February 5, 2018, 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 are displayed. The column headings in the example are edited to display clearly.

```
switch:admin> mapsdb --show details -day 2/5/2018

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

DB start time:                Thu Feb 5 19:17:13 2018
Active policy:                dflt_aggressive_policy
Configured Notifications:     RASLOG,EMAIL,FENCE,SW_CRITICAL,SW_MARGINAL,SFP_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        |
Extension 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_       |02/05/18 07:15:02|Switch      |1 Violations |
Violations(2)|      |TELNET_0            |                 |            |             |
             |1     |defSWITCHSEC_LV_0   |02/05/18 07:15:02|Switch      |1 Violations |
Fabric       |1     |defALL_PORTS_       |02/05/18 06:53:02|E-Port 12/27|IO_LATENCY_  |
Performance  |      |LATENCY_CLEAR       |                 |            |CLEAR        |
Impact  (2)  |      |                    |                 |            |             |
             |1     |defALL_PORTS_IO_    |02/05/18 06:52:02|E-Port 12/27|IO_FRAME_LOSS|
             |      |FRAME_LOSS          |                 |            |             |

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

Stats(Units)        Current     02/04/18    --/--/--    --/--/--    --/--/--
-------------------------------------------------------------------------------

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     02/04/18    --/--/--    --/--/--    --/--/--
--------------------------------------------------------------------------------

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

6 History Data for Gig Ethernet ports:
========================================
Stats(Units)        Current     02/04/18    --/--/--    --/--/--    --/--/--
-------------------------------------------------------------------------------

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 is truncated to display the output correctly. Normally there may be additional days of data.

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

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

DB start time:                Thu Feb  4 19:17:13 2018
Active policy:                dflt_aggressive_policy
Configured Notifications:     RASLOG,EMAIL,FENCE,SW_CRITICAL,SW_MARGINAL,SFP_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        |
Extension 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 |02/04/18 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 |02/04/18 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|
```

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 that is 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.

Perform the following steps to clear the stored dashboard data from a switch:

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 *Brocade Fabric OS Message Reference Manual*.

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

The following example clears only the dashboard summary data:

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

Port Monitoring Using MAPS

Monitoring a Group 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*.

There are sets of ports that behave in a similar manner and have different behavior from other sets of ports on a switch. 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. You can create a group and apply rules to the group to easily monitor these similar sets of ports using the same rules.

Perform the following steps to create a group and apply rules to the group:

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 software allows you to monitor ports based on their assigned names.

The port name is an editable attribute of a port, and 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.

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

For more information on creating dynamic user-defined groups, see User-Defined Groups.

Port Monitoring Using Device WWNs

Fabric OS software 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, see User-Defined Groups.

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 the 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, perform 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 `--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 are not automatically included.

You can specify any of the following for dynamic groups:

- 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 points should be noted 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 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 apply as described in Adding a Port to an Existing Group.

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 is 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 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 not want it to be part of the group.

Perform the following steps to remove a port from a group:

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 releases, 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 software 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 Administration 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 ports with speeds less than 10Gb/s, ITW is the sum of enc_in and enc_out.
  - For ports with speeds greater or equal to 10Gb/s, ITW is the PCS block errors.
- LF
- LOSS_SYNC

> **NOTE**
> The MAPS DPORT_ITW rule for enc_out and enc_in is not applicable to the D_Port test.

The D_Port monitoring feature is only supported for 10Gb/s, 16Gb/s, and 32Gb/s SFPs/QSPFs and 8Gb/s 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 the Fabric OS versions 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,SFP_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 |
Extension 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 *Brocade Fabric OS Command Reference Manual* for more information.

When running the `portdporttest --show <port_number>` command to see details for a 32Gb/s QSPF, the output appears similar to the results for a 16Gb/s 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 the other way around). The primary task of back-end ports is to route packets passing through ASICs of a switch. 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 their functionality is different from front-end ports, which connect to devices outside of the switch.

In Fabric OS 7.4.x and later versions, MAPSmonitors port counter errors on back-end ports in chassis-based switches and 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 improves 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 is 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 the case of fixed-port switches, the slot number is 0.The History data for the back-end ports are 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 Back-end port monitoring monitors the following back-end port errors:

- CRC
- Link Reset
- ITW
- BAD_OS
- Frame too long
- Frame truncated

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, see Back End Health and Back-end Port Monitoring Default Thresholds.

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 are edited 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)     -          -          -          -          -          -
```

Front-End Encryption Port Monitoring

From Fabric OS 8.1.0 release, encryption is supported on front-end ports. MAPS monitors error counters for these ports.

Each front-end encryption port is connected to a voided back-end port, which is used to encrypt packets and send them out. Errors that are encountered with the voided ports are reported for the front-end port by the ASIC driver.

You should note the following points when working with front-end encryption ports:

- Encryption for front-end ports is supported on fibre-channel ports in blades only on the Brocade X6-4 Director and Brocade X6-8 Director.
- Encryption is supported only on E_Ports.
- MAPS requires the API chassisGetBulkFrontEndFCStats() to obtain error statistics from the ASIC driver and update global data.

MAPS monitors the following error counters:

- **Fibre-channel block errors**: Port errors including link-reset errors.
- **Fibre-channel discard errors**: Chip errors for which action is taken on every port configured for encryption.
- **Short-frame errors**: Short frames dropped by the driver and discard filter parity errors (transient errors that do not cause any problem).

NOTE
You cannot create rules that use encryption port monitoring systems. However, you can enable any of the default policies to monitor the encryption port monitoring counters.

Example of Monitoring Violations on Encryption Ports

When an E_Port is configured as an encryption port and violations are detected by any of the monitoring systems, a rule corresponding to the monitoring system is triggered, and the RASLog and Fence actions are taken. The RASLog and Fence information displays as shown in the following example:

```
2016/10/26-12:21:00, [MAPS-1003], 402, SLOT 2 FID 44, WARNING, ALLEGIANCE_C03_44, slot5 port25,
 E-Port 5/25, Condition=ALL_PORTS(ENCR_BLK/min>0), Current Value:[ENCR_BLK, 160 Timeouts],
 RuleName=defALL_E_PORTSENCR_BLK, Dashboard Category=Port Health.

2016/10/26-12:21:01, [MAPS-1010], 403, SLOT 2 FID 44, ERROR, ALLEGIANCE_C03_44, Port(s) fenced due
 to RuleName=defALL_E_PORTSENCR_BLK, Condition=ALL_PORTS(ENCR_BLK/min>0), Obj:slot5 port25, E-Port
 5/25 [ ENCR_BLK,160 Timeouts].
```

For example, the dashboard output (shown by the `mapsdb --show` command) for the rule triggered for ENCR_BLK displays as follows:

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

DB start time:                Wed Oct 26 12:17:40 2016
Active policy:                dflt_aggressive_policy
```

```
Configured Notifications:
 RASLOG,SNMP,FENCE,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,FMS,SDDQ,UNQUAR

Fenced Ports :                  None
Decommissioned Ports :          None
Fenced circuits :               None
Quarantined Ports :             None
Top Zoned PIDs <pid(it-flows)>: 0xd30000(5) 0xd30500(2)


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

Current Switch Policy Status: HEALTHY


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                 |

GE 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        |In operating range        |

Switch Resource          |In operating range        |In operating range        |

Traffic Performance      |In operating range        |In operating range        |

Extension Health         |In operating range        |In operating range        |

Fabric Performance Impact|Out of operating range    |In operating range        |



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

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

Port Health(4)           |1          |defALL_E_PORTSENCR_BLK     |10/26/16 12:21:00|E-Port 5/25 |
160 Timeouts  |
                         |3          |defALL_TARGET_PORTSSTATE_CH|10/26/16 12:23:01|F-Port 5/17 |1
            |
                         |           |G_0                        |                 |            |
            |
                         |           |                           |                 |U-Port 5/17 |1
            |
                         |           |                           |                 |F-Port 5/17 |2
            |
Fabric Performance Impact|3          |defALL_TARGET_PORTSRX_95   |10/26/16 12:20:42|F-Port 5/17 |
95.96 %       |
(9)                      |           |                           |                 |            |
            |
                         |           |                           |                 |F-Port 5/17 |
95.95 %       |
                         |           |                           |                 |F-Port 5/17 |
95.54 %       |
                         |3          |defALL_TARGET_PORTSUTIL_95 |10/26/16 12:20:36|F-Port 5/17 |
95.94 %       |
                         |           |                           |                 |F-Port 5/17 |
95.11 %       |
                         |           |                           |                 |F-Port 5/17 |
95.13 %       |
                         |3          |defALL_TARGET_PORTSTX_95   |10/26/16 12:20:36|F-Port 5/17 |
95.91 %       |
```

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 are transparent to MAPS.

Gigabit Ethernet Port Monitoring

NOTE
Gigabit Ethernet port monitoring can be performed on the following devices:

- 7840 switch
- SX6 extension blades
- FX8-24 blades

The Fabric OS software 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 of the Fibre Channel protocol.

The Ethernet MAC counters are maintained on a 1GigE, 10GigE, and 40GigE port basis.

MAPS monitors the following error 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 the 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)|,FMS             |
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)|,FMS             |
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) |,FMS             |
defALL_EXT_GE_PORTSFRM_LEN_5     |ALL_EXT_GE_PORTS(GE_INV_LEN/MIN > 5)    |                 |
dflt_aggressive_policy           |                                        |                 |
---------------------------------------------------------------------------------------------
```

Creating a CRC Rule for GE Port Monitoring

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 a 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 a 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

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. For MAPS to be able to monitor 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 *Brocade 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 the following processes:

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).

   See MAPS Rules Overview 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.

```
switch:admin> flow --create myflow_22 -feature monitor -egrport 21 -srcdev 0x010200 -dstdev
 0x011500

switch:admin> mapsconfig --import myflow_22

switch: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

switch:admin> mapsrule --create myRule_22 -group myflow_22
                          -monitor TX_FCNT -timebase hour -op g -value 22
                          -action RASLOG -policy myPolicy_22

switch: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).

   See MAPS Rules Overview 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.

```
switch:admin> flow --create myflow_22 -feature monitor -egrport 21 -srcdev 0x010200 -dstdev
 0x011500

switch:admin> mapsconfig --import myflow_22
```

Adding Monitoring Flows after Importing

To add monitoring flows after importing, you must verify that the imported flows, define a MAPS rule, 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.

   See MAPS Rules Overview 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.

```
switch: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

switch:admin> mapsrule --create myRule_22 -group myflow22 -monitor TX_FCNT -timebase hour -op g -
value 22 -action RASLOG -policy myPolicy

switch: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 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 does not take effect until you enable the policy associated with it.

Monitoring Learned Flows

Flow Vision allows you to 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. If the 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.

If a learned flow is imported into MAPS and rules are created for the flow, MAPS evaluates 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
```

Monitoring I/O Latency

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.

For support and limitations of the IO insight feature, refer to the *Brocade Flow Vision Administrator's Guide*.

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
- **I/O Bytes**: Indicates the I/O bytes transferred on the read/write basis
- **I/O Count**: Indicates the count of I/O on the read/write basis

The Gen 6 IO Insight metrics are available only with a Flow Vision flow, and they can only be monitored by MAPS when the flow is imported into MAPS.

Table 40: 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

Matrix	Size	Monitor
		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
I/O Bytes	Less than 8K	RD_IO_RATE_LT_8K
		WR_IO_RATE_LT_8K
	8K but less 64K	RD_IO_RATE_8_64K
		WR_IO_RATE_8_64K
	64K but less than 512K	RD_IO_RATE_64_512K
		WR_IO_RATE_64_512K
	Greater than or equal to 512K	RD_IO_RATE_GE_512K
		WR_IO_RATE_GE_512K
I/O Count	Less than 8K	RD_IOPS_LT_8K
		WR_IOPS_LT_8K
	8K but less 64K	RD_IOPS_8_64K
		WR_IOPS_8_64K
	64K but less than 512K	RD_IOPS_64_512K
		WR_IOPS_64_512K
	Greater than or equal to 512K	RD_IOPS_GE_512K
		WR_IOPS_GE_512K

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.

The following sections describe the steps you need to take to monitor I/O latency:

- Create and Activate a Flow to Gather I/O Statistics.
- Display the I/O Statistics.
- Import the Flow to a Logical Group.
- Create and Activate a New Policy to Monitor the New Flow.

Creating and Activating a Flow to Gather I/O Statistics

To monitor I/O latency, you must first create a flow that gathers I/O statistics and then activate it.

1. Identify the switch identifiers (SIDs) and domain identifiers (DIDs).

```
switch:admin> flow --show sys_mon_all_fports
```

   The following information is displayed:

```
Flow Monitor (Activated):
Monitor time:  | Tue May 31 19:26:52 UTC 2016 |
---------------------------------------------------------
--------------------------------------------------------------------------------------
|Ingr(*)|SID(*)|DID(*)| Rx Frames | Rx Frames |  Rx Bytes | Rx Through- | Avg Rx Frm |
|       |      |      | Count     | per Sec.  |  Count    | put (Bps)   | Sz (Bytes) |
--------------------------------------------------------------------------------------
|303    |78a7c0|531101|  443.01M  |   1.31k   |   28.05G  |   87.64k    |     68     |
|165    |78a541|732d00|  106.06G  | 303.74k   |  204.97T  |  615.52M    |    2124    |
|361    |78e1c0|735300|  122.82G  | 348.25k   |  237.37T  |  705.71M    |    2124    |
|225    |78e140|732e00|  118.81G  | 328.53k   |  229.62T  |  665.77M    |    2124    |
|109    |786d40|735400|  119.68G  | 334.68k   |  231.29T  |  678.22M    |    2124    |
|364    |78e4c0|786c40|  883.26M  |   2.46k   |   55.93G  |  163.67k    |     68     |
|108    |786c40|78e4c0|  221.69G  | 618.28k   |  428.44T  |    1.22G    |    2124    |
```

2. Create the flow.

```
switch:admin> flow --create ios_example_flow -srcdev 78a7c0 -dstdev 531101
                    -ingrport 303 -fea mon -noact
```

3. Display all flows.

```
switch:admin> flow --show
```

   The following information is displayed:

```
-------------------------------------------------------------------------------------------------
       Flow Name     |Feat|SrcDev|DstDev|IngrPt|EgrPt|BiDir|LUN|    FrameType   |SFID|DFID|MirPt|
 ------------------------------------------------------------------------------------------------
 sys_gen_all_simports|gen |*     |*     | *    | *   | no  | - |-               |-   |-   |-    |
 sys_analytics_vtap  |mir+|-     |-     | *    | -   | yes | - |VTAP_SCSI_FRAMES|-   |-   |188,0|
 sys_mon_all_fports  |mon+|*     |*     | *    | -   | no  | - |-               |-   |-   |-    |
 gg_all_LUN          |mon |78a7c0|5a1a00| 303  | -   | yes | - |-               |-   |-   |-    |
 gg_lun_0            |mon |78a7c0|5a1a00| 303  | -   | no  | 0 |-               |-   |-   |-    |
 ios_example_flow    |mon |78a7c0|531101| 303  | -   | no  | - |-               |-   |-   |-    |
```

4. Deactivate flows.

   User-defined flow monitors cannot be activated when `sys_mon_all_fports` is active.

```
switch:admin> flow --deact sys_mon_all_fports
```

   The following information is displayed:

```
Monitor feature(s) have been deactivated.
```

5. Activate newly created flow.

```
switch:admin> flow --act ios_example_flow
```

   The following information is displayed:

```
Monitor feature(s) have been activated.
```

6. Confirm that new flow is active.

   ```
   switch:admin> flow --show
   ```

   The following information is displayed. A plus-sign (+) in the feature column indicates that the flow is active.

   ```
   ------------------------------------------------------------------------------------------------
         Flow Name     |Feat|SrcDev|DstDev|IngrPt|EgrPt|BiDir|LUN|    FrameType   |SFID|DFID|MirPt|
   ------------------------------------------------------------------------------------------------
   sys_gen_all_simports|gen |*     |*     |*     |*    |no   |-  |-               |-   |-   |-    |
   sys_analytics_vtap  |mir+|-     |-     |*     |-    |yes  |-  |VTAP_SCSI_FRAMES|-   |-   |188,0|
   sys_mon_all_fports  |mon |*     |*     |*     |-    |no   |-  |-               |-   |-   |-    |
   gg_all_LUN          |mon |78a7c0|5a1a00|303   |-    |yes  |-  |-               |-   |-   |-    |
   gg_lun              |mon |78a7c0|5a1a00|303   |-    |no   |0  |-               |-   |-   |-    |
   ios_example_flow    |mon+|78a7c0|531101|303   |-    |no   |-  |-               |-   |-   |-    |
   ------------------------------------------------------------------------------------------------
     + Denotes feature is currently activated for the flow
     The flow name with prefix sys_ denotes predefined flow
   ```

Displaying I/O Statistics

After activating a flow that gathers I/O statistics, you should allow some time to pass while traffic is running, and then you can display the statistics.

Display the I/O statistics gathered by the flow that you created.

```
switch:admin> flow --show ios_example_flow


=====================================================================================================
Name      : ios_example_flow    Features: mon(Activated)        noConfig: Off
Definition: IngrPort(303),SrcDev(0x78a7c0),DstDev(0x531101)


Flow Monitor (Activated):
Monitor time:  | Tue May 31 19:52:11 UTC 2016 |
---------------------------------------------------------------------------
---------------------------------------------------------------------------
| Rx Frames    | Rx Frames    | Rx Bytes | Rx Through- | Avg Rx Frm Sz(Bytes)|
| Count        | per Sec.     | Count    | put (Bps)   |                     |
---------------------------------------------------------------------------
| 527.25k      |  1.51k       |  67.84M  |  160.01k    |         136         |
---------------------------------------------------------------------------


-------------------------------------------------------------------------------------------------
|      I/O Count        |    I/O Per Sec.(IOPS)  | I/O bytes Transferred |   I/O bytes Per Sec.  |
| Reads / Writes/ Total | Reads / Writes/ Total | Reads / Writes/ Total | Reads / Writes/ Total |
-------------------------------------------------------------------------------------------------
|415.10k/ 56.06k/471.17k|  1.39k/    61 /  1.45k|165.46G/ 35.36M/165.50G|652.75M/ 61.60k/652.80M|
-------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------------------------------
| I/O Performance:                                                                              |
-------------------------------------------------------------------------------------------------
|       Metric        |    IO Size  |  I/O Count  |     Max(IOPS)     |       Avg(IOPS)         |
|                     |             |     All     |      All          |    6 sec   /     All    |
```

```
-------------------------------------------------------------------------------------------------
| RD IO Count        | <8K        |    75.71k   |       96          |       77     /       64   |
|                    | 8K - <64K  |         0   |        0          |        0     /        0   |
|                    | 64K - <512K|         0   |        0          |        0     /        0   |
|                    | >=512K     |   330.74k   |     1.30k         |     1.29k    /      283   |
|                    | ALL        |   406.45k   |     1.37k         |     1.36k    /      347   |
-------------------------------------------------------------------------------------------------
| WR IO Count        | <8K        |    55.62k   |       66          |       63     /       47   |
|                    | 8K - <64K  |         0   |        0          |        0     /        0   |
|                    | 64K - <512K|         0   |        0          |        0     /        0   |
|                    | >=512K     |         0   |        0          |        0     /        0   |
|                    | ALL        |    55.62k   |       66          |       63     /       47   |
-------------------------------------------------------------------------------------------------



-----------------------------------------------------------------------------------------------
| I/O Latency Metrics:                                                                        |
-----------------------------------------------------------------------------------------------
|           Metric           |   IO Size  |           Max            |            AVG         |
|                            |            |    6 sec   /     All    |    6 sec   /     All    |
-----------------------------------------------------------------------------------------------
| RD CMD -> Status Time      | <8K        |     7.81m  /     8.05m  |     3.87m  /     3.93m  |
|                            | 8K - <64K  |            /            |            /            |
|                            | 64K - <512K|            /            |            /            |
|                            | >=512K     |    13.80m  /    16.76m  |     7.08m  /     6.69m  |
|                            | ALL        |    13.80m  /    16.76m  |     6.90m  /     6.18m  |
-----------------------------------------------------------------------------------------------
| WR CMD -> Status Time      | <8K        |    16.15m  /    16.25m  |    10.40m  /    10.42m  |
|                            | 8K - <64K  |            /            |            /            |
|                            | 64K - <512K|            /            |            /            |
|                            | >=512K     |            /            |            /            |
|                            | ALL        |    16.15m  /    16.25m  |    10.40m  /    10.42m  |
-----------------------------------------------------------------------------------------------
| RD CMD -> 1st Data Time    | <8K        |     7.81m  /     8.02m  |     3.86m  /     3.92m  |
|                            | 8K - <64K  |            /            |            /            |
|                            | 64K - <512K|            /            |            /            |
|                            | >=512K     |     8.29m  /     8.30m  |     3.70m  /     3.71m  |
|                            | ALL        |     8.29m  /     8.30m  |     3.71m  /     3.75m  |
-----------------------------------------------------------------------------------------------
| WR CMD -> 1st XFER_RDY Time| <8K        |    15.74m  /    15.88m  |    10.21m  /    10.38m  |
|                            | 8K - <64K  |            /            |            /            |
|                            | 64K - <512K|            /            |            /            |
|                            | >=512K     |            /            |            /            |
|                            | ALL        |    15.74m  /    15.88m  |    10.21m  /    10.38m  |
-----------------------------------------------------------------------------------------------
| RD Pending IOs             | <8K        |        1   /        1   |        1   /        1   |
|                            | 8K - <64K  |        0   /        0   |        0   /        0   |
|                            | 64K - <512K|        0   /        0   |        0   /        0   |
|                            | >=512K     |       14   /       14   |        7   /       14   |
-----------------------------------------------------------------------------------------------
| WR Pending IOs             | <8K        |        1   /        1   |        0   /        1   |
|                            | 8K - <64K  |        0   /        0   |        0   /        0   |
|                            | 64K - <512K|        0   /        0   |        0   /        0   |
|                            | >=512K     |        0   /        0   |        0   /        0   |
----------------------------------------------------------------------------------------------
===============================================================================================
```

Importing the Flow to a Logical Group

After creating a working flow that gathers I/O statistics, you can add it to a logical group.

1. Import the flow you created into a logical group.

   ```
   switch123:admin> mapsconfig --import ios_example_flow
   ```

   The following information is displayed:

   ```
   2016/05/31-19:54:24, [MAPS-1124], 7727, SLOT 1 FID 128, INFO,
   switch123,  Flow ios_example_flow imported.
   ```

2. Confirm the flow was added to the logical group.

   ```
   switch123:admin> logicalgroup --show
   ```

   The following information is displayed. Notice the last line of data, indicating that the flow was added.

   ```
   ------------------------------------------------------------------------------------------------
   Group Name              |Predefined |Type         |Member Count |Members
   ------------------------------------------------------------------------------------------------
   ALL_PORTS               |Yes        |Port         |209          |5/6-33,5/35-47,7/0-63,8/0-63...
   ALL_E_PORTS             |Yes        |Port         |4            |5/8,5/11,5/16,5/39

   [lines removed for display purposes only in this document]

   ALL_SLOTS               |Yes        |Blade        |12           |1,2,3,4,5,6,7,8,9,10,11,12
   ALL_SW_BLADES           |Yes        |Blade        |2            |5,11
   ALL_CORE_BLADES         |Yes        |Blade        |2            |7,8
   ALL_ASICS               |Yes        |Asic         |12           |5/0-1,7/0-3,8/0-3,11/0-1
   ALL_CERTS               |Yes        |Certificate  |0            |
   ALL_LOCAL_PIDS          |Yes        |Pid          |61           |All Pids monitored
   SWITCH                  |Yes        |             |1            |0
   CHASSIS                 |Yes        |             |1            |0
   ALL_SFP                 |Yes        |Sfp          |37           |5/6-9,5/11,5/15-16,5/18-19...
   ALL_10GSWL_SFP          |Yes        |Sfp          |0            |
   ALL_10GLWL_SFP          |Yes        |Sfp          |0            |

   [lines removed for display purposes only in this document]

   ALL_CIRCUIT_IP_MED_QOS  |Yes        |Circuit Qos  |0            |
   ALL_CIRCUIT_IP_HIGH_QOS |Yes        |Circuit Qos  |0            |
   ALL_TUNNELS             |Yes        |Tunnel       |0            |
   ALL_TUNNEL_F_QOS        |Yes        |Tunnel  Qos  |0            |
   ALL_TUNNEL_LOW_QOS      |Yes        |Tunnel  Qos  |0            |
   ALL_TUNNEL_MED_QOS      |Yes        |Tunnel  Qos  |0            |
   ALL_TUNNEL_HIGH_QOS     |Yes        |Tunnel  Qos  |0            |
   ALL_TUNNEL_IP_LOW_QOS   |Yes        |Tunnel  Qos  |0            |
   ALL_TUNNEL_IP_MED_QOS   |Yes        |Tunnel  Qos  |0            |
   ```

```
ALL_TUNNEL_IP_HIGH_QOS  |Yes        |Tunnel  Qos  |0            |
gg_all_LUN              |No         |Flow         |1            |Monitored Flow
ios_example_flow        |No         |Flow         |1            |Monitored Flow
```

Creating and Activating a New Policy to Monitor the New Flow

To monitor I/O statistics, you can either create a new policy or clone an existing one. Then, add a new rule to the policy to monitor the flow that you created.

For the examples in this procedure, a clone of the default aggressive policy is cloned.

1. Ensure the policy you want to clone is active. In this example, the dflt_aggressive_policy is being cloned.

   ```
   switch123:admin> mapspolicy --show -summary
   ```

   The following information is displayed:

   ```
           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_aggressive_policy'.
   ```

2. Clone the policy.

   ```
   switch123:admin> mapspolicy --clone dflt_aggressive_policy
                             -name my_aggressive_plus_ios_rules
   ```

   The following information is displayed:

   ```
   2016/05/31-20:14:21, [MAPS-1112], 7734, SLOT 1 FID 128, INFO, switch123,
   Policy dflt_aggressive_policy cloned to my_aggressive_plus_ios_rules.
   ```

3. Create the rule and add it to the new policy.

   ```
   switch123:admin> mapsrule --create ios_mon_1st_data_less_8k -group ios_example_flow -monitor
   RD_1stDATA_TIME_LT_8K -value 10 -timebase min -op ge -action raslog,email,snmp -policy
   my_aggressive_plus_ios_rules
   ```

   The following information is displayed:

   ```
   2016/05/31-20:18:52, [MAPS-1100], 7735, SLOT 1 FID 128, INFO, switch123, Rule
   ios_mon_1st_data_less_8k is created.
   2016/05/31-20:18:52, [MAPS-1114], 7736, SLOT 1 FID 128, INFO, switch123, Rule
   ios_mon_1st_data_less_8k added to Policy my_aggressive_plus_ios_rules
   ```

4. Activate the new policy.

   ```
   switch123:admin> mapspolicy --enable my_aggressive_plus_ios_rules
   ```

   When a condition triggers the rule, a RASLog message is sent to the console if the action *raslog* is enabled.

   ```
   switch123:admin> 2016/05/31-20:23:03, [MAPS-1003], 7741, SLOT 1 FID 128, WARNING,
   switch123, Flow (ios_example_flow), Condition=ios_example_flow
   (RD_1stDATA_TIME_LT_8K/min>=10), Current Value:[ RD_1stDATA_TIME_LT_8K,3849 Microseconds],
   RuleName=ios_mon_1st_data_less_8k, Dashboard Category=Traffic Performance.
   ```

5. You can also display the rule in the MAPS dashboard output.

   ```
   switch123:admin> mapsdb --show
   ```

   The following information is displayed:

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

DB start time:                  Tue May 31 20:20:52 2018
Active policy:                  my_aggressive_plus_ios_rules
Configured Notifications:       RASLOG,SNMP,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,SDDQ
Fenced Ports :                  None
Decommissioned Ports :          None
Fenced circuits :               None
Quarantined Ports :             11/4,11/5,11/6,11/7,11/12,11/13,5/47,11/41,11/44,11/47
Top Zoned PIDs <pid(it-flows)>: 0x78a6c0(14) 0x78a3c0(14) 0x78a340(8) 0x78ae40(4) 0x78af40(4)
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              |In operating range        |In operating range        |
BE Port Health           |No Errors                 |No Errors                 |
GE Port Health           |No Errors                 |No Errors                 |
Fru Health               |In operating range        |In operating range        |
Security Violations      |In operating range        |In operating range        |
Fabric State Changes     |In operating range        |In operating range        |
Switch Resource          |In operating range        |In operating range        |
Traffic Performance      |Out of operating range    |In operating range        |
Extension 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 Value  |
(Rule Count)    |Count |                   |                 |                 |(units)          |
--------------------------------------------------------------------------------------------------
Traffic         |4     |ios_mon_1st_data_  |05/31/16 20:26:03|Flow (ios_example|3766 Microseconds|
Performance (4) |      |less_8k            |                 |_flow)           |                 |
                |      |                   |                 |Flow (ios_example|3826 Microseconds|
                |      |                   |                 |_flow)           |                 |
                |      |                   |                 |Flow (ios_example|3826 Microseconds|
                |      |                   |                 |_flow)           |                 |
                |      |                   |                 |Flow (ios_example|3849 Microseconds|
                |      |                   |                 |_flow)           |                 |
Fabric Perform- |1     |defALL_PORTS_IO_   |05/31/16 20:23:03|F-Port 5/47      |IO_PERF_IMPACT   |
ance Impact (2) |       |PERF_IMPACT        |                  |                 |                 |
                |       |                   |                  |                 |                 |
                |1      |defALL_PORTS_IO_   |05/31/16 20:22:03|F-Port 5/47       |IO_LATENCY_CLEAR |
                |       |LATENCY_CLEAR      |                  |                 |                 |
```

Monitoring Flows Created Using VM Insight

Starting with Fabric OS 8.1.x release, the Monitoring and Alerting Policy Suite (MAPS) can monitor applications supported by VM Insight by monitoring flows created using VM Insight. MAPS 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 evaluates various conditions and generates RASLogs, SNMP alerts, or e-mails if the application violates MAPS rules.

Perform the following steps for MAPS to monitor flows created using VM Insight:

1. Create the flow. For example:
   ```
   switch:admin> flow --create vmMonitor --srceid 00112233-4455-6677-8899-aabbccddeeff --
   dstdev 0x123456 –feat mon –ingr 10
   ```
2. Import the flow into MAPS. For example:
   ```
   switch:admin> mapsconfig -import vmMonitor
   ```
3. Verify that the flow was imported as a group. The group name *vmMonitor* should be listed. For example:
   ```
   switch:admin> logicalgroup -show
   ```
4. Define a MAPS rule to monitor the created flow or group. For example:
   ```
   switch:admin> mapsrule -create vmMonitorRule -group vmMonitor -monitor
   WR_PENDING_IO_8_64K -timebase minute -value 10 -op GE -action raslog,snmp,email
   ```

   See MAPS Rules Overview for information on creating and using rules.

Statistics Used to Monitor Applications

MAPS monitors the following statistics when monitoring applications:

Table 41: Statistics Used to Monitor Applications

TX_FCNT RX_FCNT TX_THPUT RX_THPUT IO_RD IO_WR IO_RD_BYTES IO_WR_BYTES RD_STATUS_TIME_LT_8K RD_STATUS_TIME_8_64K RD_STATUS_TIME_64_512K RD_STATUS_TIME_GE_512K WR_STATUS_TIME_LT_8K WR_STATUS_TIME_8_64K	WR_STATUS_TIME_64_512K WR_STATUS_TIME_GE_512K WR_1stXFER_RDY_LT_8K WR_1stXFER_RDY_8_64K WR_1stXFER_RDY_64_512K WR_1stXFER_RDY_GE_512K RD_PENDING_IO_LT_8K RD_PENDING_IO_8_64K RD_PENDING_IO_64_512K RD_PENDING_IO_GE_512K WR_PENDING_IO_LT_8K WR_PENDING_IO_8_64K WR_PENDING_IO_64_512K WR_PENDING_IO_GE_512K

NOTE
For more information about each statistic, see Traffic Performance and Monitoring I/O Latency.

Fabric Performance Impact Monitoring Using MAPS

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 other elements) and consequently multiple places that can be the cause of fabric performance impacts (bottlenecks). Because the behavior of each application is unique, the impact of a bottleneck on one individual application may 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 *Brocade Fabric OS Administration Guide*.

NOTE
In Fabric OS 8.0.1 and later releases, the 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 the data retrieved from the port on the switch (just one element in the fabric), which is used to determine the potential impact on 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 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 is 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

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         |
Extension 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
```

After the latency is cleared on the F_Port 8/8, the MAPS dashboard report changes to the following:

```
switch:admin> mapsdb --show

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         |
Extension 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      |
```

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.

Transmit Queue Latency Counter Monitoring

MAPS monitors port transmit 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, transmit 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 and Slow-Drain Device Quarantining 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 the transmit 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
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         |
Extension 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_vcare the problem ports.

```
switch:admin> portstatsshow 1
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
```

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 release, 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.

```
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, the creation of a rule that specifies 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 release, the bottleneck detection functionality is 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. Note that the column headings in the example are edited 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 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. Note that the headings in the example are edited 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 continues to log events whether RASLogs are set to on or off in the bottleneck configuration.

Refer to the *Brocade Fabric OS Administration Guide* for specific command details and bottleneck monitoring parameters.

Port Toggling Support

MAPS supports port toggling, which allows Fabric OS software to recover a port that is 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. 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 is switched over to a redundant path. MAPS also assumes that while the port is being toggled, the operational state of the port is not 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 is 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,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL
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 is 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 Channel/Virtual Circuit (VC). However, the credits used to send traffic across the link are common to all the flows using the same link. Therefore, a slow-draining device can 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 the figure, the inability of the slow-draining device (1) to clear frames quickly enough is causing back-pressure not just to the edge switch or AG device (2), but also to the core switch (3) and the target (4) that is trying to send data to the slow-draining device.

Figure 1: Slow Drain Dataflow

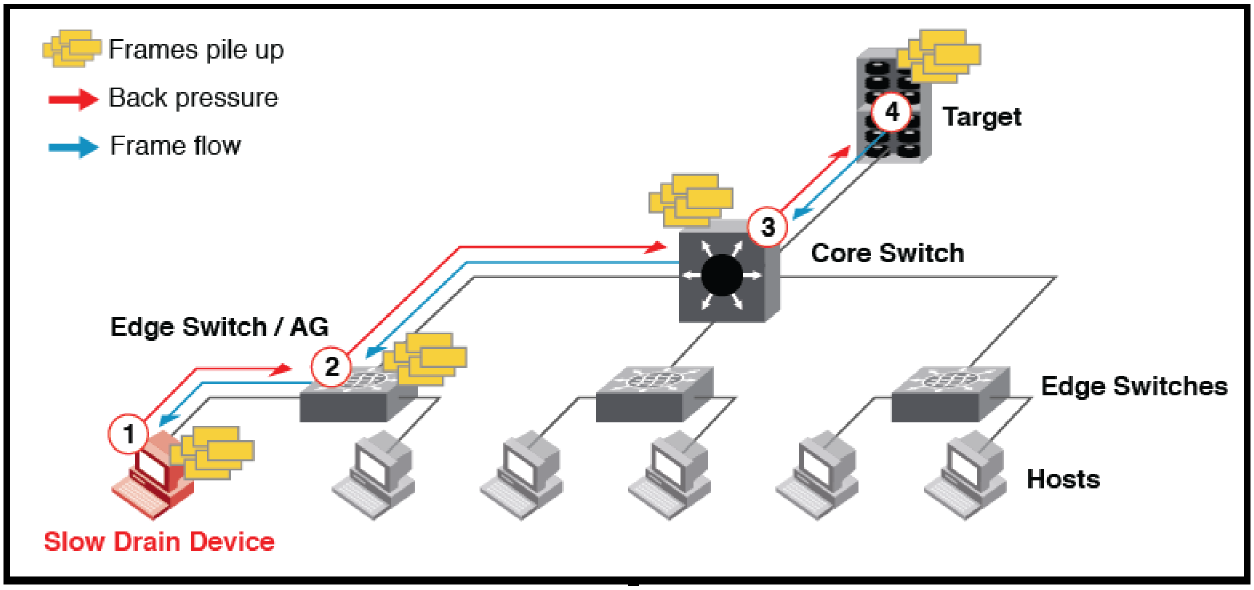

To isolate the traffic destined to slow drain the device or port, the Slow-Drain Device Quarantine (SDDQ) feature of Fabric Performance Impact (FPI) quarantines the port. SDDQ moves all the traffic destined for the slow-drain devices to a low priority VC and adds the port to the predefined *ALL_QUARANTINED_PORTS MAPS* logical group. The RASLog messages are displayed to notify that the traffic is successfully moved to a lower priority VC.

The following example displays the RASLog messages for SDDQ:

```
[NS-1017], 6, SLOT 6 | FID 21, INFO, DCX_120_21, Domain 34, Port index 171, All devices zoned with the slow
 drain device 0x220d00 have been quarantined.
[MAPS-1022], 7, SLOT 6 | FID 21, WARNING, DCX_120_21, Port 3/27 (Port index 171) has been marked as Slow Drain
 Device.
[AG-1086], 513, FID 128, INFO, G610, Slow drain device 0x011205 connected to F-port (8) have been quarantined.
[AG-1087], 528, FID 128, INFO, G610, Slow drain device 0x011205 connected to F-port (8) have been
 unquarantined.
```

For more information on RASLog messages, refer to the *Brocade FOS Message Reference Guide*.

Once MAPS quarantines a port, it continues to monitor the slow-drain port and if the condition gets cleared, then MAPS automatically performs the un-quarantine action (if configured). The user can also manually un-quarantine the port. See Manually Clearing Quarantined Ports or Automatically Clearing Quarantined Ports for more details about this procedure.

Prerequisites

The following points must be taken care before using SDDQ feature:

- The Quality of Service (QoS) must be enabled to use the SDDQ feature.
- The fabric vision license must be installed on the switch to use the SDDQ feature. Intermediate switches do not need the fabric vision license for SDDQ, but they must be enabled with QoS on all Inter-Switch Links (ISL).

  NOTE
  The Quality of Service (QoS) feature license, also known as the *Adaptive Networking (AN)* feature, does not need to be installed to use the QoS feature. The QoS feature can be configured without the licenses in Fabric OS 7.2.0 and later releases.

Important Points on Slow-Drain Device Quarantining

The following notes of Slow-Drain Device Quarantining (SDDQ) are:

- SDDQ is disabled by default on installation.
- SDDQ is supported:
  - On all Fabric OS platforms that are running Fabric OS 7.4.0 or later releases, including FICON environments.
  - For N_Port ID Virtualization (NPIV) devices connected to F_Ports.
  - On switches running in Access Gateway mode from Fabric OS 8.2.1 or a later release.
- SDDQ is not supported:
  - On switches running Fabric OS versions earlier than 7.4.0.
  - On switches running in Access Gateway mode earlier than Fabric OS 8.2.1 release.
  - In Fibre Channel Routing (FCR) backbone fabrics. Although, SDDQ can be supported on devices within an FCR edge fabric, the edge fabric will not apply SDDQ to the flows for any FCR-imported devices.
  - On F_Port trunks.
- All the actions SDDQ and UNQUAR are associated with monitoring system *DEV_LATENCY_IMPACT* but not with the state. In this case, SDDQ and UNQUAR actions are configured with all *frame_loss, perf_impact*, and latency clear.
- Although MAPS allows you to enable and disable SDDQ on switches, it is recommended that you enable SDDQ for the entire fabric.
- No default rules are defined for the *ALL_QUARANTINED_PORTS* group, but you can create custom rules to monitor the ports in this group.
- When a port is marked as slow-drain, only the flows destined to it are shifted to the low-priority VC. Flows in the reverse direction are not affected.
- Quarantine action takes the precedence over QoS zone configuration. The traffic is moved to the low priority VC irrespective of the zone QoS priority (high, medium or low).
- If device latency is due to Class 3 timeouts (C3TXTO) and the active C3TXTO rule has port fencing (port disabling) as an action, then it 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 might not be possible.
- The maximum number of ports that can be isolated per unit (chassis or fixed-port switch) is 32. The default value is 10. This is controlled by the chassis-wide *Chassis SDDQ limit* user-configurable key, which is set using the `configurechassis` command. Refer to the *Brocade Fabric OS Command Reference Manual* for more information.
- To restore the default value as 10 for the chassis SDDQ limit, you must perform `switchdisable` . Then you must perform `chassisdefault -all` to restore it to 10.
- SDDQ action is blocked in the following scenarios to prevent quarantining a large number of source ports:

– If the total number of ports zoned to the slow-drain port is set to 33 or more (This does not include slow-drain port).
– If the defzone is labeled as *all access*.

NOTE
The Port Toggling action should not be used in conjunction with SDDQ, because this could result in unpredictable behavior.

Un-quarantine Timeout Behavior Across HA Failover

After an HA failover completes, MAPS continues to monitor for quarantined ports for the remainder of the time specified in the rule. Also, MAPS continues to perform an automatic un-quarantine if the `UNQUAR` action is specified in the FPI rule and the UNQUAR action is globally enabled in `mapsconfig` , when the SDDQ state is clear. For example, if a rule specified an un-quarantine action for 2.0 hours and an HA failover event occurs after 30 minutes, MAPS will monitor the state for an additional 1.5 hours after the HA failover completes.

All the SDDQ information is persistent across HA failover, which includes quarantine ports configuration, un-quarantine timeout and monitoring window, and rules. For example, if a rule specified with an un-quarantine timeout of 2.0 hours and an HA failover event occurs in 30 minutes after the SDDQ is cleared on a port, MAPS monitors the state for an additional 1.5 hours after the HA failover completes.

Enabling Slow-Drain Device Quarantining

Slow-Drain Device Quarantining (SDDQ) is enabled by having a valid Fabric Vision license, and it is activated by including the SDDQ action in the configured MAPS action list.

To enable SDDQ, perform 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.
3. Enable a policy with rules that have the SDDQ action.

   The following example enables SDDQ as an available action:

   ```
   switch:admin> mapsconfig --actions SNMP,RASLOG,EMAIL,SDDQ
   ```

Scalability Limits of SDDQ

SDDQ has the following scalability limits:

1. Maximum quarantine ports per chassis are 32
2. Maximum ports quarantine per slow-drain port is 32

As per above requirements, the MAPS must not quarantine more than 32 ports per chassis or it cannot move the traffic to low priority VC if the port is zoned with 33.

In most cases, you should not enable the Slow-Drain Device Quarantining (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 can help with slow-draining device issues.

Typically, in a FICON setup or other setup where the system has a single zone, it is recommended not to use the SDDQ feature.

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 is called out in this example. If there is no slow-draining device, the line does 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
    Device link speed: 8G
    Connected through AG: Yes
    Real device behind AG: Yes
    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
    Device link speed: 8G
    Connected through AG: Yes
    Real device behind AG: Yes
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
    Device link speed: 8G
    Connected through AG: Yes
    Real device behind AG: Yes
    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
    Device link speed: 8G
    Connected through AG: Yes
    Real device behind AG: Yes
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 is called out in this example. If there was no slow-draining device, the line does 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 is called out in this example. If there was no slow-draining device, the line does 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 are quarantined in the *ALL_QUARANTINED_GROUP*.

Perform the following steps to display a list of ports which are quarantined by MAPS:

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: 3
---------------------------------------------------------------------

Online Quarantined Devices across the fabric
---------------------------------------------------------------------
 Port Index |   PID   |            PWWN          |  Locality
---------------------------------------------------------------------
     3      | 051100  | 30:10:00:05:33:ac:c6:13  |  LOCAL
     3      | 051101  | 30:10:01:05:33:ac:c6:13  |  LOCAL
---------------------------------------------------------------------
```

Clearing Quarantined Ports

MAPS allows you to remove ports which are quarantined from the quarantine group (*ALL_QUARANTINED_GROUP*). You can also configure MAPS to automatically remove the quarantined ports from the quarantine group if slow-draining behavior is not detected for a period of time.

Manually Clearing Quarantined Ports

To manually remove ports from the quarantine group that is quarantined by MAPS, perform the following steps:

1. Enter `sddquarantine --show` to view a list of the quarantined ports.
2. Enter `sddquarantine --clear` followed by either:
   - The slot/port ID that needs to be removed from the quarantine group.
   - The keyword `all` to clear all quarantined ports.

   A port is not allowed to be cleared from quarantine while the latency or frame loss condition that caused it to be quarantined persists, because the port will not be flagged again until the condition is cleared and then retriggered. However, you can override this restriction 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 if the number of devices zoned to slow-drain port exceeds to zoned limits of 32.

   The first part of the following example uses the `sddquarantine --show` command to display all quarantined local ports and the online quarantined device information across the fabric. The second part shows the use of the `sddquarantine --clear` command to remove the ports from quarantine. Notice that port 3 was not able to be cleared using the `all` option, because it was still in either the *IO_FRAME_LOSS* or the *IO_PERF_IMPACT* condition. This port was 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          |  Locality
---------------------------------------------------------------------
   17      | 051100  | 30:10:00:05:33:ac:c6:13   |  LOCAL
   17      | 051101  | 30:10:01:05:33:ac:c6:13   |  LOCAL
---------------------------------------------------------------------

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
```

Automatically Clearing Quarantined Ports

You can configure MAPS to automatically remove (un-quarantine) a port from the quarantine group after the port changes to the *IO_LATENCY_CLEAR* state and remains in this state for a specified amount of un-quarantine timeout. To achieve this, you create a rule using the *UNQUAR* action and set the un-quarantine timeout value using the `-uqrt` and `–uqrt_unit` options.

Perform the following steps:

1. Create or configure a FPI rule to monitor the *IO_FRAME_LOSS/PERF_IMPACT* state. Specify the *UNQUAR* action and set the duration using `-uqrt` and specify the time unit (*min*, *hour*, or *day*) with `-uqrt_unit` .
2. Enable the *UNQUAR* action in `mapsconfig` .

Example of Clearing Quarantined Ports

For example, to set the un-quarantine action for one hour after the port has changed to the *IO_LATENCY_CLEAR* state, you might create a rule as follows:

```
switch:admin> mapsrule --create test_PORTS_IO_FRAME_LATENCY -group ALL_PORTS -monitor
 DEV_LATENCY_IMPACT
                       -timebase none -op eq -value IO_PERF_IMPACT
                       -action RASLOG,SDDQ,UNQUAR -uqrt 1 -uqrt_unit hour
```

Example of Removing Automatic Un-quarantine Action

To remove the automatic un-quarantine action from a rule, configure the rule by removing the *`UNQUAR`* action from the action list and using the `-uqrt_clear` option as follows:

```
switch:admin> mapsrule --config test_PORTS_IO_FRAME_LATENCY
                        -action RASLOG,SDDQ -uqrt_clear
```

Considerations When Using the UNQUAR Action

When using the *UNQUAR* action to automatically un-quarantine ports, you must note the following information:

- You can use the *UNQUAR* action only with FPI rules that monitor the *IO_PERF_IMPACT* or *IO_FRAME_LOSS* state.
- To use the *UNQUAR* action, use the `mapsconfig --actions` command and specify *UNQUAR* as one of the actions.
- The *UNQUAR* action can be specified in an FPI rule or a global configuration only with the SDDQ action. However, the SDDQ action can be specified without the *UNQUAR* action if automatic unquarantining is not required.
- The *IO_PERF_IMPACT* and *IO_FRAME_LOSS FPI* rules in the moderate and conservative default policies are configured with the *UNQUAR* action with an unquarantine timeout of 10 minutes.
- Unquarantine actions are aborted when the disruptive operations result during a configuration download or when new policies are enabled. You must manually clear ports from the quarantine group. Refer to Manually Clearing Quarantined Ports.
- Do not have more than one rule to un-quarantine but if multiple rules are configured with the *UNQUAR* action and un-quarantine timeouts, then the timeout value of the first rule that triggered determines when the port is un-quarantined.

NOTE
MAPS FPI reports *IO Latency_Clear UNQUAR* even when the action of SDDQ and *UNQUAR* are not configured. The dashboard update happens irrespective of the configured actions.

Cases of Using the UNQUAR Action

The following cases demonstrate how a port is unquarantined when using the *UNQUAR* action with a timeout value specified. In both cases, use an FPI rule that includes the *UNQUAR* action with a timeout of 1 hour.

Case 1:

- The port is quarantined because slow-drain behavior was detected at 7:00 am and continued until 7:30 am.
- The state of the port is clear at 7:30 am, so MAPS starts the timer to monitor this state for 1 hour.
- The state of the port is clear from 7:30 am until 8:30 am without any slow-drain activity.
- The port is unquarantined automatically at 8:30 am.

Case 2:

- The port is quarantined, because slow-drain behavior was detected at 7:00 am and continued until 7:30 am.
- The state of the port is clear at 7:30 am, so MAPS starts the timer to monitor this state for 1 hour.
- Slow-drain behavior is detected again at 8:00 am and MAPS stops the timer for this port.
- The state of the port clears again at 8:30 am, so MAPS start the timer again to monitor this state for 1 hour.
- At 9:30 am, the state is still clear (no additional slow-drain behavior is detected), and the port is automatically unquarantined.

Delaying the Quarantining of Ports

Using the capabilities of the rule-on-rule (RoR) rules, you can delay the quarantining of a port until after a specified number of rule violations.

For example, you can delay the quarantining of a port until after a rule is violated five times:

1. Create a policy.

```
switch:admin> mapspolicy --create test_policy1
```

2. Create a base FPI rule, but do **not** set an SDDQ action.

```
switch:admin> mapsrule --create test_PORTS_IO_FRAME_LOSS -group ALL_PORTS
                  -monitor DEV_LATENCY_IMPACT -timebase none -op eq
                  -value IO_FRAME_LOSS -action raslog,snmp,email -policy test_policy1
```

3. Create a rule-on-rule (RoR) rule, including an SDDQ action.

NOTE
See Creating RoR Rules to Monitor Other Rules for more details about RoR rules.

```
switch:admin> mapsrule --createRoR test_ror_PORTS_IO_FRAME_LOSS -group test_PORTS_IO_FRAME_LOSS
                  -monitor test_PORTS_IO_FRAME_LOSS -timebase hour -op g -value 5
                  -action raslog,sddq,unquar -uqrt 1 -uqrt_unit hour
```

This RoR rule monitors the base FPI rule to determine how long the port remains in the congestion state. The above rule gets triggered if the port remains in the congestion state for five minutes. It only triggers the SDDQ action if the state changes to *IO_FRAME_LOSS* more than five times within an hour. Therefore, the port is not quarantined on the first state change, only on the sixth state change within an hour.

Congestion Dashboard

Congestion Dashboard displays the list of potentially congested ports determined by monitoring the following aspects:

NOTE
Congestion Dashboard is supported only for Gen 6 platforms.

- **Transient Queue Latency (TQL)** — TQL is the maximum time a frame spent in the transmit queue.

  - If the TQL is 80 or more milliseconds, it is considered as frame loss due to congestion (*IO_FRAME_LOSS* state).
  - If the TQL is 10 or more milliseconds, it is considered as traffic performance impact due to congestion (*IO_PERF_IMPACT* state).
- **Credit zero statistics** — Credit zero statistics indicates the time taken between a frame ready to be transmitted and the port made buffer credits available for the transmission. The following are the allowed latency per time intervals in the order of checking done. If a port exceeds any of these latency values, then the port is considered to be impacting traffic performance (*IO_PERF_IMPACT* state).
  a. 700 milliseconds in 1 second.
  b. 2500 milliseconds in 5 seconds.
  c. 3000 milliseconds in 10 seconds.
- **Class 3 discard errors** — If a port drops frames, it is considered to be frame loss due to congestion (IO_FRAME_LOSS state).

Congested Port States

In the Congestion Dashboard, the ports are categorized into the following states based on the specified criteria:

- **Frame Loss** — This is the highest severity congestion state. In this state, MAPS generates a frame loss alert.
- **I/O Perf Impact** — This is the second highest severity congestion state. In this state, MAPS generates a performance impact alert.
- **Medium** — A port is marked in medium congestion state if any or both of the following conditions are met:
  - If the TQL is 5 or more milliseconds but less than 10 milliseconds.
  - If the credit zero statistics indicates a latency of 100 or more milliseconds but less than 700 milliseconds in 1 second.
- **Low** — A port is marked in low congestion state if any or both of the following conditions are met:
  - If the TQL is 3 or more milliseconds but less than 5 milliseconds.
  - If the credit zero statistics indicates a latency of 50 or more milliseconds but less than 100 milliseconds in 1 second.
- **Info** — A port is marked in informative congestion state if any or both of the following conditions are met:
  - If the TQL is 1 or more milliseconds but less than 3 milliseconds.
  - If the credit zero statistics indicates a latency of 10 or more milliseconds but less than 50 milliseconds in 1 second.

Congestion State Calculation

The *Medium*, *Low*, and *Info congestion* states of a port are determined using the area weight principle by considering 60 samples of statistics each one second apart. For example, consider the following scenarios:

- In a minute, if a port remains 8 times in medium, 6 times in low, 3 times in info congestion state, then the area weight is equal to 8*3+ 6*2+3*1 = 39
- In a minute, if a port remains 12 times in medium, 7 times in low, 2 times in info congestion state, then the area weight is equal to 12*3+7*2+2*1 = 52

Table 42: Congestion State Calculation

Duration	State	Threshold Range Based on Area Weight
Per Second	Frame Loss	5
	I/O Perf Impact	4
	Medium	3
	Low	2
	Info	1

Duration	State	Threshold Range Based on Area Weight
Per Minute	Frame Loss	>240 and <=300
	I/O Perf Impact	>180 and <=240
	Medium	>120 and <=180
	Low	>60 and <=120
	Info	>0 and <=60
Per Hour	Frame Loss	>14400 and <=18000
	I/O Perf Impact	>10800 and <=14400
	Medium	>7200 and <=10800
	Low	>3600 and <=7200
	Info	>0 and <=3600

Displaying the Congestion Dashboard

To display the Congestion Dashboard of the top 10 ports for the last one hour, run the `mapsdb --show congestion` command. You can also list more than 10 ports in the Dashboard using the `-top` option. To display the Congestion Dashboard at the specified hour, use the *-hr* option.

```
Switch#admin> mapsdb --show congestion
-----------------------------------------
DB start time:  Thu Sep  7 20:45:05 2017
-----------------------------------------
                                 |Frequency details for time window (22:48 - 23:48)|
Port        |Current Min State   |Frame Loss  |Perf Impact |Medium |Low    |Info   |
-----------------------------------------------------------------------------------
E-Port 5/44 |Low St              |0           |0           |19     |41     |0      |
E-Port 11/38|Low St              |0           |0           |16     |44     |0      |
E-Port 5/40 |Low St              |0           |0           |17     |43     |0      |

Switch#admin> mapsdb --show congestion -state -hr 20
-----------------------------------------
DB start time:  Thu Sep  7 20:45:05 2017
-----------------------------------------
                                 |       Frequency details for hour 20:00:00       |
Port        |Current Min State   |Frame Loss  |Perf Impact |Medium |Low    |Info   |
-----------------------------------------------------------------------------------
E-Port 5/9  |No Congestion       |0           |0           |0      |0      |5      |
E-Port 5/40 |Info St             |0           |0           |0      |0      |7      |
E-Port 5/44 |No Congestion       |0           |0           |0      |0      |2      |
E-Port 10/16|No Congestion       |0           |0           |0      |0      |2      |
E-Port 11/38|No Congestion       |0           |0           |0      |0      |1      |
F-Port 12/23|No Congestion       |0           |0           |0      |0      |2      |
F-Port 12/21|No Congestion       |0           |0           |0      |0      |2      |
```

To display the Congestion Dashboard with frequency details, use the `-freq` option. Frequency table displays the frequency of the congestion for the past 10 hours. It is sorted based on congestion frequency count for a port in a given hour. MAPS samples port state every second. This indicates number of seconds in the particular hour bucket the port was in congestion state. First column represents current hour. So if the current switch time is 20:30:00, then the first column in the following output will have 30 minutes frequency value (20:00:00 to 20:30:00). All the other columns represent complete one-hour data.

```
Switch:admin> mapsdb --show congestion -freq
-----------------------------------------
DB start time:  Thu Sep  7 20:45:05 2017
-----------------------------------------
23:00:00         |22:00:00         |21:00:00         |20:00:00         |19:00:00         |18:00:00         |
17:00:00         |16:00:00         |15:00:00         |14:00:00         |
-----------------------------------------------------------------------------------------------------------------------
E-Port 11/38,(293|E-Port 11/38,(357|E-Port 5/40,(47) |E-Port 5/9,(18)  |F-Port 12/21,(15)|F-Port 12/23,(3) |
               |                 |                 |                 |
0)               |9)               |                 |                 |                 |                 |
               |                 |                 |                 |
E-Port 5/40,(2928|E-Port 5/44,(3576|E-Port 11/38,(27)|E-Port 5/40,(9)  |F-Port 12/23,(12)|F-Port 12/21,(3) |
               |                 |                 |                 |
)                |)                |                 |                 |                 |                 |
               |                 |                 |                 |
E-Port 5/44,(2923|E-Port 5/40,(3575|E-Port 5/44,(26) |E-Port 5/44,(6)  |                 |                 |
               |                 |                 |                 |
)                |)                |                 |                 |                 |                 |
               |                 |                 |                 |
                 |                 |                 |E-Port 10/16,(4) |                 |                 |
               |                 |                 |                 |
                 |                 |                 |E-Port 11/38,(3) |                 |                 |
               |                 |                 |                 |
                 |                 |                 |F-Port 12/23,(2) |                 |                 |
               |                 |                 |                 |
                 |                 |                 |F-Port 12/21,(2) |                 |                 |
               |                 |                 |                 |
```

Monitoring UCS Uplink Distribution to the Brocade SAN

Starting with Fabric OS 8.2.x release, MAPS monitors the distribution of the Cisco Unified Computing System (UCS) server vHBA logins on the physical FC uplinks from the Fabric Interconnect (FI) to the Brocade SAN switch.

Feature Overview

With multiple parallel FC uplinks from the UCS FI to the Brocade SAN switch, the FI will evenly distribute the server vHBA logins (NPIV) among the available uplinks and keep the links in a balanced state. In case of any failures on any of the uplinks, the server logins will failover to the remaining available uplinks; but once the failed uplink is restored, the server logins do not failback that results in an oversubscription of some of the uplinks and causing an imbalanced state.

Figure 2: UCS Rebalancing

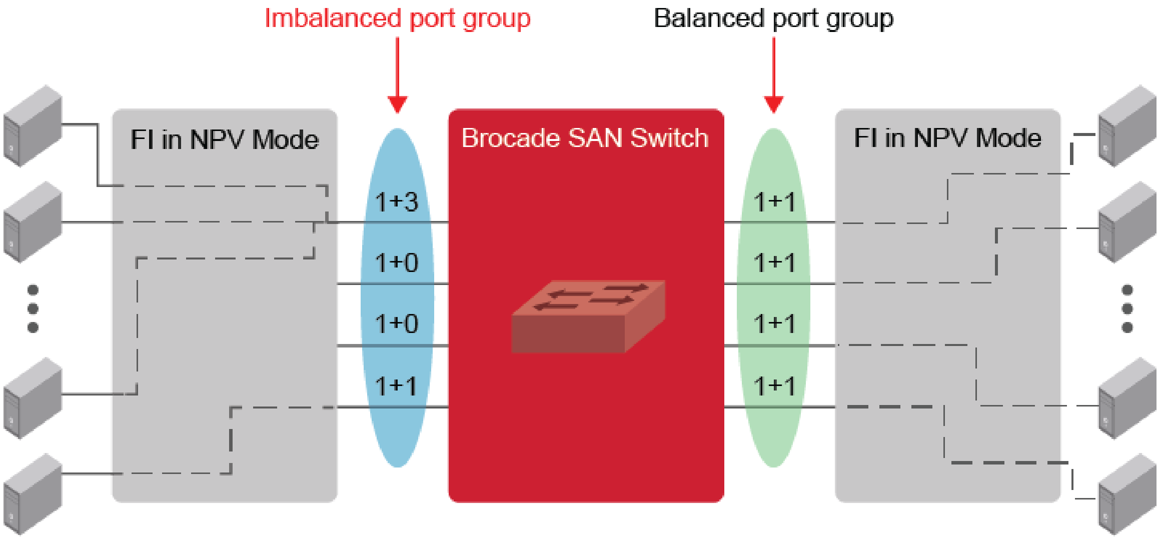

The MAPS `Monitoring UCS Uplink Distribution` feature automatically detects and group the parallel UCS FI uplinks to the Brocade switch and monitor the UCS server logins over these links for over-subscription conditions. On detecting an imbalance, MAPS generates the alerts to notify the administrators. Administrators can configure automatic rebalance action or perform a manual rebalance action.

Formation of the Dynamic Port Group

MAPS forms the dynamic port group based on the vendor OUI and node WWN of the uplink ports. Ports are dynamically added or removed from the group when they are online or offline respectively. The port groups are formed automatically, and no user configuration is required to create the port groups.

Verifying the Port Group State

You can verify the port group state using the `devicelogin --show` CLI command.

```
switch:admin> devicelogin --show
------------------------------------------------------------------------------
Node WWN                        |State      |Ports Count|Ports (Number of devices) |
------------------------------------------------------------------------------
20:07:00:de:fb:a2:87:01|BALANCED |4         |7/39(9),7/36(9),7/37(9),7/38(9)       |
```

Detecting Imbalance

MAPS monitor all the F-ports connected through the same FI as a port group. MAPS monitors the number of devices logged in on individual ports in the group, and if the device count between two ports differs by more than one, it sets the port group to IMBALANCED state.

The `DEV_LOGIN_DIST` monitoring system monitors the port group state and has the following rules:

- `BALANCED` - This means that all the ports in the group have close to an equal number of devices.
- `IMBALANCED` - This means that the difference between the number of devices logged into any two ports in the port group is greater than one. This indicates that the device logins are not distributed as evenly as possible on the ports.
- `BALANCED_FAILED` - This means that MAPS acts and fails to rebalance the group.

Once MAPS takes the rebalance action, it expects the FI to redistribute the devices among existing ports to bring back the port group into balance state. MAPS waits for sometime before it decides to set the port group state to *BALANCE* or *BALANCE_FAILED*. If it finds that the port group is still in the imbalance state, then the group state is set to *BALANCE_FAILED*. If the rebalance fails, then you can manually rebalance the ports using the `deviceLogin --rebalance` CLI command.

NOTE

- The automatic rebalance action must be disabled before performing manual rebalance.
- This feature is not supported when SAN PIN groups are configured on the UCS.

Default Rules for UCS Uplink Distribution Monitoring

If you are using custom policies to monitor UCS Uplink Distribution, use the default rules associated with the port group state. The following table shows the default rules that can be used to monitor UCS Uplink Distribution with the conditions and configurable actions.

Table 43: Default Rules for UCS Uplink Distribution Monitoring

Default Rules	Conditions	Actions
defALL_F_PORTSDEV_LOGIN_DIST_BALANCED	ALL_F_PORTS(DEV_LOGIN_DIST/NONE==BALANCED)	RASLog, SNMP, EMAIL
defALL_F_PORTSDEV_LOGIN_DIST_BALANCE_FAILED	ALL_F_PORTS(DEV_LOGIN_DIST/NONE==BALANCE_FAILED)	RASLog, SNMP, EMAIL
defALL_F_PORTSDEV_LOGIN_DIST_IMBALANCED	ALL_F_PORTS(DEV_LOGIN_DIST/NONE==IMBALANCED)	RASLog, SNMP, EMAIL, RE_BALANCE

UCS Uplink Distribution Monitoring Configuration

This section describes the method to configure the UCS uplink distribution monitoring feature.

By default, the Cisco UCS FI ports are identified and grouped based on the OUI and the node WWN of the physical ports.

1. Verify the port group and the balance state using the `devicelogin --show`

```
switch:admin> devicelogin --show
------------------------------------------------------------------------------
Node WWN                     |State             |Ports Count|Ports (Number of devices)      |
------------------------------------------------------------------------------
20:07:00:de:fb:a2:87:01|BALANCED |4          |7/39(9),7/36(9),7/37(9),7/38(9)|
```

2. Enable the MAPS monitoring of UCS FI ports using either a MAPS dashboard or activating through RASLogs.
   a) Enable the MAPS monitor by activating the default conservative/aggressive/moderate policies (not supported in base policy) or a custom policy with the `DEV_LOGIN_DIST` rules. Run the `mapsdb --show` command to note any change in the UCS port group state.

```
switch:admin> mapsdb --show
1 Dashboard Information:
------------------------
DB start time:                                      Tue Sep 26 07:12:16 2018
Active policy:                                   dflt_conservative_policy
Configured Notifications:   SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL
```

```
3.2 Rules Affecting Health:
-------------------------------

Category(Violation Count)         | RepeatCount|Rule Name
                                                                  |Execution Time         |Object
         |Triggered Value(Units)|
-----------------------------------------------------------------------------------------------------------------------
Fabric Performance Impact(2)|1                                     |
defALL_F_PORTSDEV_LOGIN_DIST_BALANCED   |09/26/18 09:54:04|F-Port 7/36 |BALANCED              |
                            |1           |defALL_F_PORTSDEV_LOGIN_DIST_IMBALANCED |09/26/18 09:44:04|F-
Port 7/36 |IMBALANCED            |
```

b) Enable the RASLog notification by activating the MAPS RASLog action by using the `mapsconfig --actions` *<RASLOG>* command.

```
switch:admin> mapsconfig --actions RASLOG
switch:admin> mapsconfig --show
Configured Notifications: SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,RASLOG

switch:admin> errshow
[MAPS-1003],1173,SLOT 2|FID 7,WARNING,X6-4_140_061,slot7 port36,F-Port 7/36,
 Condition=ALL_F_PORTS(DEV_LOGIN_DIST==IMBALANCED),Current Value:[DEV_LOGIN_DIST,IMBALANCED],
RuleName=defALL_F_PORTSDEV_LOGIN_DIST_IMBALANCED,Dashboard Category=Fabric Performance Impact.

[MAPS-1003],1174,SLOT 2|FID 7,WARNING,X6-4_140_061,slot7 port36,F-Port 7/36,
 Condition=ALL_F_PORTS(DEV_LOGIN_DIST==BALANCED),Current Value:[DEV_LOGIN_DIST,BALANCED],
RuleName=defALL_F_PORTSDEV_LOGIN_DIST_BALANCED,Dashboard Category=Fabric Performance Impact.
```

3. Enable the rebalancing operation using one of the following methods.

   a) Enable the MAPS auto-rebalancing operation by using `mapsconfig --actions RE_BALANCE` command.

```
switch:admin> mapsconfig --actions RE_BALANCE
switch:admin> mapsconfig --show
Configured Notifications: SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,RASLOG,RE_BALANCE
```

   b) Alternatively, run the `devicelogin --rebalance` command to manually initiate the rebalance operation.

```
switch:admin> devicelogin --rebalance 20:07:00:de:fb:a2:87:01
```

4. Add the new OUI to the database to support the UCS rebalance feature on ports that are not identified by the default OUI database. The OUI database must be up-to-date before executing the command. (This feature is supported only with Cisco UCS FI ports).

```
switch:admin> nodewwn --add -vendor <vendor_name> <vendor_wwn>
                    vendor_name - Only vendor name supported is Cisco.
                    vendor_wwn  - Vendor node WWN, in ':' separated format.
```

Other MAPS Monitoring Capabilities

Email Delivery Monitoring

MAPS allows monitoring of emails that are sent out of the system, and it sends a notification if an email is not delivered.

There is only one email failure RASLog generated within an hour for every email address configured. The RASlog contains the email 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 an email server problem or network issues.

An email can be sent from a system, but it might not be successfully delivered to the destination email server. In MAPS, emails are sent using the `sendmail` command in Linux. The initial steps to debug the problem are to check 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 email. 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 email 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@company.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 software allows you to monitor the air-flow direction of the fans. If two fans are running in opposite directions, then the switch status is specified 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. Fans in the mixed mode are permitted on some systems; therefore, the fan FRU is not faulted. Mixed mode is permitted for fans in the following devices:

- Brocade 6505
- Brocade 6510
- Brocade 6520
- Brocade 7840
- Brocade G620
- Brocade X6-4
- Brocade X6-8

You can check the state of the switch in the dashboard with the following rule, which is included in the 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:

```
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:                  Thu Jan 12 17:39:09 2018
Active policy:                  test_1
Configured Notifications:       RASLOG,SW_CRITICAL,SW_MARGINAL,SFP_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        |
Extension 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 |01/12/17 17:42:54|Chassis   |TRUE
        |
```

RASLog Example for Fan Air-Flow Direction Monitoring

The following is an example of a RASLog message that is sent when MAPS detects that a mismatch has occurred in the air-flow direction on the switch:

```
2017/01/12-17:39:30, [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.

2017/01/12-17:39:30, [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.

2017/01/12-17:39:30, [MAPS-1020], 1509, FID 128, WARNING, sw128_top,
Switch wide status has changed from HEALTHY to MARGINAL.
```

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 the 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 help 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 an IPEXT enabled VE tunnel

The following table explains the mapping between monitors and corresponding groups:

Table 44: Brocade IPEXT Monitoring Parameters and Groups

Monitor	Logical Groups
UTIL PKTLOSS	ALL_TUNNEL_IP_HIGH_QOS ALL_TUNNEL_IP_MED_QOS ALL_TUNNEL_IP_LOW_QOS ALL_CIRCUIT_IP_HIGH_QOS 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 Monitor 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        |
Extension 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    |
--------------------------------------------------------------------------------------
Extension   |2     |ip_cir_high_qos_pktloss|08/10/1508:58:01|Circuit/Qos 24| 0 %     |
Health (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 Monitor Dashboard Output Sample

```
Category    |Repeat |Rule Name    |Execution Time   |Object        |Triggered     |
(Rule Count)|Count  |             |                 |              |Value(Units)  |
----------------------------------------------------------------------------------
Extension   |8      |ipcktjitter  |08/10/15 09:11:01|Port/Cir 24/0 |0 %           |
Health(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 Monitor Dashboard Output Sample

```
Category      |Repeat |Rule Name    |Execution Time   |Object        |Triggered   |
(Rule Count)  |Count  |             |                 |              |Value(Units)|
----------------------------------------------------------------------------------
Extension     |1      |iptnlutil    |08/10/15 09:15:36|Tunnel 24     |47.16 %     |
Health(1)     |
----------------------------------------------------------------------------------
```

MAPS Monitoring for Extension Platforms

FCIP Quality of Service (QoS) monitoring uses policies based on a combination of data characteristics and delivery requirements to prioritize data traffic appropriately. 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 had to be retransmitted. MAPS monitors the state changes and throughput of each 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*, and 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 Extension 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 45: Use of Brocade Extension 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 Extension Health for descriptions of these parameters.
IP_EXTN_FLOW	ALL_DP

Table 46: Extension 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)
Brocade 7840 Extension Switch Brocade 7810 Extension Switch Brocade SX6 Extension Blade	• IP_EXTN_FLOW Not supported on other platforms.

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 is 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 appears 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 are edited 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  |
Extension 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)|
-----------------------------------------------------------------------------------------------------
Extension Health (1)|1     |defALL_CIRCUIT_HIGH_QOS_UTIL_75      |8/11/14 06:19:6|Circuit Qos 23/0
  |
Extension Health (1)|1     |defALL_CIRCUIT_HIGH_QOS_PKTLOSS_PER_1|8/11/14 07:02:5|Circuit Qos 23/0
  |
```

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 following table.
- `--clear` : Clears the SAM report.
- `--help` : Displays the information on how to use the commands.

Only `mapsSam --show` has additional option parameters. These parameters are listed in the following table and illustrated in the following examples:

Table 47: 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 is down. This enables you to see if a particular port is failing more often than the others. The report does not distinguish why a port is recorded as *down*, it only reports how long the port is 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 `mapsSam --show` , the report lists the following information for each port:

- Port Number
- Port type
  - AE (AE_Port)
  - DIS (disabled port)
  - DIA (D_Port)
  - DP (persistently disabled port)
  - E (E_Port)
  - F (F_Port)
  - G (G_Port)
  - N (N_Port)
  - T (Trunk port)
  - TF (F_Port trunk)
  - U (U_Port)

NOTE
Trunk ports are members of a trunk that is not the trunk master. A slave port will always be part of the trunk, and will display as a T-Port or TF-Port and master E-Port or F-Port. The *T* designation is for the slave port of the E port trunk, whereas *TF* is for F port trunk.

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%
```

Scalability Limit Monitoring

MAPS monitors changes of fabric-level monitoring systems. These systems have all 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) are triggered. The column headings in the example are edited 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        |
Extension 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 do not participate in a metaSAN. In such 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 the user 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 fabric contains 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 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 a maximum number of 16 Fibre Channel routers (FCRs) for Fabric OS 8.1.x and later releases. 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 reach the configured threshold; MAPS triggers the action configured for the rule and caches the data in the dashboard. See Default Rules for Scalability Limit Monitoring 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>16), Current Value:[BB_FCR_CNT,17], RuleName= defSWITCHBB_FCR_CNT_MAX,
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:

```
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  |
Extension 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_MAX|05/27/14 17:02:00|Switch|17
   |
(1)                 |        |                        |                  |      |
   |
```

Zone Configuration Size Monitoring

In Fabric OS 7.3.x and later releases, 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. See Default Rules for Scalability Limit Monitoring 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 to keep the switch from overloading 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. Note that the column headings in the example are edited 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
Extension 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 theinteger value in percentage, and the timebase is *NONE*.

Table 48: 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, FMS
dflt_moderate_policy	defALL_NPIV_LOGIN_PER_75	ALL_F_PORTS(DEV_NPIV_LOGINS/NONE > 75)	SNMP, RASLog, EMAIL, FMS
dflt_conservative_policy	defALL_NPIV_LOGIN_PER_90	ALL_F_PORTS(DEV_NPIV_LOGINS/NONE > 90)	SNMP, RASLog, EMAIL, FMS

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. Therefore, 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 49: Scalability Monitoring Default Rules

Policy Name	Rule Name	Rule Condition	Rule Action
dflt_conservative_policy	defSWITCHL2_DEVCNT_PER_90 defSWITCHLSAN_DEVCNT_PER_90 defSWITCHZONE_CFGSZ_PER_90 defSWITCHBB_FCR_CNT_MAX	L2_DEVCNT_PER greater than 90 LSAN_DEVCNT_PER greater than 90 ZONE_CFGSZ_PER greater than 90 BB_FCR_CNT greater than 16	RASLog, SNMP, EMAIL, FMS
dflt_moderate_policy	defSWITCHL2_DEVCNT_PER_75 defSWITCHLSAN_DEVCNT_PER_75 defSWITCHZONE_CFGSZ_PER_80 defSWITCHBB_FCR_CNT_MAX	L2_DEVCNT_PER greater than 75 LSAN_DEVCNT_PER greater than 75 ZONE_CFGSZ_PER greater than 80 BB_FCR_CNT greater than 16	RASLog, SNMP, EMAIL, FMS
dflt_aggressive_policy	defSWITCHL2_DEVCNT_PER_60 defSWITCHLSAN_DEVCNT_PER_60 defSWITCHZONE_CFGSZ_PER_70 defSWITCHBB_FCR_CNT_MAX	L2_DEVCNT_PER greater than 60 LSAN_DEVCNT_PER greater than 60 ZONE_CFGSZ_PER greater than 70 BB_FCR_CNT greater than 16	RASLog, SNMP, EMAIL, FMS

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.

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.

See Security Health for a list and descriptions of parameters that are used for monitoring certificates in MAPS.

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 signature of the certificate issuer.
- 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 not 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 the HEALTHY state.

The following certificates can be imported into the system:

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

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

Default Policies 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:

Table 50: Default Rules Added to Default Policies

Default Policy	Rule Name	Condition	Actions
Conservative	`defCHASSISCERT_VALIDITY_15`	`ALL_CERTS(DAYS_TO_EXPIRE/NONE)<15)`	RASLog, SNMP, EMAIL, FMS

Default Policy	Rule Name	Condition	Actions
	defCHASSISCERTS_EXPIRED	CHASSIS (EXPIRED_CERTS/NONE>0)	RASLog, SNMP, EMAIL, FMS, SW_CRITICAL, SW_MARGINAL
Moderate	defCHASSISCERT_VALIDITY_20	ALL_CERTS(DAYS_TO_EXPIRE/NONE)<20)	RASLog, SNMP, EMAIL, FMS
	defCHASSISCERTS_EXPIRED	CHASSIS (EXPIRED_CERTS/NONE>0)	RASLog, SNMP, EMAIL, FMS, SW_CRITICAL, SW_MARGINAL
Aggressive	defCHASSISCERT_VALIDITY_30	ALL_CERTS(DAYS_TO_EXPIRE/NONE)<30)	RASLog, SNMP, EMAIL, FMS
	defCHASSISCERTS_EXPIRED	CHASSIS (EXPIRED_CERTS/NONE>0)	RASLog, SNMP, EMAIL, FMS, SW_CRITICAL, SW_MARGINAL

Certificate Monitor Rule Creation

The rule creation for the 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.
```

Updating Monitoring Policies for Devices with Four PSUs

If you have a chassis that supports more than two power supply units (PSUs), then 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

Perform the following steps to change the active MAPS power supply *switchstatus* policy settings:

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 to dial into a support center to report system problems.

MAPS Threshold Values

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, perform 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 and 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* .

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 51: Default Fabric State Change Monitoring Thresholds and Actions

Monitoring Statistic	Fabric State Change Monitoring Threshold Values by Policy			Actions	RASLog IDs
	Aggressive	Moderate	Conservative		
Domain ID change (DID_CHG)	1	1	1	RASLog, SNMP, EMAIL	MAPS-2324 to 2327
Fabric logins (FLOGI)	4	6	8	RASLog, SNMP, EMAIL	MAPS-2340 to 2343
Fabric reconfiguration (FAB_CFG)	1	2	4	RASLog, SNMP, EMAIL	MAPS-2320 to 2323
E_Ports down (EPORT_DOWN)	1	2	4	RASLog, SNMP, EMAIL	MAPS-2316 to 2319
Segmentation changes (FAB_SEG)	1	2	4	RASLog, SNMP, EMAIL	MAPS-2328 to 2331
Zone changes (ZONE_CHG)	2	5	10	RASLog, SNMP, EMAIL	MAPS-2336 to 2339
L2 Device Count (L2_DEVCNT_PER)	60	75	90	RASLog, SNMP, EMAIL	MAPS-2344 to 2347
LSAN Device Count (LSAN_DEVCNT_PER)	60	75	90	RASLog, SNMP, EMAIL	MAPS-2348 to 2351
Zone Configuration size (ZONE_CFGSZ_PER)	70	80	90	RASLog, SNMP, EMAIL	MAPS-2352 to 2355
FCR Count (BB_FCR_CNT)	16	16	16	RASLog, SNMP, EMAIL	MAPS-2356 to 2359

Extension Monitoring Thresholds

These Extension 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 52: Default Extension Monitoring Thresholds and Actions

Monitoring Statistic	Units	Extension Monitoring Threshold Values by Policy			Actions	RASLog IDs
		Aggressive	Moderate	Conservative		
CIR_STATE	Circuit state changes per minute	0	3	5	RASLog, SNMP, EMAIL, FENCE, DECOM, FMS	MAPS-2140 to 2143
CIR_UTIL	Circuit utilization percentage per hour	60	75	90	RASLog, SNMP, EMAIL, FMS	MAPS-2144 to 2147
CIR_PKTLOSS	Circuit packet loss percentage per minute	0.05	0.1	0.5	RASLog, SNMP, EMAIL, FMS	MAPS-2148 to 2151
RTT	Circuit round-trip times delay in milliseconds	250	250	250	RASLog, SNMP, EMAIL, FMS	MAPS-2428 to 2431

Monitoring Statistic	Units	Extension Monitoring Threshold Values by Policy			Actions	RASLog IDs
		Aggressive	Moderate	Conservative		
JITTER	Circuit 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, FMS	MAPS-2432 to 2435
TUNNEL_STATE	Tunnel state changes per minute	0	1	3	RASLog, SNMP, EMAIL, FENCE, DECOM, FMS	MAPS-2028 to 2031
TUNNEL_UTIL	Tunnel QoS percentage per hour	50	75	90	RASLog, SNMP, EMAIL, FMS	MAPS-2048 to 2051
PKTLOSS	QoS Packet loss percentage per minute	0.05	0.1	0.5	RASLog, SNMP, EMAIL, FMS	MAPS-2424 to 2427
IP_EXTN_FLOW	Number of IP extension flows through a DP	325	400	475	RASLog, SNMP, EMAIL, FMS	MAPS-2956 to 2959
QOS_UTIL	QoS utilization percentage	0	1	1	RASLog, SNMP, EMAIL, FMS	MAPS-2420 to 2423
CIR_QOS_UTIL	Circuit QoS utilization percentage	0.05	0.1	0.5	RASLog, SNMP, EMAIL, FMS	MAPS-2440 to 2443
CIR_QOS_PKTLOSS	Circuit QoS packet loss percentage	0.05	0.1	0.5	RASLog, SNMP, EMAIL, FMS	MAPS-2444 to 2447
TUNNEL_IP_UTIL	Tunnel IP utilization percentage	50	75	90	RASLog, SNMP, EMAIL, FMS	MAPS-2700 to 2703
CIR_IP_UTIL	Circuit IP utilization percentage	60	75	90	RASLog, SNMP, EMAIL, FMS	MAPS-2704 to 2707
CIR_IP_PKTLOSS	Circuit IP packet loss percentage	0.05	0.1	0.5	RASLog, SNMP, EMAIL, FMS	MAPS-2708 to 2711
IP_RTT	IP circuit round-trip times	250	250	250	RASLog, SNMP, EMAIL, FMS	MAPS-2712 to 2715
IP_JITTER	IP circuit jitter	5	15	20	RASLog, SNMP, EMAIL, FMS	MAPS-2716 to 2719

Table 53: Extension Gigabit Ethernet (GE) Port Health Monitoring Thresholds and Actions

		Extension Monitoring Threshold Values by Policy				
Monitoring Statistic	Units	Aggressive	Moderate	Conservative	Actions	RASLog IDs
GE_CRC	The number of times an invalid cyclic redundancy check (CRC) error occurs on a GE port or a frame that computes to an invalid CRC.	0	1	1	RASLog, SNMP, EMAIL, FMS	MAPS-2060 to 2063
GE_LOS_OF_SIG	The number of times that a signal loss occurs in offline GE ports.	0	1	1	RASLog, SNMP, EMAIL, FMS	MAPS-2064 to 2067

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 54: FRU Monitoring Statistics, States, and Actions

Monitored Statistic	Supported States	Actions	RASLog IDs
Power Supply (PS_STATE)	ON, OUT, FAULTY	RASLog, SNMP, EMAIL, FMS	MAPS-2168 to MAPS-2171
Fan (FAN_STATE)	ON, OUT, FAULTY	RASLog, SNMP, EMAIL, FMS	MAPS-2172 to MAPS-2175
Slot (BLADE_STATE)	ON, OFF, OUT, FAULTY	RASLog, SNMP, EMAIL, FMS	MAPS-2184 to MAPS-2187
SFP (SFP_STATE) for port SFP (SFP_STATE) for the Ethernet port	IN, OUT, FAULTY IN, OUT, FAULTY	RASLog, SNMP, EMAIL, FMS RASLog, SNMP, EMAIL, FMS	MAPS-2180 to MAPS-2183 MAPS-3000 to MAPS-3003
WWN (WWN)	ON, OUT, FAULTY	RASLog, SNMP, EMAIL, FMS	MAPS-2188 to MAPS-2191

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.

Some monitors support *time-duration* timebases, and some support the *None* timebase as shown in the following table:

Monitors that Support Minute, Hour, and Day Timebases	Monitors that Support Only None Timebase
C3 Time-Outs	SFP Current
CRC Errors	SFP Power On Hours
Invalid Transmit Words	SFP Receive Power
Link Failure	SFP Temperature
Link Reset	SFP Transmit Power
Loss of Signal	SFP Voltage

Monitors that Support Minute, Hour, and Day Timebases	Monitors that Support Only None Timebase
Loss of Sync	
Protocol Errors	
State Change	

Back-End Port Monitoring Thresholds

All back-end port monitors support the *Minute*, *Hour*, *Day*, and *Week* timebases.

The following table lists the errors for which MAPS monitors on back-end ports, the trigger thresholds, and the default actions to be taken when the threshold is crossed.

Table 55: Back-End Port Monitoring Default Thresholds and Actions

Errors	Thresholds	Actions	RASLog IDs
CRC	• 10 per 5 minutes • 100 per day	RASLog, SNMP, EMAIL, FENCE, DECOM, FMS	MAPS-2108 to MAPS-2111
LR	• 10 per 5 minutes • 100 per day	RASLog, SNMP, EMAIL, FMS	MAPS-2116 to MAPS-2119
ITW	• 10 per 5 minutes • 100 per day	RASLog, SNMP, EMAIL, FENCE, DECOM, FMS	MAPS-2112 to MAPS-2115
BAD_OS	• 10 per 5 minutes • 100 per day	RASLog, SNMP, EMAIL, FMS	MAPS-2120 to MAPS-2123
FRM_LONG	• 10 per 5 minutes • 100 per day	RASLog, SNMP, EMAIL, FMS	MAPS-2124 to MAPS-2127
FRM_TRUNC	• 10 per 5 minutes • 100 per day	RASLog, SNMP, EMAIL, FMS	MAPS-2128 to MAPS-2131

D_Port Monitoring Thresholds

The following tables list the *default D_Port monitoring* threshold values and actions for each default policy:

Aggressive Policy Defaults for D_Port Monitoring

Table 56: Aggressive Policy Default D_Port Monitoring Threshold Values and Actions

Monitoring Statistic	Unit	Threshold	Actions	RASLog IDs
CRC Errors (defALL_D_PORTSCRC_1)	Min	1	EMAIL, SNMP, RASLog	MAPS-2004 to 2007
Invalid Transmit Words (defALL_D_PORTSITW_1)	Min	1	EMAIL, SNMP, RASLog	MAPS-2008 to 2011
Link Failure (defALL_D_PORTSLF_1)	Min	1	EMAIL, SNMP, RASLog	MAPS-2016 to 2019
Sync Loss (defALL_D_PORTSLOSS_SYNC_1)	Min	1	EMAIL, SNMP, RASLog	MAPS-2012 to 2015

Monitoring Statistic	Unit	Threshold	Actions	RASLog IDs
CRC Errors (defALL_D_PORTSCRC_H30)	Hour	30	EMAIL, SNMP, RASLog	MAPS-2004 to 2007
Invalid Transmit Words (defALL_D_PORTSITW_H30)	Hour	30	EMAIL, SNMP, RASLog	MAPS-2008 to 2011
Link Failure (defALL_D_PORTSLF_H30)	Hour	30	EMAIL, SNMP, RASLog	MAPS-2016 to 2019
Sync Loss (defALL_D_PORTSLOSS_SYNC_H30)	Hour	30	EMAIL, SNMP, RASLog	MAPS-2012 to 2015
CRC Errors (defALL_D_PORTSCRC_D500)	Day	500	EMAIL, SNMP, RASLog	MAPS-2004 to 2007
Invalid Transmit Words (defALL_D_PORTSITW_D500)	Day	500	EMAIL, SNMP, RASLog	MAPS-2008 to 2011
Link Failure (defALL_D_PORTSLF_D500)	Day	500	EMAIL, SNMP, RASLog	MAPS-2016 to 2019
Sync Loss (defALL_D_PORTSLOSS_SYNC_D500)	Day	500	EMAIL, SNMP, RASLog	MAPS-2012 to 2015

Moderate Policy Defaults for D_Port Monitoring

Table 57: Moderate Policy Default D_Port Monitoring Threshold Values and Actions

Monitoring Statistic	Unit	Threshold	Actions	RASLog IDs
CRC Errors (defALL_D_PORTSCRC_2)	Min	2	EMAIL, SNMP, RASLog, FMS	MAPS-2004 to 2007
Invalid Transmit Words (defALL_D_PORTSITW_2)	Min	2	EMAIL, SNMP, RASLog, FMS	MAPS-2008 to 2011
Link Failure (defALL_D_PORTSLF_2)	Min	2	EMAIL, SNMP, RASLog, FMS	MAPS-2016 to 2019
Sync Loss (defALL_D_PORTSLOSS_SYNC_2)	Min	2	EMAIL, SNMP, RASLog, FMS	MAPS-2012 to 2015
CRC Errors (defALL_D_PORTSCRC_H60)	Hour	60	EMAIL, SNMP, RASLog, FMS	MAPS-2004 to 2007
Invalid Transmit Words (defALL_D_PORTSITW_H60)	Hour	60	EMAIL, SNMP, RASLog, FMS	MAPS-2008 to 2011
Link Failure (defALL_D_PORTSLF_H60)	Hour	60	EMAIL, SNMP, RASLog, FMS	MAPS-2016 to 2019
Sync Loss (defALL_D_PORTSLOSS_SYNC_H60)	Hour	60	EMAIL, SNMP, RASLog, FMS	MAPS-2012 to 2015
CRC Errors (defALL_D_PORTSCRC_D1000)	Day	1000	EMAIL, SNMP, RASLog, FMS	MAPS-2004 to 2007
Invalid Transmit Words (defALL_D_PORTSITW_D1000)	Day	1000	EMAIL, SNMP, RASLog, FMS	MAPS-2008 to 2011

Monitoring Statistic	Unit	Threshold	Actions	RASLog IDs
Link Failure (defALL_D_PORTSLF_D1000)	Day	1000	EMAIL, SNMP, RASLog, FMS	MAPS-2016 to 2019
Sync Loss (defALL_D_PORTSLOSS_SYNC_D1000)	Day	1000	EMAIL, SNMP, RASLog, FMS	MAPS-2012 to 2015

Conservative Policy Defaults for D_Port Monitoring

Table 58: Conservative Policy Default D_Port Monitoring Threshold Values and Actions

Monitoring Statistic	Unit	Threshold	Actions	RASLog IDs
CRC Errors (defALL_D_PORTSCRC_3)	Min	3	EMAIL, SNMP, RASLog, FMS	MAPS-2004 to 2007
Invalid Transmit Words (defALL_D_PORTSITW_3)	Min	3	EMAIL, SNMP, RASLog, FMS	MAPS-2008 to 2011
Link Failure (defALL_D_PORTSLF_3)	Min	3	EMAIL, SNMP, RASLog, FMS	MAPS-2016 to 2019
Sync Loss (defALL_D_PORTSLOSS_SYNC_3)	Min	3	EMAIL, SNMP, RASLog, FMS	MAPS-2012 to 2015
CRC Errors (defALL_D_PORTSCRC_H90)	Hour	90	EMAIL, SNMP, RASLog, FMS	MAPS-2004 to 2007
Invalid Transmit Words (defALL_D_PORTSITW_H90)	Hour	90	EMAIL, SNMP, RASLog, FMS	MAPS-2008 to 2011
Link Failure (defALL_D_PORTSLF_H90)	Hour	90	EMAIL, SNMP, RASLog, FMS	MAPS-2016 to 2019
Sync Loss (defALL_D_PORTSLOSS_SYNC_H90)	Hour	90	EMAIL, SNMP, RASLog, FMS	MAPS-2012 to 2015
CRC Errors (defALL_D_PORTSCRC_D1500)	Day	1500	EMAIL, SNMP, RASLog, FMS	MAPS-2004 to 2007
Invalid Transmit Words (defALL_D_PORTSITW_D1500)	Day	1500	EMAIL, SNMP, RASLog, FMS	MAPS-2008 to 2011
Link Failure (defALL_D_PORTSLF_D1500)	Day	1500	EMAIL, SNMP, RASLog, FMS	MAPS-2016 to 2019
Sync Loss (defALL_D_PORTSLOSS_SYNC_D1500)	Day	1500	EMAIL, SNMP, RASLog, FMS	MAPS-2012 to 2015

E_Port Monitoring Thresholds

The following table lists the *default E_Port* monitoring threshold values and actions:

Table 59: Default E_Port Monitoring Threshold Values and Actions

	E_Port Monitoring Threshold Values by Policy				
Monitoring Statistic	Aggressive	Moderate	Conservative	Actions	RASLog IDs
C3 Time out (C3TX_TO)	5	10	20	RASLog, SNMP, EMAIL, FENCE, DECOM, FMS	MAPS-2036 to 2039
CRC Errors (CRC)	Low: 0 High: 2	Low: 10 High: 20	Low: 21 High: 40	Low: EMAIL, SNMP, RASLog High: EMAIL, SNMP, FENCE, DECOM	MAPS-2004 to 2007
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	MAPS-2008 to 2011
Link Reset (LR)	Low: 2 High: 4	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLog High: EMAIL, SNMP, FENCE, DECOM	MAPS-2031 to 2035
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	MAPS-2028 to 2031
Loss of signal (LOSS_SIGNAL)	0	3	5	EMAIL, SNMP, RASLog, FMS	MAPS-2020 to 2023
Link Failure (LF)	0	3	5	EMAIL, SNMP, RASLog, FMS	MAPS-2016 to 2019
Sync Loss (LOSS_SYNC)	0	3	5	EMAIL, SNMP, RASLog, FMS	MAPS-2012 to 2015
Transmit errors per minute (TX)	95	95	95	EMAIL, SNMP, RASLog, FMS	MAPS-2044 to 2047
Receive errors per minute (RX)	95	95	95	EMAIL, SNMP, RASLog, FMS	MAPS-2040 to 2043
Utility errors per minute (UTIL)	95	95	95	EMAIL, SNMP, RASLog, FMS	MAPS-2048 to 2051

F_Port Monitoring Thresholds

The following tables list the default monitoring thresholds and actions for host F_Ports and target F_Ports:

NOTE
If an F_Port cannot be identified as either a host or a target, then the thresholds for it are the same as those for Host F_Ports.

Host F_Port Default Monitoring Thresholds

Table 60: Default Host F_Port Monitoring Threshold Values and Actions

Monitoring Statistic	Host F_Port Monitoring Threshold Values by Policy			Actions	RASLog IDs
	Aggressive	Moderate	Conservative		
C3 Time out (C3TXTO) C3TXTO is not monitored for resident N_Ports or F_Ports with any default policy in AG mode.	Low: 2 High: 4	Low: 3 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLog, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS	MAPS-2036 to 2039
CRC Errors (CRC)	Low: 0 High: 2	Low: 10 High: 20	Low: 21 High: 40	Low: EMAIL, SNMP, RASLog, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS	MAPS-2004 to 2007
Invalid Transmit Words (ITW)	Low: 15 High: 20	Low: 21 High: 40	Low: 41 High: 80	Low: EMAIL, SNMP, RASLog, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS	MAPS-2008 to 2011
Link Reset (LR)	Low: 2 High: 4	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLog, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS	MAPS-2032 to 2035
State Change (STATE_CHG)	Low: 2 High: 4	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLog, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS	MAPS-2028 to 2031
Loss of signal (LOSS_SIGNAL)	0	3	5	EMAIL, SNMP, RASLog, FMS	MAPS-2020 to 2023
Link Failure (LF)	0	3	5	EMAIL, SNMP, RASLog, FMS	MAPS-2016 to 2019
Sync Loss (LOSS_SYNC)	0	3	5	EMAIL, SNMP, RASLog, FMS	MAPS-2012 to 2015
Transmit errors per minute (TX)	95	95	95	EMAIL, SNMP, RASLog, FMS	MAPS-2044 to 2027
Receive errors per minute (RX)	95	95	95	EMAIL, SNMP, RASLog, FMS	MAPS-2040 to 2043
Utility errors per minute (UTIL)	95	95	95	EMAIL, SNMP, RASLog, FMS	MAPS-2048 to 2051

Target F_Port Default Monitoring Thresholds

Table 61: Default Target F_Port Monitoring Threshold Values and Actions

Monitoring Statistic	Target F_Port Monitoring Threshold Values by Policy			Actions	RASLog IDs
	Aggressive	Moderate	Conservative		
C3 Time out (C3TX_TO) C3TXTO is not monitored for resident N_Ports or F_Ports with any default policy in AG mode.	Low: 0 High: 2	Low: 3 High: 5	Low: 6 High: 10	Low: EMAIL, SNMP, RASLog, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS	MAPS-2036 to 2039
CRC Errors (CRC)	Low: 0 High: 2	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLog, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS	MAPS-2004 to 2007
Invalid Transmit Words (ITW)	Low: 5 High: 10	Low: 11 High: 20	Low: 21 High: 40	Low: EMAIL, SNMP, RASLog, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS	MAPS-2008 to 2011
Link Reset (LR)	Low: 0 High: 2	Low: 3 High: 5	Low: 6 High: 10	Low: EMAIL, SNMP, RASLog, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS	MAPS-2032 to 2035
State Change (STATE_CHG)	Low: 0 High: 2	Low: 3 High: 7	Low: 8 High: 15	Low: EMAIL, SNMP, RASLog, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS	MAPS-2028 to 2031
Loss of signal (LOSS_SIGNAL)	0	3	5	EMAIL, SNMP, RASLog, FMS	MAPS-2020 to 2023
Link Failure (LF)	0	3	5	EMAIL, SNMP, RASLog, FMS	MAPS-2016 to 2019
Sync Loss (LOSS_SYNC)	0	3	5	EMAIL, SNMP, RASLog, FMS	MAPS-2012 to 2015
Transmit errors per minute (TX)	95	95	95	EMAIL, SNMP, RASLog, FMS	MAPS-2044 to 2027
Receive errors per minute (RX)	95	95	95	EMAIL, SNMP, RASLog, FMS	MAPS-2040 to 2043
Utility errors per minute (UTIL)	95	95	95	EMAIL, SNMP, RASLog, FMS	MAPS-2048 to 2051

Front-End Encryption Port Monitoring Thresholds

The following tables list the error-port monitoring threshold values and actions:

Table 62: Default Front-End Encryption Port Monitoring Threshold Values and Actions

Monitoring Statistic	Error Port Monitoring Threshold Values by Policy			Actions	RASLog IDs
	Aggressive	Moderate	Conservative		
Block errors per minute (def_E_PORTS_ENCR_BLK_ERR)	0	0	0	RASLog, SNMP, EMAIL, FENCE, DECOM, FMS	MAPS-2092 to 2095
Discard errors per minute (def_E_PORTS_ENCR_DISC_ERR)	30	30	30	RASLog, SNMP, EMAIL, FENCE, DECOM, FMS	MAPS-2096 to 2099
Short frame errors per minute (def_E_PORTS_ENCR_SHRT_FRM)	0	0	0	RASLog, SNMP, EMAIL, FMS	MAPS-2100 to 2103

NOTE
You cannot create user-defined rules with these monitoring statistics.

Non-F_Port Monitoring Thresholds

The following table lists the default non-F_Port monitoring threshold values and actions:

Table 63: Default Non-F_Port Monitoring Threshold Values and Actions

Monitoring Statistic	Non-F_Port Monitoring Threshold Values by Policy			Actions	RASLog IDs
	Aggressive	Moderate	Conservative		
C3 Time out (C3TX_TO) C3TXTO is not monitored for resident N_Ports or F_Ports with any default policy in AG mode.	N/A	N/A	N/A	N/A	MAPS-2036 to 2039
CRC Errors (CRC)	Low: 0 High: 2	Low: 10 High: 20	Low: 21 High: 40	Low: EMAIL, SNMP, RASLog, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS	MAPS-2004 to 2007
Invalid Transmit Words (ITW)	Low: 15 High: 20	Low: 21 High: 40	Low: 41 High: 80	Low: EMAIL, SNMP, RASLog, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS	MAPS-2008 to 2011
Link Reset (LR)	Low: 2 High: 4	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLog, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS	MAPS-2032 to 2035
State Change (STATE_CHG)	Low: 2 High: 4	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLog, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS	MAPS-2028 to 2031

	Non-F_Port Monitoring Threshold Values by Policy				
Monitoring Statistic	Aggressive	Moderate	Conservative	Actions	RASLog IDs
Loss of signal (LOSS_SIGNAL)	0	3	5	EMAIL, SNMP, RASLog, FMS	MAPS-2020 to 2023
Link Failure (LF)	0	3	5	EMAIL, SNMP, RASLog, FMS	MAPS-2016 to 2019
Sync Loss (LOSS_SYNC)	0	3	5	EMAIL, SNMP, RASLog, FMS	MAPS-2012 to 2015
Transmit errors per minute (TX)	95	95	95	EMAIL, SNMP, RASLog, FMS	MAPS-2044 to 2027

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 64: Default Resource Monitoring Thresholds and Actions

	Resource Monitoring Threshold Values by Policy				
Monitoring Statistic	Aggressive	Moderate	Conservative	Actions	RASLog IDs
Flash memory percentage used (FLASH_USAGE)	90	90	90	RASLog, SNMP, EMAIL, FMS	MAPS-2152 to 2155
CPU percentage used (CPU)	80	80	80	RASLog, SNMP, EMAIL, FMS	MAPS-2156 to 2159
Memory percentage used (MEMORY_USAGE)	75	75	75	RASLog, SNMP, EMAIL, FMS	MAPS-2160 to 2163
Ethernet management port state (ETH_MGMT_PORT_STATE)	Up/Down	Up/Down	Up/Down	RASLog, SNMP, EMAIL, FMS	MAPS-2408 to 2411
Temperature Sensor (TEMP)	OUT_OF_RANGE	OUT_OF_RANGE	OUT_OF_RANGE	RASLog, SNMP, EMAIL, FMS	MAPS-2164 to 2167
VTAP I/Os per second (VTAP_IOPS)	250000 seconds	250000 seconds	250000 seconds	RASLog, SNMP, EMAIL, FMS, UNINSTALL_VTAP	MAPS-2668 to 2671

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 65: Default Security Monitoring Thresholds and Actions

Monitoring Statistic	Security Monitoring Threshold Values by Policy			Actions	RASLog IDs
	Aggressive	Moderate	Conservative		
DCC violations (SEC_DCC)	0	2	4	RASLog, SNMP, EMAIL	MAPS-2208 to 2211
HTTP violation* (SEC_HTTP)	0	2	4	RASLog, SNMP, EMAIL	MAPS-2200 to 2203
Illegal command (SEC_CMD)	0	2	4	RASLog, SNMP, EMAIL	MAPS-2236 to 2239
Incompatible security DB (SEC_IDB)	0	2	4	RASLog, SNMP, EMAIL	MAPS-2232 to 2235
Login violations* (SEC_LV)	0	2	4	RASLog, SNMP, EMAIL	MAPS-2212 to 2215
Invalid certifications (SEC_CERT)	0	2	4	RASLog, SNMP, EMAIL	MAPS-2216 to 2219
No-FCS (SEC_FCS)	0	2	4	RASLog, SNMP, EMAIL	MAPS-2228 to 2231
SCC violations (SEC_SCC)	0	2	4	RASLog, SNMP, EMAIL	MAPS-2204 to 2207
SLAP failures (SEC_AUTH_FAIL)	0	2	4	RASLog, SNMP, EMAIL	MAPS-2224 to 2227
Telnet violations* (SEC_TELNET)	0	2	4	RASLog, SNMP, EMAIL	MAPS-2196 to 2199
TS out of sync* (SEC_TS)	1 per hour 2 per day	2 per hour 4 per day	4 per hour 10 per day	RASLog, SNMP, EMAIL	MAPS-2220 to 2223
Expired certifications (EXPIRED_CERTS)	0	0	0	RASLog, SNMP, EMAIL, SW_CRITICAL, SW_MARGINAL, FMS	MAPS-2244 to 2247
Number of days to expire (DAYS_TO_EXPIRE)	0	0	0	RASLog, SNMP, EMAIL, FMS	MAPS-2240 to 2243

NOTE
The monitoring systems with a (*) are monitored only in the default switch.

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, 16, 25, and 32Gb/s SFP Monitoring Threshold Defaults

The following table lists the default thresholds for 10, 16, 25, and 32Gb/s SFPs. It also lists thresholds for 10Gb/s and 25Gb/s in the FCoE or Ethernet port mode.

NOTE
The 8th and 9th columns in the following table are applicable to the Brocade FC32-64 port blade and Brocade G630 switch platforms in the FCoE or Ethernet port mode.

NOTE
The threshold values in the following table are for minimum (lower) and maximum (upper) thresholds; where there is one value, it is the maximum threshold.

Table 66: Default SFP Monitoring Thresholds and Actions for 10, 16, 25, and 32Gb/s SFPs

	SFP Monitoring Thresholds for All Policies										
Monitoring Statistic	**10Gb/s SWL**	**10Gb/s LWL**	**16Gb/s SWL**	**16Gb/s LWL**	**25Km 16Gb/s LWL**	**32Gb/s SWL**	**32Gb/s LWL**	**10Gb/s FCoE or Ethernet**	**25Gb/s FCoE or Ethernet**	**Actions**	**RASLog IDs**
Current (CURRENT) (mA)	10	95	12	70	90	12	60	10	10	SFP_MARGINAL, RASLog, SNMP, EMAIL, FMS	MAPS-2269 to 2271
Receive Power (RXP) (µW)	1999	2230	1259	1995	2338	1259	1995	31 to 2000	60 to 2187	SFP_MARGINAL, RASLog, SNMP, EMAIL, FMS	MAPS-2272 to 2275
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	-5 to 90	-5 to 75	SFP_MARGINAL, RASLog, SNMP, EMAIL, FMS	MAPS-2260 to 2263
Transmit Power (TXP) (µW)	1999	2230	1259	1995	4466	1259	1584	3000 to 3600	1259	SFP_MARGINAL, RASLog, SNMP, EMAIL, FMS	MAPS-2276 to 2279
Voltage (VOLTAGE) (mV)	3000 to 3600	2970 to 3600	3000 to 3600	3000 to 3600	2850 to 3750	3000 to 3600	3000 to 3600	3000 to 3600	2970 to 3630	SFP_MARGINAL, RASLog, SNMP, EMAIL, FMS	MAPS-2264 to 2267

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.

> **NOTE**
> The 6th and 7th columns in the following table are applicable to the Brocade FC32-64 port blade and Brocade G630 switch platforms in the FCoE or Ethernet port mode.

> **NOTE**
> The threshold values in the following table are for minimum and maximum thresholds; where there is one value, it is the maximum threshold.

Table 67: Default SFP Monitoring Thresholds and Actions for QSFPs and All Other SFPs

	QSFP and Other SFP Monitoring Thresholds for All Policies							
Monitoring Statistic	**100M 16Gb/s SWL QSFP**	**2K QSFP**	**32Gb/s SWL QSFP**	**QSFP**	**All Other SFPs**	**40Gb/s QSFP FCoE or Ethernet**	**100Gb/s QSFP FCoE or Ethernet**	**Actions**	**RASLog IDs**
Current (CURRENT) (mA)	10	39	10	10	50	10	2 to 10	RASLog, SNMP,EMAIL, FMS	MAPS-2269 to 2271
Receive Power (RXP) (µW)	2187	2000	3400	2180	5000	44 to 2188	60 to 2187	RASLog, SNMP, EMAIL, FMS	MAPS-2272 to 2275
Temperature (TEMP) (°C)	-5 to 85	-15 to 85	-5 to 75	-5 to 85	-13 to 85	-5 to 75	-5 to 75	RASLog, SNMP, EMAIL, FMS	MAPS-2260 to 2263

	QSFP and Other SFP Monitoring Thresholds for All Policies								
Monitoring Statistic	100M 16Gb/s SWL QSFP	2K QSFP	32Gb/s SWL QSFP	QSFP	All Other SFPs	40Gb/s QSFP FCoE or Ethernet	100Gb/s QSFP FCoE or Ethernet	Actions	RASLog IDs
Transmit Power (TXP) (µW)	—	—	—	—	5000	2970 to 3630	48 to 3467	RASLog, SNMP, EMAIL, FMS	MAPS-2276 to 2279
Voltage (VOLTAGE) (mV)	2970 to 3630	2900 to 3600	2970 to 3630	2940 to 3600	2960 to 3630	2970 to 3630	3970 to 3630	RASLog, SNMP, EMAIL, FMS	MAPS-2264 to 2267

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* .

> **NOTE**
> Whenever *IO_PERF_IMPACT* is used, *IO_LATENCY_CLEAR* must also be included in the active policy to clear the latency record.

Table 68: Default Fabric Performance Impact Latency Monitoring Threshold Values and Actions

Monitoring Statistic	Value (Y/N) for All Policies	Actions	RASLog IDs
Fabric Performance Impact (DEV_LATENCY_IMPACT)	IO_FRAME_LOSS IO_PERF_IMPACT IO_LATENCY_CLEAR	RASLog, SNMP, EMAIL, SDDQ, TOGGLE RASLog, SNMP, EMAIL, SDDQ, TOGGLE RASLog, SNMP, EMAIL	MAPS-2068 to 2071

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 69: Default Fabric Performance Impact RX, TX, and UTIL Monitoring Threshold Values and Actions

	FPI Monitoring Threshold Values by Policy				
Monitoring Statistic	Aggressive	Moderate	Conservative	Actions	RASLog IDs
Receive bandwidth usage percentage (RX)	60	75	90	EMAIL, SNMP, RASLog	MAPS-2040 to 2043
Transmit bandwidth usage percentage (TX)	60	75	90	EMAIL, SNMP, RASLog	MAPS-2044 to 2047
Trunk utilization percentage (UTIL)	60	75	90	EMAIL, SNMP, RASLog	MAPS-2048 to 2051

Monitoring Statistic	FPI Monitoring Threshold Values by Policy			Actions	RASLog IDs
	Aggressive	Moderate	Conservative		
Back-end port latency impact (BE_LATENCY_IMPACT)	0	0	0	EMAIL, SNMP, RASLog , FMS	MAPS-2672 to 2675
Initiator to target flow ratio (IT_FLOW)	8	16	32	EMAIL, SNMP, RASLog , FMS	MAPS-2936 to 2939
Device logins distribution (DEV_LOGIN_DIST)	BALANCED	IMBALANCED	BALANCE_FAILED	EMAIL, SNMP, RASLog , FMS, RE_BALANCE	MAPS-3004 to 3007

Switch Status Policy Monitoring Thresholds

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. The *Switch Status Monitor* supports only the *None* timebase.

The switch status policy sends alerts every minute. Therefore, if a monitoring statistic is faulty for less than a minute, the alert for that statistic is not sent. 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 for each policy:

Switch Status Monitoring Thresholds and Actions for the Default Aggressive Policy

The following table lists the default *Switch Status Monitoring* thresholds and actions for the default aggressive policy:

Monitoring Statistic	Description	Switch Status Threshold Values			Actions	RASLog IDs
		Devices	Marginal Threshold	Critical Threshold		
BAD_FAN	Number of fans that are absent or faulty.	Brocade 6505	N/A	N/A	Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2292 to 2295
		All other platforms	N/A	1		
BAD_PWR	Power supply is absent or faulty.	Brocade DCX 8510-4 and Brocade DCX 8510-8	N/A	3	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2288 to 2291
		Brocade X6-4 Director	N/A	1		
		Brocade X6-8 Director	2	3		
		Brocade 6505	N/A	N/A		
		All other platforms	N/A	1		

Monitoring Statistic	Description	Switch Status Threshold Values			Actions	RASLog IDs
		Devices	Marginal Threshold	Critical Threshold		
BAD_TEMP	Sensors indicate temperature is outside of range.	All supported platforms	1	2	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2284 to 2287
DOWN_CORE	Number of core blades that are not functioning.	Only chassis are supported	1	2	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2360 to 2363
ERR_PORTS	Percentage of ports that cannot operate because they are fenced or decommissioned, or unable to operate due to security issues and so on. (A loopback port is considered an error port).	All supported platforms	N/A	5	N/A	MAPS-2312 to 2315
FAN_AIRFLOW_MISMATCH	Air flows of fans are mismatched.	Brocade 6505, Brocade 6510, Brocade 6520, Brocade 7840, Brocade G620, Brocade X6-4, Director, and Brocade X6-8 Director	TRUE	N/A	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS	MAPS-2192 to 2195
FAULTY_BLADE	The number of blades with a faulty status.	Only chassis are supported	1	N/A	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS	MAPS-2304 to 2307

Monitoring Statistic	Description	Switch Status Threshold Values			Actions	RASLog IDs
		Devices	Marginal Threshold	Critical Threshold		
FAULTY_PORTS	Percentage of ports that are faulty.	All supported platforms	N/A	5	Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2296 to 2299
FLASH_USAGE	Percentage of storage space being used.	All supported platforms	90	N/A	RASLog, SNMP, EMAIL, SW_MARGINAL, SW_CRITICAL, FMS	MAPS-2152 to 2155
HA_SYNC	High Availability is not synchronized.	Only chassis are supported	sync=0	N/A	RASLog, SNMP, EMAIL, SW_MARGINAL, SW_CRITICAL, FMS	MAPS-2364 to 2367
MARG_PORTS	Percentage of ports that are marginal.	All supported platforms	N/A	5	Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2308 to 2311
MISSING_SFP	Percentage of ports without SFPs.	All supported platforms	N/A	N/A	N/A	MAPS-2300 to 2303
WWN_DOWN	Number of faulty WWN systems.	Only chassis are supported	N/A	1	Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2368 to 2371

Switch Status Monitoring Thresholds and Actions for the Default Moderate Policy

The following table lists the default *Switch Status Monitoring* thresholds and actions for the default moderate policy:

Monitoring Statistic	Description	Switch Status Threshold Values			Actions	RASLog IDs
		Devices	Marginal Threshold	Critical Threshold		
BAD_FAN	Number of fans that are absent or faulty.	Brocade 6505	N/A	N/A	Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2292 to 2295
		All other platforms	N/A	1		
BAD_PWR	Power supply is absent or faulty.	Brocade DCX 8510-4 and Brocade DCX 8510-8	N/A	3	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2288 to 2291
		Brocade X6-4 Director	N/A	1		
		Brocade X6-8 Director	2	3		

Monitoring Statistic	Description	Switch Status Threshold Values			Actions	RASLog IDs
		Devices	Marginal Threshold	Critical Threshold		
		Brocade 6505	N/A	N/A		
		All other platforms	N/A	1		
BAD_TEMP	Sensors indicate temperature is outside of range.	All supported platforms	1	2	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2284 to 2287
DOWN_CORE	Number of core blades that are not functioning.	Only chassis are supported	1	2	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2360 to 2363
ERR_PORTS	Percentage of ports that cannot operate because they are fenced or decommissioned, or unable to operate due to security issues and so on.	All supported platforms	6	10	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL	MAPS-2312 to 2315
FAN_AIRFLOW_MISMATCH	Air flows of fans are mismatched.	Brocade 6505, Brocade 6510, Brocade 6520, Brocade 7840, Brocade G620, Brocade X6-4 Director, and Brocade X6-8 Director	TRUE	N/A	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS	MAPS-2192 to 2195
FAULTY_BLADE	The number of blades with a faulty status.	Only chassis are supported	1	N/A	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS	MAPS-2304 to 2307

Monitoring Statistic	Description	Switch Status Threshold Values			Actions	RASLog IDs
		Devices	Marginal Threshold	Critical Threshold		
FAULTY_PORTS	Percentage of ports that are faulty.	All supported platforms	6	10	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2296 to 2299
FLASH_USAGE	Percentage of storage space being used.	All supported platforms	90	N/A	RASLog, SNMP, EMAIL, SW_MARGINAL, SW_CRITICAL, FMS	MAPS-2152 to 2155
HA_SYNC	High Availability is not synchronized.	Only chassis are supported	sync=0	N/A	RASLog, SNMP, EMAIL, SW_MARGINAL, SW_CRITICAL, FMS	MAPS-2364 to 2367
MARG_PORTS	Percentage of ports that are marginal.	All supported platforms	6	10	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2308 to 2311
MISSING_SFP	Total number of ports without SFPs.	All supported platforms	N/A	N/A	N/A	MAPS-2300 to 2303
WWN_DOWN	Number of faulty WWN systems.	Only chassis are supported	N/A	1	Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2368 to 2371

Switch Status Monitoring Thresholds and Actions for the Default Conservative Policy

The following table lists the default *Switch Status Monitoring* thresholds and actions for the default conservative policy:

Monitoring Statistic	Description	Switch Status Threshold Values			Actions	RASLog IDs
		Devices	Marginal Threshold	Critical Threshold		
BAD_FAN	Number of fans that are absent or faulty.	Brocade 6505	N/A	N/A	Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2292 to 2295
		All other platforms	N/A	1		
BAD_PWR	Power supply is absent or faulty.	Brocade DCX 8510-4 and Brocade DCX 8510-8	N/A	3	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2288 to 2291
		Brocade X6-4 Director	N/A	1		
		Brocade X6-8Director	2	3		
		Brocade 6505	N/A	N/A		
		All other platforms	N/A	1		
BAD_TEMP	Sensors indicate temperature is outside of range.	All supported platforms	1	2	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2284 to 2287
DOWN_CORE	Number of core blades that are not functioning.	Only chassis are supported	1	2	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2360 to 2363
ERR_PORTS	Percentage of ports that cannot operate because they are fenced or decommissioned, or unable to operate due to security issues and so on.	All supported platforms	11	25	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2312 to 2315

Monitoring Statistic	Description	Switch Status Threshold Values			Actions	RASLog IDs
		Devices	Marginal Threshold	Critical Threshold		
FAN_AIRFLOW_MISMATCH	Air flows of fans are mismatched.	Brocade 6505, Brocade 6510, Brocade 6520, Brocade 7840, Brocade G620, Brocade X6-4 Director, and Brocade X6-8 Director	TRUE	N/A	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS	MAPS-2192 to 2195
FAULTY_BLADE	The number of blades with a faulty status.	Only chassis are supported	1	N/A	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS	MAPS-2304 to 2307
FAULTY_PORTS	Percentage of ports that are faulty.	All supported platforms	11	25	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2296 to 2299
FLASH_USAGE	Percentage of storage space being used.	All supported platforms	90	N/A	RASLog, SNMP, EMAIL, SW_MARGINAL, SW_CRITICAL, FMS	MAPS-2152 to 2155
HA_SYNC	High Availability is not synchronized.	Only chassis are supported	sync=0	N/A	RASLog, SNMP, EMAIL, SW_MARGINAL, SW_CRITICAL, FMS	MAPS-2364 to 2367
MARG_PORTS	Percentage of ports that are marginal.	All supported platforms	11	25	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2308 to 2311
MISSING_SFP	Percentage of ports without SFPs.	All supported platforms	N/A	N/A	N/A	MAPS-2300 to 2303
WWN_DOWN	Number of faulty WWN systems.	Only chassis are supported	N/A	1	Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2368 to 2371

Switch Status Monitoring Thresholds and Actions for the Default Base Policy

The following table lists the default *Switch Status Monitoring* thresholds and actions for the default base policy:

Monitoring Statistic	Description	Switch Status Threshold Values			Actions	RASLog IDs
		Devices	Marginal Threshold	Critical Threshold		
BAD_FAN	Number of fans that are absent or faulty.	Brocade 6505	N/A	N/A	Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2292 to 2295
		All other platforms	N/A	1		
BAD_PWR	Power supply is absent or faulty.	Brocade DCX 8510-4 and Brocade DCX 8510-8	N/A	3	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2288 to 2291
		Brocade X6-4 Director	N/A	1		
		Brocade X6-8Director	2	3		
		Brocade 6505	N/A	N/A		
		All other platforms	N/A	1		
BAD_TEMP	Sensors indicate temperature is outside of range.	All supported platforms	1	2	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2284 to 2287
DOWN_CORE	Number of core blades that are not functioning.	Only chassis are supported	1	2	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2360 to 2363

Monitoring Statistic	Description	Switch Status Threshold Values			Actions	RASLog IDs
		Devices	Marginal Threshold	Critical Threshold		
FAN_AIRFLOW_MISMATCH	Air flows of fans are mismatched.	Brocade 6505, Brocade 6510, Brocade 6520, Brocade 7840, Brocade G620, Brocade X6-4 Director, and Brocade X6-8 Director	TRUE	N/A	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS	MAPS-2192 to 2195
FAULTY_BLADE	The number of blades with a faulty status.	Only chassis are supported	1	N/A	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS	MAPS-2304 to 2307
FLASH_USAGE	Percentage of storage space being used.	All supported platforms	90	N/A	RASLog, SNMP, EMAIL, SW_MARGINAL, SW_CRITICAL, FMS	MAPS-2152 to 2155
HA_SYNC	High Availability is not synchronized.	Only chassis are supported	sync=0	N/A	RASLog, SNMP, EMAIL, SW_MARGINAL, SW_CRITICAL, FMS	MAPS-2364 to 2367
MARG_PORTS	Percentage of ports that are marginal.	All supported platforms	11	25	Marginal: RASLog, SW_MARGINAL, SNMP, EMAIL, FMS Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2308 to 2311
MISSING_SFP	Percentage of ports without SFPs.	All supported platforms	N/A	N/A	N/A	MAPS-2300 to 2303
WWN_DOWN	Number of faulty WWN systems.	Only chassis are supported	N/A	1	Critical: RASLog, SW_CRITICAL, SNMP, EMAIL, FMS	MAPS-2368 to 2371

Traffic Performance Thresholds

Traffic performance monitors are used to monitor imported Flow Vision flows. The default policies do not include preset thresholds.

Monitor	Description	Group Type	Supported Timebases				Actions	RASLog IDs
			Minute	Hour	Day	None		
RX_THPUT	Receive throughput	Flow	Yes	Yes	No	No	RASLog, SNMP, EMAIL, NONE, FMS	MAPS-2388 to 2391
TX_FCNT	Transmit frame count	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, NONE, FMS	MAPS-2376 to 2379
RX_FCNT	Receive frame count	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, NONE, FMS	MAPS-2380 to 2383
TX_THPUT	Transmit throughput	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, NONE, FMS	MAPS-2384 to 2387
IO_RD	I/O read command count (not supported for VM Insight)	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, NONE, FMS	MAPS-2392 to 2395
IO_WR	I/O write command count (not supported for VM Insight)	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, NONE, FMS	MAPS-2396 to 2399
IO_RD_BYTES	I/O read data (not supported for VM Insight)	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, NONE, FMS	MAPS-2400 to 2403
IO_WR_BYTES	I/O write data (not supported for VM Insight)	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, NONE, FMS	MAPS-2404 to 2407
RD_STATUS_TIME_LT_8K	Time for read I/O request less than 8K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2724 to 2727
RD_STATUS_TIME_8_64K	Time for read I/O request greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2728 to 2731
RD_STATUS_TIME_64_512K	Time for read I/O request greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2732 to 2735
RD_STATUS_TIME_GE_512K	Time for read I/O request greater than or equal to 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2736 to 2739

Monitor	Description	Group Type	Supported Timebases				Actions	RASLog IDs
			Minute	Hour	Day	None		
WR_STATUS_TIME_LT_8K	Time for write I/O request less than 8K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2740 to 2743
WR_STATUS_TIME_8_64K	Time for write I/O request greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2744 to 2747
WR_STATUS_TIME_64_512K	Time for write I/O request greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2748 to 2751
WR_STATUS_TIME_GE_512K	Time for write I/O request greater than or equal to 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2752 to 2755
RD_1stDATA_TIME_LT_8K	Time for first data read less than 8K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2756 to 2759
RD_1stDATA_TIME_8_64K	Time for first data read greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2760 to 2763
RD_1stDATA_TIME_64_512K	Time for first data read greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2764 to 2767
RD_1stDATA_TIME_GE_512K	Time for first data read greater than or equal to 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2768 to 2771
WR_1stXFER_RDY_LT_8K	First data transfer in the ready state less than 8K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2772 to 2775
WR_1stXFER_RDY_8_64K	First data transfer in ready state greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2776 to 2779
WR_1stXFER_RDY_64_512K	First data transfer in ready state greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2780 to 2783
WR_1stXFER_RDY_GE_512K	First data transfer in ready state greater than or equal to 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2784 to 2787
RD_PENDING_IO_LT_8K	Pending read I/O requests less than 8K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2788 to 2791

Monitor	Description	Group Type	Supported Timebases				Actions	RASLog IDs
			Minute	Hour	Day	None		
RD_PENDING_IO_8_64K	Pending read I/O requests greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2792 to 2795
RD_PENDING_IO_64_512K	Pending read I/O requests greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2796 to 2799
RD_PENDING_IO_GE_512K	Pending read I/O requests greater than or equal to 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2800 to 2803
WR_PENDING_IO_LT_8K	Pending write I/O requests less than 8K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2804 to 2807
WR_PENDING_IO_8_64K	Pending write I/O requests greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2808 to 2811
WR_PENDING_IO_64_512K	Pending write I/O requests greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2812 to 2815
WR_PENDING_IO_GE_512K	Pending write I/O requests greater than or equal to 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2816 to 2819
RD_IO_RATE_LT_8K	Read I/O bytes less than 8K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2820 to 2823
RD_IO_RATE_8_64K	Read I/O bytes greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2824 to 2827
RD_IO_RATE_64_512K	Read I/O bytes greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2828 to 2831
RD_IO_RATE_GE_512K	Read I/O bytes greater than or equal to 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2832 to 2835
WR_IO_RATE_LT_8K	Written I/O bytes less than 8K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2836 to 2839

Monitor	Description	Group Type	Supported Timebases				Actions	RASLog IDs
			Minute	Hour	Day	None		
WR_IO_RATE_8_64K	Written I/O bytes greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2840 to 2843
WR_IO_RATE_64_512K	Written I/O bytes greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2844 to 2847
WR_IO_RATE_GE_512K	Written I/O bytes greater than or equal to 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2848 to 2851
RD_IOPS_LT_8K	Read I/O count less than 8K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2852 to 2855
RD_IOPS_8_64K	Read I/O count greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2856 to 2859
RD_IOPS_64_512K	Read I/O count greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2860 to 2863
RD_IOPS_GE_512K	Read I/O count greater than or equal to 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2864 to 2867
WR_IOPS_LT_8K	Written I/O count less than 8K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2868 to 2871
WR_IOPS_8_64K	Written I/O count greater than or equal to 8K but less than 64K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2872 to 2875
WR_IOPS_64_512K	Written I/O count greater than or equal to 64K but less than 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2876 to 2879
WR_IOPS_GE_512K	Written I/O count greater than or equal to 512K	Flow	Yes	Yes	Yes	No	RASLog, SNMP, EMAIL, FMS	MAPS-2880 to 2883

Revision History

FOS-82x-MAPS-UG106; October 14, 2022

Updated the Minimum Quiet Time values of the following monitoring systems in the Quieting a Rule section:

- CURRENT
- PWR_HRS
- RXP
- SFP_TEMP
- TXP
- VOLTAGE

FOS-82x-MAPS-UG105; February 05, 2021

All references to the Fabric OS 8.2.x version have been standardized.

FOS-82x-MAPS-UG104; March 26, 2020

- Updated the default rules for UCS Uplink Distribution Monitoring.
- Updated the ALL_OTHER_SFP value in the tables in Predefined Groups.

FOS-821-MAPS-UG103; June 07, 2019

- Updated the SFP_MARGINAL action in the output of the `mapsconfig --show` and `mapsdb --show` commands.

FOS-821-MAPS-UG102; March 11, 2019

- Updated the marginal threshold, critical threshold, and actions values to N/A in the tables in Switch Status Policy Monitoring.
- Updated all monitoring system records.

FOS-821-MAPS-UG101; December 04, 2018

- Revised the publication number.
- Updated the document for stylistic changes.

FOS-821-MAPS-UG100; October 17, 2018

- Updated MAPS rules and groups altered in this release.
- Updated Firmware Downgrade Considerations for MAPS.
- Added the `DEV_LOGIN_DIST` monitoring statistic in Fabric Performance Impact.
- Updated the contents of Slow-Drain Device Quarantining.
- Added UCS Fabric Interconnect Rebalance.
- Added the minimum quiet time value for the IP_EXTN_FLOW monitoring system.
- Updated all threshold values tables with the latest actions for each monitoring system in MAPS Threshold Values.

Documentation Legal Notice

This notice provides copyright and trademark information as well as legal disclaimers.

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

The product described by this document may contain open source software covered by the GNU General Public License or other open source license agreements. To find out which open source software is included in Brocade products, to view the licensing terms applicable to the open source software, and to obtain a copy of the programming source code, please download the open source disclosure documents in the Broadcom Customer Support Portal (CSP). If you do not have a CSP account or are unable to log in, please contact your support provider for this information.

BROADCOM®