53-1003941-01
24 September 2015

Monitoring and Alerting Policy Suite

Administrator's Guide

Supporting Fabric OS v7.4.1

BROCADE

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

The authors and Brocade Communications Systems, Inc. assume no liability or responsibility to any person or entity with respect to the accuracy of this document or any loss, cost, liability, or damages arising from the information contained herein or the computer programs that accompany it.

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

Contents

Preface

Document conventions

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

Text formatting conventions

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

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

Command syntax conventions

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

Convention	Description
bold text	Identifies command names, keywords, and command options.
italic text	Identifies a variable.
value	In Fibre Channel products, a fixed value provided as input to a command option is printed in plain text, for example, **--show** WWN.

Convention	Description
[ ]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ **x** | **y** | **z** }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options. In Fibre Channel products, square brackets may be used instead for this purpose.
x | **y**	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, for example, passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, *member*[*member*...].
\	Indicates a "soft" line break in command examples. If a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Notes, cautions, and warnings

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

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

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

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

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

Brocade resources

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

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

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

Release notes are available on MyBrocade under Product Downloads.

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

Contacting Brocade Technical Support

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

Brocade customers

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

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

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

Brocade OEM customers

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

- OEM/Solution Providers are trained and certified by Brocade to support Brocade® products.
- Brocade provides backline support for issues that cannot be resolved by the OEM/Solution Provider.

- Brocade Supplemental Support augments your existing OEM support contract, providing direct access to Brocade expertise. For more information, contact Brocade or your OEM.
- For questions regarding service levels and response times, contact your OEM/Solution Provider.

Document feedback

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

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

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

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

About This Document

Supported hardware and software

In those instances in which procedures or parts of procedures documented here apply to some switches but not to others, this list identifies exactly which switches are supported and which are not.

Although many different software and hardware configurations are tested and supported by Brocade Communications Systems, Inc. for Fabric OS 7.4.1, documenting all possible configurations and scenarios is beyond the scope of this document.

Fabric OS v7.4.1 provides features that interoperate with the Brocade Analytics Monitoring Platform running Fabric OS v7.4.0_amp. For more information about Brocade Analytics Monitoring Platform and Fabric OS v7.4.0, you can find content on MyBrocade.com.

The following hardware platforms are supported by this release of Fabric OS.

Brocade Gen 4 platform (8-Gpbs) fixed-port switches

- Brocade 300 switch
- Brocade 5100 switch
- Brocade 5300 switch
- Brocade 5410 blade server SAN I/O module
- Brocade 5424 blade server SAN I/O module
- Brocade 5430 blade server SAN I/O module
- Brocade 5431 blade server SAN I/O module
- Brocade 5432 blade server SAN I/O module
- Brocade 5450 blade server SAN I/O module
- Brocade 5460 blade server SAN I/O module
- Brocade 5470 blade server SAN I/O module
- Brocade 5480 blade server SAN I/O module
- Brocade NC-5480 blade server SAN I/O module
- Brocade 7800 extension switch
- Brocade VA-40FC switch
- Brocade Encryption Switch

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

- Brocade 6543 blade server SAN I/O module
- Brocade 6505 switch
- Brocade M6505 Blade Server SAN I/O Module
- Brocade 6510 switch

- Brocade 6520 switch
- Brocade 6547 blade server SAN I/O module
- Brocade 6548 blade server SAN I/O module
- Brocade 7840 extension switch

Brocade Gen 4 platform (8-Gpbs) DCX Backbone family

- Brocade DCX
- Brocade DCX-4S

Brocade Gen 5 platform (16-Gbps) DCX Backbone family

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

What's new in this document

This document includes new and modified information for the Fabric OS 7.4.1 release of MAPS.

Changes made for this release

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

- Supported hardware and software on page 11 was updated to indicate that the Brocade 6453 blade server SAN I/O module is supported in this release.
- Downgrade considerations were updated for this release.
- The description of SFP marginal on page 64 in the MAPS rules overview on page 59 has been updated to include examples of how the **sfpshow -health** command is affected when the **SFP_MARGINAL** action has been triggered.

Glossary

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

TABLE 1 MAPS-related terminology

Term	Description
Action	Activity performed by MAPS if a condition defined in a rule evaluates to true.
ASIC	Application-Specific Integrated Circuit; the chip within a switch or blade that controls its operation, including packet switching and routing.
Back-end port	Port that connects a core switching blade to an edge switching blade (and vice versa).
Condition	Combination of count, timebase, threshold value, and the logical operation (examples: =, >, <, >=) that needs to be evaluated; for example: CRC/MIN > 10.
Congestion	Circumstance in which the offered load exceeds the physical capacity of the circuit but does not exceed the rate at which the other end of the circuit can continuously accept traffic.

TABLE 1 MAPS-related terminology (Continued)

Term	Description
CRC	Cyclic redundancy check; an error-detecting code used to detect accidental changes to raw data.
Flow	A set of Fibre Channel frames that share similar traits.
Flow Monitor	A Flow Vision application that can monitor all the traffic passing through E_Ports, EX_Ports, F_Ports, and XISL_Ports. Refer to the *Flow Vision Administrator's Guide* for details.
Front-end port	Port that connects a switch to another device, such as a storage unit.
FRU	Field-Replaceable Unit; Power supply or other physical element of a chassis-based device that can be replaced by an end-user.
HBA	Host Bus Adapter; Fibre Channel interface that allows a host system to connect to other network or storage devices.
ISL	Inter-Switch Link; a link joining two Fibre Channel switches using E_Ports.
ITW	Invalid transmission word; if a data word does not transmit successfully, encoding errors may result.
Logical group	A group of ports that behave in a similar manner. MAPS uses such a group to monitor the ports using a single set of rules and thresholds.
NPIV	N_Port ID Virtualization; a feature permitting multiple Fibre Channel node port (N_Port) IDs to share a single physical N_Port.
Policy	Set of rules which are activated at the same time.
Rule	Setting that associates a condition with actions that need to be triggered when the specified condition is evaluated to true.
Root flow	Instance of a static flow used to create learned flows.
QoS	Quality of Service; mechanism for assigning priorities to data flows. See Traffic Management.
RSCN	Registered State Change Notification; a Fibre Channel fabric notification sent to all specified switches and nodes identifying significant fabric changes.
SAN	Storage Area Network; a dedicated network that provides access to consolidated, block level data storage.
Slow-drain(ing) device	Device that does not process frames at the rate generated by the source.
Static flow	Flow created when learning is not used and all the parameters required are contained in the flow definition.
Sub-flow	System auto-created flow based on a root flow. A root flow can have more than one sub-flow.
Threshold	Value used in evaluating a condition.
Timebase	The time period across which the change in a counter is to be monitored.
Traffic isolation	Fabric OS feature that reserves ISLs in a fabric for traffic between a particular set of ports. This is a best effort service and the ISLs are not guaranteed to be reserved exclusively for the specified traffic
Traffic management	Fabric OS mechanism that assigns high, medium, or low priority to data flows.

TABLE 1 MAPS-related terminology (Continued)

Term	Description
VC	Virtual Circuit; mechanism for transporting data over a packet-switched network so that it appears as a dedicated physical layer link between the source and destination.
Virtual channel	Synonym for Virtual Circuit.
XISL	An eXtended ISL; a type of ISL connection that carries traffic for multiple logical fabrics on the same physical link while maintaining fabric isolation.

Monitoring and Alerting Policy Suite Overview

MAPS overview

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

In Fabric OS 7.4.0, MAPS is no longer optional; it replaces Fabric Watch, and provides a set of monitors that work even if the MAPS license is not activated. Refer to Feature monitors not requiring a Fabric Vision license on page 22 for information on these monitors. Refer to Differences between Fabric Watch and MAPS configurations on page 21 and Fabric Watch to MAPS migration on page 20 for information on the transition from Fabric Watch. MAPS is controlled though the Fabric Vision license.

MAPS is implicitly activated when you upgrade to Fabric OS 7.4.0; however, doing so without a license provides reduced functionality. Refer to MAPS activation on page 16 and Firmware upgrade considerations on page 18 for specific details on activating MAPS. When you upgrade to Fabric OS 7.4.0, MAPS starts monitoring the switch as soon as the switch is active.

When the Fabric Vision license is activated, MAPS provides you with a set of predefined monitoring policies that allow you to use it immediately. These are described in Predefined policies on page 50. In addition, MAPS provides you with the ability to create customized monitoring rules and policies, allowing you to define specific groups of ports or other elements that will share a common threshold. MAPS lets you define how often to check each switch and fabric measure and specify notification thresholds. This allows you to configure MAPS to provide notifications before problems arise, for example, when network traffic through a port is approaching a bandwidth limit. Whenever fabric measures exceed these thresholds, MAPS can automatically notify you through e-mail messages, SNMP traps, log entries, or combinations. For information on using these features, refer to MAPS groups overview on page 41, MAPS rules overview on page 59, and MAPS policies overview on page 49.

MAPS provides a dashboard that allows you to quickly view what is happening on the switch (for example, the kinds of errors, the error count, and so on). Refer to the MAPS dashboard overview on page 89 for more information.

MAPS license requirements

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

In order to provide full functionality, MAPS requires an active and valid Fabric Vision license. Alternatively if you are upgrading and already have a license for Fabric Watch plus a license for Advanced Performance Monitoring, you will automatically get MAPS functionality without having to obtain a separate license. Refer to the *Fabric OS Software Licensing Guide* for more information about licensing and how to obtain the necessary license keys. If you only have one of these licenses, you will need to acquire the other in order to use all the MAPS features. Without the correct licenses from the previous releases or the current release, neither converting Fabric Watch custom policies to MAPS nor migrating to MAPS is allowed.

MAPS activation

The Fabric Vision license must be activated (enabled) in Fabric OS 7.4.0 for you to use the full set of MAPS options or any legacy Fabric Watch thresholds. In addition, to reuse the Fabric Watch thresholds, you must enable MAPS in Fabric OS 7.3.x before upgrading to Fabric OS 7.4.0.

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

- Activating MAPS is a chassis-specific process, and you can activate only one chassis at a time.
- On any given chassis there can be multiple logical switches.
- Activating MAPS enables it for all logical switches in the chassis, but each logical switch can have its own MAPS configuration.

Cautions

MAPS activation is a non-reversible process. After you activate MAPS for Fabric OS 7.4.0, it is enabled for any version of Fabric OS from 7.2.0 or later.

Fabric Watch is not part of Fabric OS 7.4.0, and only becomes active if you downgrade to a version of Fabric OS prior to 7.2.0. This will enable Fabric Watch with its last configured settings. These settings will be the defaults if no custom settings are available.

Refer to Firmware upgrade and downgrade considerations for MAPS on page 18 for additional firmware upgrade and downgrade considerations.

MAPS configuration files

The MAPS user configuration is persistent across reboot and can be uploaded or downloaded. A configuration upload or download affects only the user-created configuration files. You cannot upload or download the default MAPS configuration file. Only one user configuration file can exist for each logical switch.

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

Deleting a user-created MAPS configuration file

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

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

MAPS interaction with other Fabric OS features

MAPS interacts in different ways with different Fabric OS features, including Admin Domains, High Availability, and Virtual Fabrics.

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

TABLE 2 Interactions between Fabric OS features and MAPS

Feature	MAPS interaction
Admin Domains	MAPS is supported on switches that have Admin Domains. There can only be one MAPS configuration that is common to all the Admin Domains on the chassis. Users with Administrator privileges can modify the MAPS configuration from any Admin Domain. **ATTENTION** If MAPS is enabled, do not download configuration files that have Admin Domains defined in them, as this may cause unpredictable behavior.
High Availability	MAPS configuration settings are maintained across a High Availability (HA) failover or HA reboot; however, MAPS will restart monitoring after a HA failover or HA reboot and the cached MAPS statistics are not retained.
Virtual Fabrics	When using Virtual Fabrics, different logical switches in a chassis can have different MAPS configurations.

Restrictions on MAPS monitoring

The following restrictions apply globally to MAPS monitoring:

- Small form-factor pluggable (SFP) transceivers on simulated mode (SIM) ports cannot be monitored using MAPS.
- If a SCN pertaining to the FRU (such as PS_FAULTY or FAN_FAULTY) occurs before the dashboard starts monitoring, then it might not be shown in the dashboard.

Refer to Restrictions on Flow Vision flow monitoring for additional restrictions on monitoring Flow Vision flows.

Firmware upgrade and downgrade considerations for MAPS

The following firmware upgrade and downgrade considerations apply to the Monitoring and Alerting Policy Suite (MAPS) in this version of Fabric OS.

Firmware upgrade considerations

When upgrading to Fabric OS 7.4.1, there are four main scenarios:

- Upgrading if Fabric Watch is installed but not enabled (activated) before the upgrade to Fabric OS 7.4.1.
- Upgrading if Fabric Watch is installed and enabled before the upgrade to Fabric OS 7.4.1.
- Upgrading if MAPS was installed before the upgrade to Fabric OS 7.4.1, but not enabled.
- Upgrading if MAPS was installed before the upgrade to Fabric OS 7.4.1, and is enabled.

Each of these scenarios is discussed in the following sections. In addition, MAPS Fabric Performance Impact monitoring and the legacy bottleneck monitoring feature are mutually exclusive. If the legacy bottleneck monitoring feature was enabled before the upgrade, MAPS will not monitor fabric performance.

Upgrading if Fabric Watch is installed but not enabled

If Fabric Watch is installed on the switch but is not enabled at the time of the upgrade, then MAPS is installed and will monitor the fabric using the policy named “dflt_base_policy” which has the basic monitoring capabilities and will not permit you to migrate any custom Fabric Watch thresholds as part of the upgrade.

NOTE
If the Fabric Vision license is not active before you upgrade, then you cannot enable any policies and MAPS will monitor the fabric using only the basic monitoring capabilities until the Fabric Vision license is installed.

Upgrading if Fabric Watch was installed and is enabled

If Fabric Watch is installed on the switch and is enabled, then MAPS is installed with the basic monitoring capabilities, you can migrate any existing Fabric Watch thresholds as part of the change to MAPS.

Upgrading if MAPS was installed, but not enabled

If MAPS is installed on the switch but is not enabled at the time the firmware is upgraded, after the upgrade MAPS will be enabled either with the active policy named in the configuration, or if there is no active policy named, then using the default conservative policy “dft_conservative_policy”. If MAPS was licensed but not enabled, you can re-enable MAPS with one of the default policies or a converted Fabric Watch policy (if Fabric Watch was licensed before upgrade) based on the Fabric Watch custom thresholds that have been converted to MAPS policies before the upgrade.

If MAPS gets enabled implicitly (without the Fabric Vision license), then 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 was installed and is enabled

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

Firmware downgrade considerations

There are three primary downgrade scenarios from Fabric OS 7.4.1 for MAPS.

- If you are downgrading from Fabric OS 7.4.0 and do not have an active Fabric Vision license, you will still have the options to use Fabric Watch or activate a license for the previous version of MAPS (if you have one) after the downgrade.
- If you are downgrading from Fabric OS 7.4.0 or 7.4.1 and have an active Fabric Vision license, when you downgrade to Fabric OS 7.3.x or 7.2.x, Fabric Watch will not be available, but you can still use MAPS at the earlier functionality level. In this case, the Fabric Watch customer thresholds are available.
- If you are downgrading from Fabric OS 7.4.0 or 7.4.1 and have an active Fabric Vision license, when you downgrade to Fabric OS 7.1.x or earlier, Fabric Watch will be available (it will require you to activate the license for it), but MAPS will not be available.

When downgrading from Fabric OS 7.4.x to a previous version of the operating system, the following MAPS-related behaviors should be expected:

- If Fabric Watch was licensed before upgrading to Fabric OS 7.4.x, downgrading to Fabric OS 7.1.x or earlier recovers all the Fabric Watch custom configurations. If Fabric Watch was not licensed before the upgrade, then downgrading to 7.1.x or earlier means that only the non-licensed Fabric Watch monitors are restored.
- When an active Control Processor (CP) is running Fabric OS 7.2.x or later with MAPS disabled, and the standby CP has an earlier version of Fabric OS, High Availability will be synchronized, but MAPS will not be allowed to be enabled until the firmware on the standby device is changed to the matching version.
- When an active CP is running Fabric OS 7.4.x, 7.3.x, or 7.2.x and MAPS is enabled, but the standby CP is running Fabric OS 7.1.0 or earlier, then High Availability will not be synchronized until the standby CP is upgraded to a version of Fabric OS that matches the one on the active CP.
- On Backbone platforms that have a single CP, there is no change in behavior when downgrading to an earlier version of Fabric OS.
- Downgrading to previous versions of Fabric OS will fail if some features are not supported in the earlier firmware, and their loss could impact MAPS functionality. In this case, MAPS provides instructions on how to disable these features before firmware downgrade. An example of this is if either MAPS actions or rules include Fabric Performance Impact monitoring or port decommissioning.

  The following is a list of the rules that will be affected by a downgrade from Fabric OS 7.4.x:

  - Any rules triggered by DEV_LATENCY_IMPACT with the SDDQ action.
  - Any rule that has the action SDDQ configured or if this action is enabled globally, downgrading to any version of Fabric OS that does not support the SDDQ action is blocked.
  - Any rule triggered by DEV_LATENCY_IMPACT with the FPI state IO_LATENCY_CLEAR.
  - Any rule that has the action FMS.
  - Any rule that has the action TOGGLE.

  - Any rule that has the quiet time (QT) option set.
  - Any rules that operate on the predefined groups ALL_CIRCUIT_HIGH_QOS, ALL_CIRCUIT_MED_QOS, ALL_CIRCUIT_LOW_QOS, and ALL_CIRCUIT_F_QOS.
- Downgrading to versions of Fabric OS prior to version 7.4.0 will trigger a warning message if any feature is not supported in the earlier firmware and keeping the feature configuration has no impact. In this case, the downgrade will not be blocked; however, MAPS will display a warning message similar to the following:

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

  Examples of this condition include MAPS having any user-created rules pertaining to monitoring the following: IO_LATENCY_CLEAR, ALL_CIRCUIT_HIGH_QOS, ALL_CIRCUIT_MED_QOS, ALL_CIRCUIT_LOW_QOS, ALL_CIRCUIT_F_QOS, DEV_NPIV_LOGINS, or QT.
- Downgrading to versions of Fabric OS prior to version 7.3.0 is not allowed if the MAPS Fabric Performance Impact monitoring feature is enabled. You must disable FPI using the **mapsconfig –disableFPImon** command before starting the firmware downgrade.
- Downgrading is not allowed if automatic VC quarantining has been enabled and there are slow-drain isolations already performed. The isolated slow-drain flows need to be restored to their original virtual circuits before proceeding with the downgrade because the commands to restore them will not be available in the downgraded firmware version.
- Downgrading is not allowed if there is any quarantined port in the logical group ALL_QUARANTINED_PORTS. Before you can downgrade the switch firmware, you must clear the ports from the quarantined state using the **sddquarantine --clear** *slot/port* or **all** command.

NOTE
There may be other factors affecting downgrading the firmware version; these are only the ones that MAPS affects or is affected by.

Fabric Watch to MAPS migration

Fabric Watch is no longer present in Fabric OS 7.4.0; if you want that monitoring functionality, you must migrate from Fabric Watch to MAPS.

MAPS provides the option for a seamless migration of any existing customized Fabric Watch thresholds to MAPS when you activate it, thus allowing you to take advantage of your existing configurations. MAPS provides additional advanced monitoring, such as monitoring for the same error counters across different time periods, or having more than two thresholds for any error counters. MAPS also provides support for you to monitor the statistics provided by the Flow Monitor feature of Flow Vision. Refer to Differences between Fabric Watch and MAPS configurations on page 21.

When you start MAPS for the first time, you can have it automatically convert any existing Fabric Watch configurations to ones that are compatible with MAPS so you do not need to recreate all of the thresholds and rules. However, you must do this before you activate MAPS. It is not possible to do an automated conversion after the initial activation; to enable them later you will need to create the thresholds and rules individually. Refer to Upgrading firmware on Fabric Watch-licensed switches on page 21 for details.

Refer to Enabling MAPS using Fabric Watch rules on page 25 for instructions on starting MAPS using your existing Fabric Watch rules.

When you upgrade a switch without a Fabric Watch license, the monitors described in Feature monitors not requiring a Fabric Vision license on page 22 are activated. Upgrading a switch without a Fabric Watch license will not allow you to convert any Fabric Watch configurations on that switch to MAPS custom policies.

NOTE
With the removal of Fabric Watch, a number of commands have been removed from Fabric OS or have been modified. Refer to the *Fabric OS Upgrade Guide* and *Fabric OS Command Reference* for information on the specific commands that have been removed or changed.

Upgrading firmware on Fabric Watch-licensed switches

On switches that have Fabric Watch licensed and are explicitly using Fabric Watch monitoring, upgrading to Fabric OS 7.4.0 is possible with or without migrating to MAPS. However, you must migrate to MAPS prior to upgrading in order to use any custom Fabric Watch threshold configurations, and to continue to be able use the monitoring service non-disruptively after the upgrade.

ATTENTION
Upgrading without converting Fabric Watch configurations to MAPS custom policies will result in the loss of any custom threshold configurations.

To migrate from Fabric Watch thresholds to MAPS policies, complete the following steps before upgrading to Fabric OS 7.4.0.

NOTE
Without installing the Fabric Vision license, neither converting to MAPS custom policies nor migrating to MAPS is allowed.

1. Enter **licenseshow** to check if the Fabric Vision license or Fabric Watch plus Advanced Performance Management (APM) licenses are available.
   - If only the Fabric Watch license is available, obtain and install an APM license.
   - If only APM license available, obtain and install a Fabric Watch license.
   - If neither is available, obtain and install a Fabric Vision license.
2. Convert Fabric Watch configurations to MAPS custom policies using the **mapsconfig --fwconvert** command.
3. Enable MAPS with one of the custom policies or default policies available using the **mapsconfig --enablemaps -policy** *policy_name* command.

The following example displays the available licences (in this case, there is a Fabric Vision license installed) and enables the “my_custom_FW_threshold” policy.

```
switch:admin> licenseshow
pSD43LDSQpFYB9Qg44ttRfN3JGXrtyYAB89Gp:
    Fabric Vision license
switch:admin> mapsconfig --enablemaps -policy my_custom_FW_threshold
```

Differences between Fabric Watch and MAPS configurations

The monitoring and alerting configurations available in the MAPS are not as complex as those available in Fabric Watch; as a consequence MAPS lacks some of the functionality available in Fabric Watch.

The following table shows the differences between Fabric Watch and MAPS configurations and functionality.

TABLE 3 Comparison of Fabric Watch and MAPS configurations and functionality

Configuration	Fabric Watch behavior	MAPS behavior
End-to-End monitoring (Performance Monitor class)	Supported	Supported through flows.
Frame monitoring (Performance Monitor class)	Supported	Supported through flows.
RX, TX monitoring	Occurs at the individual physical port level.	Occurs at the trunk or port level as applicable.
Pause/Continue behavior	Occurs at the element or counter level. For example, monitoring can be paused for CRC on one port and for ITW on another port.	Occurs at the element level. Monitoring can be paused on a specific port, but not for a specific counter on that port.
CPU/Memory polling interval	Can configure the polling interval as well as the repeat count.	This configuration can be migrated from Fabric Watch, but cannot be changed.
E-mail notification configuration	Different e-mail addresses can be configured for different classes.	E-mail configuration supported globally.
Temperature sensor monitoring	Can monitor temperature values.	Can monitor only if the sensor state is in or out of range.

Feature monitors not requiring a Fabric Vision license

In releases of Fabric OS prior to 7.4.0, Fabric Watch monitored some Fabric OS features even if the Fabric Watch license was not active. In Fabric OS 7.4.0 and later, many of these features are monitored by MAPS even if the Fabric Vision license is not active.

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

- Switch status policies
- Switch resource changes (FRU state, Flash memory space, Temperature, CPU usage, Memory usage, and Ethernet management port)

For more information, refer to MAPS commands that do not need a Fabric Vision license on page 22.

MAPS commands that do not need a Fabric Vision license

The following table lists the MAPS commands that are accessible without an installed Fabric Vision license. This provides similar functionality to that available as an unlicensed Fabric Watch feature in versions of Fabric OS earlier than 7.4.0.

TABLE 4 MAPS commands accessible without a Fabric Vision license

Command	Effect	Description
logicalgroup --show	Displays unlicensed feature groups details.	Viewing group information on page 41
mapsdb --show	Displays the MAPS dashboard.	MAPS dashboard sections on page 89
mapspolicy --show –summary	Displays a summary of all the policies on the switch.	Viewing policy information on page 51
mapspolicy --show *policy_name*	Displays the rules for the specified MAPS policy on the switch.	Viewing policy information on page 51
mapsrule --show *rule_name*	Displays the details of the specified MAPS rule on the switch.	Viewing MAPS rules on page 66
mapsrule --show -all	Displays all the MAPS rules on the switch.	Viewing MAPS rules on page 66
mapssam --show flash	Displays the flash memory usage as a percentage.	MAPS Service Availability Module on page 125
mapssam --show cpu	Displays the CPU usage as a percentage.	MAPS Service Availability Module on page 125
mapssam --show memory	Displays the general RAM memory usage as a percentage, along with total, used, and free memory values.	MAPS Service Availability Module on page 125
mapshelp	Displays list of MAPS commands.	

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

MAPS Setup and Operation

Initial MAPS setup

The Brocade Monitoring and Alerting Policy Suite (MAPS) is installed by default, but enabling more than basic functionality requires that you activate the license.

If you want to use your existing Fabric Watch threshold rules in MAPS, you must convert them before installing Fabric OS 7.4.0. Once you have done this, you can enable and configure MAPS to use the converted rules. If you have already done this as part of an upgrade to Fabric OS 7.3.x, you do not need to convert them again.

Enabling MAPS using Fabric Watch rules

If you are already using Fabric Watch and would like MAPS to use the same thresholds, you must convert the Fabric Watch policies into MAPS policies before you install Fabric OS 7.4.0. MAPS will then automatically use the policy named “fw_active_policy” to provide the same monitoring functionality as you had using Fabric Watch.

To monitor a switch in this manner, complete the following steps.

ATTENTION
To retain the Fabric Watch policies and thresholds, you must do step 1 *before* installing Fabric OS 7.4.0.

1. Migrate from Fabric Watch by entering **mapsconfig --fwconvert** to import the Fabric Watch rules.
2. Enable MAPS by entering **mapsconfig --enablemaps -policy fw_active_policy**.

   Upon successful completion of this command, the following happens:

   - The Fabric Watch configurations are converted to MAPS policies.
   - Fabric Watch monitoring and commands are disabled.
   - MAPS commands are enabled.
   - The MAPS “fw_active_policy” policy is enabled.

   After you convert the rules and activate MAPS, you can enable the MAPS “fw_active_policy” policy, and then monitor your switch while fine-tuning the policy to fit your environment. When you are satisfied with the configuration settings, you can specify the actions you want to occur when thresholds are crossed.
3. Set global actions on the switch to “none” by entering **mapsconfig --actions none**.

Setting the global actions to “none” allows you to test the configured thresholds before enabling the actions.

4. Monitor the switch by entering **mapsdb --show** or **mapsdb --show all**.
5. Fine-tune the rules used by the policy as necessary.
6. Set global actions on the switch to the allowed actions by using **mapsconfig --actions** and specifying all of the actions that you want to allow on the switch.

The following example enables MAPS, loads the policy “dflt_conservative_policy”, sets the actions to “none”, and then sets approved actions.

```
switch:admin> mapsconfig --fwconvert
switch:admin> mapsconfig --enablemaps -policy dflt_conservative_policy

WARNING:
This command enables MAPS and replaces all Fabric Watch configurations and
monitoring. Once MAPS is enabled, the Fabric Watch configuration can't be converted
to MAPS. If you wish to convert your Fabric Watch configuration into MAPS policies,
select NO to this prompt and first issue the "mapsconfig --fwconvert" command. Once
the Fabric Watch configuration is concerted into MAPS policies, you may reissue the
"mapsconfig --enablemaps" command to continue this process.  If you do not use
Fabric Watch or need the configuration, then select YES to enable MAPS now.
Do you want to continue? (yes, y, no, n): [no] yes
...
MAPS is enabled.

switch:admin> mapsconfig --actions none
switch:admin> mapsconfig --actions raslog,fence,snmp,email,sw_marginal
```

Enabling MAPS without using Fabric Watch rules

If you are not already using Fabric Watch, or do not wish to continue using the Fabric Watch policies, you can quickly start monitoring your switch using MAPS with one of the predefined policies delivered with MAPS.

ATTENTION
If you follow these instructions, the Fabric Watch configurations are not converted to MAPS policies as part of the migration, Fabric Watch commands are disabled, and MAPS commands are enabled.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **mapspolicy --enablemaps -policy** followed by the name of the policy you want to enable. You must include an existing policy name in this command to enable MAPS. The default policies are:
   - dflt_conservative_policy
   - dflt_aggressive_policy
   - dflt_moderate_policy
   - dflt_base_policy (This policy is only for basic monitoring, similar to using MAPS without a license. Refer to Feature monitors not requiring a Fabric Vision license on page 22 for details.)
3. Set global actions on the switch to “none” by entering **mapsconfig --actions none**.

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

The following example enables MAPS, loads the policy “fw_aggressive_policy”, sets the actions to “none”, and then sets approved actions.

```
switch:admin> mapsconfig --enablemaps -policy fw_aggressive_policy
WARNING: Fabric Watch is discontinued in FOS 7.4 and will not run after firmware
upgrade. To continue with monitoring capability, it is recommended to migrate to MAPS
prior to firmware upgrade. Users can convert existing Fabric Watch thresholds into
MAPS policies by using mapsConfig --fwconvert CLI command and continue monitoring
with the same settings. Fabric Watch thresholds cannot be converted to MAPS policies
after firmware upgrade. Please refer to MAPS Administrator's Guide for further
information."

Do you want to continue [Y/N]: Y
...
MAPS is enabled.

switch:admin> mapsconfig --actions none
switch:admin> mapsconfig --actions raslog,fence,snmp,email,sw_marginal
```

For additional information, refer to the following topics:

- Refer to Predefined policies on page 50 for details on the default MAPS policies.
- Refer to Viewing the MAPS dashboard on page 92 for details on the **mapsdb** command output.
- Refer to MAPS rule actions on page 59 for details on configuring MAPS rule actions.

Monitoring across different time windows

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

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

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

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

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

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

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

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

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

Setting the active MAPS policy to a default policy

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

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

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

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

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

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

For more information, refer to Predefined policies on page 50.

Pausing MAPS monitoring

You can stop monitoring a port or other element in MAPS. You might do this during maintenance operations such as device or server upgrades.

To temporarily stop monitoring a port or other element in MAPS, complete the following steps. This suspends MAPS monitoring for the specified element member.

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

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

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

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

Resuming MAPS monitoring

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

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

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

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

The following example resumes MAPS monitoring for port 5.

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

MAPS Elements and Categories

MAPS structural elements

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

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

TABLE 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, refer to Working with MAPS rules and actions on page 66.
Category	A grouping of similar elements that can be monitored (for example, "Security Violations"). For more information, refer to MAPS monitoring categories on page 31.
Condition	A true or false trigger created by the combination of a timebase and a threshold value. For more information, refer to MAPS conditions on page 55.
Element	A value (measure or statistic) that can be monitored. This includes switch conditions, data traffic levels, error messages, and other values.
Group	A collection of similar objects that you can monitor as a single entity. For example, a collection of ports can be assembled as a group. For more information, refer to MAPS groups overview on page 41.
Rule	A direction associating a condition with one or more actions that must occur when the specified condition is evaluated to be true. For more information, refer to MAPS rules overview on page 59.
Policy	A set of rules defining thresholds for triggering actions MAPS is to take when that threshold is triggered. When a policy is enabled, all of the rules in the policy are in effect. For more information, refer to MAPS policies overview on page 49.

MAPS monitoring categories

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

MAPS provides you with the following monitorable categories:

- Port Health on page 32
- Back-end Health on page 34
- FRU Health on page 34

- Security Violations on page 35
- Fabric State Changes on page 35
- Switch Resource on page 36
- Traffic Performance on page 37
- FCIP Health on page 38
- Fabric Performance Impact on page 39
- Switch Policy Status on page 40

In addition to being able to set alerts and other actions based on these categories, the MAPS dashboard displays their status. Refer to MAPS dashboard overview on page 89 for information on using the MAPS dashboard.

Port Health

The Port Health category monitors port statistics and takes action based on the configured thresholds and actions. You can configure thresholds per port type and apply the configuration to all ports of the specified type. Ports whose thresholds can be monitored include physical ports, D_Ports, E_Ports, F_Ports, and Virtual E_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), transmit power (TXP), and state changes in physical ports, D_Ports, E_Ports, and F_Ports.

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

TABLE 6 Port Health category parameters

Monitored parameter	Description
Cyclic redundancy check (CRC with good EOF (crc g_eof) markers)	The number of times an invalid cyclic redundancy check error occurs on a port or a frame that computes to an invalid CRC. Invalid CRCs can represent noise on the network. Such frames are recoverable by retransmission. Invalid CRCs can indicate a potential hardware problem.
Invalid transmission words (ITW)	The number of times an invalid transmission word error occurs on a port. A word did not transmit successfully, resulting in encoding errors. Invalid word messages usually indicate a hardware problem. **NOTE** The ITW counter includes any physical coding sub-layer (PCS) violations. ITW violations can occur due to an “encoding in” or “encoding out” violation, a PCS violation, or all of these. Encoding violations occur only at slow (8 Gbps or lower) speeds, and PCS violations occur only at high (10 Gbps or higher) speeds.
Sync loss (LOSS_SYNC)	The number of times a synchronization error occurs on the port. Two devices failed to communicate at the same speed. Synchronization errors are always accompanied by a link failure. Loss of synchronization errors frequently occur due to a faulty SFP transceiver or cable.
Link failure (LF)	The number of times a link failure occurs on a port or sends or receives the Not Operational Primitive Sequence (NOPS). Both physical and hardware problems can cause link failures. Link failures also frequently occur due to a loss of synchronization or a loss of signal.
Signal loss (LOSS_SIGNAL)	The number of times that a signal loss occurs in a port. Signal loss indicates that no data is moving through the port. A loss of signal usually indicates a hardware problem.

TABLE 6 Port Health category parameters (Continued)

Monitored parameter	Description
Protocol errors (PE)	The number of times a protocol error occurs on a port. Occasionally, protocol errors occur due to software glitches. Persistent errors generally occur due to hardware problems.
Power-on time (PWR_HRS)	The number of hours that an SFP transceiver has been powered up.
NPIV logins (DEV_NPIV_LOGINS)	The number of NPIV logins to the device. Refer to Monitoring NPIV logins to F_Ports on page 121 for details.
Link reset (LR)	The ports on which the number of link resets exceed the specified threshold value.
Class 3 timeouts (C3TXTO)	The number of Class 3 discard frames because of timeouts.
State changes (STATE_CHG)	The state of the port has changed for one of the following reasons: • The port has gone offline. • The port has come online. • The port is faulty.
SFP current (CURRENT)	The amperage supplied to the SFP transceiver in milliamps (mA). Current area events indicate hardware failures.
SFP receive power (RXP)	The power of the incoming laser in microwatts (µW). This is used to help determine if the SFP transceiver is in good working condition. If the counter often exceeds the threshold, the SFP transceiver is deteriorating.
SFP transmit power (TXP)	The power of the outgoing laser in microwatts (µW). This is used to help determine if the SFP transceiver is in good working condition. If the counter often exceeds the threshold, the SFP transceiver is deteriorating.
SFP voltage (VOLTAGE)	The voltage supplied to the SFP transceiver in millivolts (mV). If this value exceeds the threshold, the SFP transceiver is deteriorating.
SFP temperature (SFP_TEMP)	The temperature of the SFP transceiver in degrees Celsius. A high temperature indicates that the SFP transceiver may be in danger of damage.

Port health and CRC monitoring

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

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

Back-end Health

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

The following table lists the monitored parameters in this category.

TABLE 7 Back-end Health category parameters

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

For more information on back-end health monitoring, refer to Back-end port monitoring on page 82.

FRU Health

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

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

TABLE 8 FRU Health category parameters

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

Security Violations

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

The following table lists the monitored parameters in this category.

TABLE 9 Security Violations category parameters

Monitored parameter	Description
DCC violations (SEC_DCC)	An unauthorized device attempts to log in to a secure fabric.
HTTP violations (SEC_HTTP)	A browser access request reaches a secure switch from an unauthorized IP address.
Illegal command (SEC_CMD)	Commands permitted only to the primary Fibre Channel Switch (FCS) are executed on another switch.
Incompatible security DB (SEC_IDB)	Secure switches with different version stamps have been detected.
Login violations (SEC_LV)	Login violations which occur when a secure fabric detects a login failure.
Invalid Certifications (SEC_CERT)	Certificates are not valid.
No-FCS (SEC_FCS)	The switch has lost contact with the primary FCS.
SCC violations (SEC_SCC)	SCC violations which occur when an unauthorized switch tries to join a secure fabric. The WWN of the unauthorized switch appears in the ERRLOG.
SLAP failures (SEC_AUTH_FAIL)	SLAP failures which occur when packets try to pass from a non-secure switch to a secure fabric.
Telnet violations (SEC_TELNET)	Telnet violations which occur when a Telnet connection request reaches a secure switch from an unauthorized IP address.
TS out of sync (SEC_TS)	Time Server (TS) violations, which occur when an out-of-synchronization error has been detected.

Fabric State Changes

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

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

TABLE 10 Fabric State Changes category parameters

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

Switch Resource

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

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

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

The following table lists the monitored parameters in this category.

TABLE 11 Switch Resource category parameters

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

Traffic Performance

The Traffic Performance category groups areas that track the source and destination of traffic. You can use traffic thresholds and alarms to determine traffic load and flow and to reallocate resources appropriately.

The following table lists the monitored parameters in this category. These parameters are only applicable to Flow Monitor flows.

TABLE 12 Traffic Performance category parameters

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

TABLE 12 Traffic Performance category parameters (Continued)

Monitored parameter	Description
SCSI frames written (IO_WR)	The number of SCSI I/O write command frames recorded for the flow.
SCSI frames read (IO_RD_BYTES)	The number of SCSI I/O bytes read as recorded for the flow.
SCSI frames written (IO_WR_BYTES)	The number of SCSI I/O bytes written as recorded for the flow.

FCIP Health

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

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

TABLE 13 FCIP Health category parameters

Monitored parameter	Description
FCIP circuit state changes (CIR_STATE)	The state of the circuit has changed for one of the following reasons: • The circuit has gone offline. • The circuit has come online. • The circuit is faulty.
FCIP circuit utilization (CIR_UTIL)	The percentage of circuit utilization in the configured time period (this can be minute, hour, or day).
FCIP circuit packet loss (CIR_PKTLOSS)	The percentage of the total number of packets that have had to be retransmitted.
FCIP QoS utilization (UTIL)	The percentage of FCIP circuit QoS groups utilization.
FCIP packet loss (PKTLOSS)	The percentage of the total number of packets that have had to be retransmitted in each QoS level. This applies to each FCIP QoS group only.
FCIP circuit round trip time (RTT)	The circuit round-trip latency. This is an absolute value, and only applies to the circuit group.
FCIP circuit jitter (JITTER)	The amount of jitter in a circuit. This is a calculated percentage, and only applies to the circuit group.

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

TABLE 14 Brocade 7840 FCIP Health category parameters

Monitored parameter	Description
FCIP tunnel state change (STATE_CHG)	The count of FCIP tunnel state changes. This applies to the tunnel group only.

TABLE 14 Brocade 7840 FCIP Health category parameters (Continued)

Monitored parameter	Description
FCIP tunnel or tunnel QoS utilization (UTIL)	The percentage of FCIP utilization. This applies to both the tunnel and the tunnel QoS groups.

Refer to the FCIP monitoring thresholds on page 133 for the default values.

Fabric Performance Impact

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

The following table lists the monitored parameters in this category.

TABLE 15 Fabric Performance Impact category parameters

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

For more information on Fabric Performance Impact monitoring, refer to Fabric performance monitoring using MAPS on page 103.

Switch Policy Status

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

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

NOTE
Not all switches support all monitors.

TABLE 16 Switch Policy Status category parameters

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

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

MAPS Groups, Policies, Rules, and Actions

MAPS groups overview

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

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

Viewing group information

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

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

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

```
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
```

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

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

You can also use this command to display the state of flows from a MAPS perspective. The state of a flow is shown in the output in the “Members” column. The following example shows the output of

logicalgroup --show fpm1 for the active Flow Vision flow “fpm1” that has been imported into, and being monitored through, MAPS.

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

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

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

Predefined groups

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

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

TABLE 17 Predefined MAPS groups

Predefined group name	Object type	Description
ALL_PORTS	FC Port	All ports in the logical switch.
ALL_BE_PORTS	N/A	All back-end ports in the physical switch.
ALL_D_PORTS	FC Port	All D_Ports in the logical switch.
ALL_E_PORTS	FC Port	All E_Ports and EX_Ports in the logical switch. This includes all the ports in E_Port and EX_Port trunks as well.
ALL_F_PORTS	FC Port	All F_Ports in the logical switch. This includes all the ports in F_Port trunks as well.
ALL_HOST_PORTS	FC Port	All ports in the logical switch connected to hosts. MAPS automatically detects if a device connected on this port is a server device and adds it to this set.
ALL_TARGET_PORTS	FC Port	All logical switch ports connected to targets. MAPS automatically detects if a device connected on this port is a target device and adds it to this set. If the device connected to the port is identified as both a Host and a Target device, MAPS treats the port as a Target port.
ALL_OTHER_F_PORTS	FC Port	All F_Ports in the logical switch which are neither Host nor Target ports.

TABLE 17 Predefined MAPS groups (Continued)

Predefined group name	Object type	Description
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. **NOTE** You cannot add or delete members from the ALL_QUARANTINED_PORTS group. No default rules are defined for the group, but you can create custom rules to monitor the ports in this group.
ALL_SFP	SFP	All small form-factor pluggable (SFP) transceivers.
ALL_10GSWL_SFP	SFP	All 10-Gbps Short Wavelength (SWL) SFP transceivers on FC Ports in the logical switch.
ALL_10GLWL_SFP	SFP	All 10-Gbps Long Wavelength (LWL) SFP transceivers on FC Ports in the logical switch.
ALL_16GSWL_SFP	SFP	All 16-Gbps SWL SFP transceivers in the logical switch.
ALL_16GLWL_SFP	SFP	All 16-Gbps LWL SFP transceivers in the logical switch.
ALL_OTHER_SFP	SFP	All SFP transceivers that do not belong to one of the following groups: • ALL_10GSWL_SFP • ALL_10GLWL_SFP • ALL_16GSWL_SFP • ALL_16GLWL_SFP • ALL_QSFP
ALL_QSFP	SFP	All quad small form-factor pluggable (QSFP) transceivers in the logical switch.
ALL_2K_QSFP	SFP	All 2 kilometer-capable quad small form-factor pluggable (QSFP) transceivers 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_CIRCUITS	Circuit	All Fibre Channel over Internet Protocol (FCIP) circuits in the logical switch.
ALL_CIRCUIT_HIGH_QOS	QoS	These are available for Brocade extension switches (Brocade 7800, Brocade 7840, Brocade FX8-24) only. QoS monitoring for circuits is based on pre-defined QoS priorities for throughput and lost packets, and is done at the
ALL_CIRCUIT_MED_QOS	QoS	

TABLE 17 Predefined MAPS groups (Continued)

Predefined group name	Object type	Description
ALL_CIRCUIT_LOW_QOS	QoS	circuit level. These groups correspond to the circuits on the Brocade 7800, Brocade 7840, Brocade FX8-24.
ALL_CIRCUIT_F_QOS	QoS	
ALL_FAN	Fan	All fans in the chassis.
ALL_FLASH	Flash	The flash memory card in the chassis.
ALL_PS	Power Supply	All power supplies in the chassis.
ALL_TS	Temperature Sensor	All temperature sensors in the chassis.
ALL_WWN	WWN	All WWN cards in the chassis.
ALL_TUNNELS	Tunnel	All FCIP tunnels in the switch. These are available for Brocade extension switches (Brocade 7800, Brocade 7840, Brocade FX8-24) only.
ALL_TUNNEL_HIGH_QOS	QoS	These are available for Brocade 7840 devices only. QoS monitoring for tunnels is based on pre-defined QoS priorities for throughput and lost packets, and is done at the tunnel level. These groups correspond to the tunnels on Brocade 7840 devices.
ALL_TUNNEL_MED_QOS	QoS	
ALL_TUNNEL_LOW_QOS	QoS	
ALL_TUNNEL_F_QOS	QoS	
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.

For more information on groups, refer to:

- MAPS groups overview on page 41
- User-defined groups on page 44

User-defined groups

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

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

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

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

- Dynamic groups can only be defined based on the port type.
- The device node WWN information is fetched from the FDMI database, group membership is validated against this database.
- On an Access Gateway device, if a group is created with the feature specified as “device node WWN”, the ports on the switch that are connected to the Access Gateway are not taken into account. However, on the Access Gateway itself, the ports where devices are connected will be part of the group.
- A port or device can be a member of multiple groups.
- A maximum of 64 user-defined groups and imported flows combined is permitted per logical switch.
- All operations on a dynamic group are similar to those for static groups.
- Group names are not case sensitive; My_Group and my_group are considered to be the same.

Creating a static user-defined group

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

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

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

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

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

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

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

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

Modifying a static user-defined group

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

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

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

You need to specify every element in the command. When adding or removing ports, you can specify either a single port (2/15), or specify multiple ports as either individual IDs separated by commas (2/15, 5/16), or a range of ports with the IDs separated by hyphens (2/15-16/15).

3. Optional: Enter **logicalGroup --show** *group_name* **-details** to view the group membership.

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

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

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

Creating a dynamic user-defined group

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

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

TABLE 18 Group-definition operators

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

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

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

   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" whose membership is defined as ports belonging to a device whose node WWN starts with 30:08:00:05.

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

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

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

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

Modifying a dynamic user-defined group

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

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

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

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

3. Use the following commands to add or delete specific ports from the group, or modify the group definition.
   - 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*.
   - To modify the group definition, enter **logicalGroup --update** *group_name* **-feature** *feature_name* **-pattern** *pattern*.

   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).
4. Optional: Enter **logicalGroup --show** *group_name* **-details** to view the group membership.

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

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

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

Restoring a group to its default membership

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

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

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

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

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

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

switch:admin> logicalgroup --restore GOBLIN_PORTS

switch:admin> logicalgroup --show GOBLIN_PORTS -detail

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

Cloning a group

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

To clone a group, complete the following steps.

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

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

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

Deleting groups

The **logicalgroup --delete** *group_name* command allows you to remove any logical group of monitoring elements 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 all the rules that are configured with the given group and then deletes the group. If a logical group is present in user-defined rules, the **-force** option deletes all the rules that are configured with the given group and then deletes the group.

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

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

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

MAPS policies overview

A MAPS policy is a set of rules that define thresholds for measures and action to take when a threshold is triggered. When you enable a policy, all of the rules in the policy are in effect.

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

The following restrictions apply to policies:

- Only one policy can be active at a time.
- When you enable a policy, it becomes the active policy and the rules in the active policy take effect.
- One policy must always be active on the switch.
  - You can have an active policy with no rules, but you must have an active policy.
  - You cannot disable the active policy. You can only change the active policy by enabling a different policy.

Viewing policy values

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

The following example displays all the thresholds for D_Ports in the dflt_conservative_policy 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)
```

Predefined policies

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

The four 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.
- 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

  This policy contains rules that monitor the unlicensed features which were made available earlier through Fabric Watch. Refer to Feature monitors not requiring a Fabric Vision license on page 22 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. For more information, refer to the following links.

- Modifying a default policy on page 55
- Creating a policy on page 53
- User-defined policies on page 51

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

For System z and FICON environments, Brocade recommends that you start with the Aggressive policy. For Open Systems environments and other environments, Brocade recommends that you start with the Moderate policy.

Default MAPS policy rules

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

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

User-defined policies

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

Refer to Working with MAPS policies on page 51 for information on working with user-defined policies.

Fabric Watch legacy policies

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

- fw_custom_policy

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

- fw_default_policy

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

- fw_active_policy

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

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

The following factors also apply to Fabric Watch conversions:

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

Working with MAPS policies

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

Viewing policy information

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

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

1. Connect to the switch and log in using an account with admin permissions.
2. Choose from the following options:

- To view a summary of all the policies on the switch, enter **mapspolicy --show -summary**. This displays the policy names and the number of rules in each policy.
- To view the features of all the policies on the switch, enter **mapspolicy --show -all**. This displays the policy names, rule names, actions, and conditions for each policy for all policies.
- To view the features of a specific policy on the switch, enter **mapspolicy --show** *policy_name*. This displays the policy names, rule names, actions, and conditions for the named policy.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Creating a policy

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

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

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

1. Create a new policy or clone a policy from one of your existing policies.

- To create a new policy, enter **mapsPolicy --create** *policy_name* to create a policy.
- To clone an existing policy, enter **mapsPolicy --clone** *policy_name* **-name** *clone_policy_name*.

2. Create or modify rules to configure the required thresholds in the new policy.

   - To create a rule, enter **mapsRule --create** *rule_name* **-group** *group_name* **-monitor** *ms name* **-timebase** *timebase* **-op** *op_value* **-value** *value* **-action** *action* **-policy** *policy_name*.
   - To clone an existing rule, enter **mapsRule --clone** *rule_name* **-name** *clone_rule_name*.
   - To modify existing rules, enter **mapsRule --config** *rule_name parameters*.

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

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

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

Enabling a policy

Only one policy can be enabled at a time, and it must be enabled before it takes effect. If the active policy is changed, or if the rules in the active policy are changed, the active policy must be re-enabled for the changes to take effect.

To enable a policy, complete the following steps.

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

The following example enables the “dflt_aggressive_policy” policy.

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

Modifying a user-defined policy

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

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

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

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

   - Use **mapspolicy** to add rules to and delete rules from the policy.
   - Use **mapsrule** to modify rules or to create rules and add them to the policy.

2. Optional: Even if the policy is the active policy, you must re-enable the policy using the **mapspolicy --enable** *policy_name* command for the changes to take effect. Adding a new rule or changing an existing rule in the active policy does not take effect until you re-enable the policy.

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

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

switch:admin> mapspolicy --show daily_policy

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

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

Modifying a default policy

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

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

1. Create a copy of the default policy.
   ```
   switch:admin> mapspolicy --clone dflt_conservative_policy -name my_policy
   ```
2. Modify the rules in the policy based on your requirements.

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

   Use **mapsPolicy** to add and delete rules to and from the policy. Use **mapsRule** to create rules and add them to the policy.
3. Enable the policy.
   ```
   switch:admin> mapspolicy --enable my_policy
   ```
   The previously enabled policy is disabled, and the specified policy is enabled.

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

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

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

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

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

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

MAPS conditions

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

For example, if you specified a rule to be triggered if the CRC counter in the last minute is greater than 10, the threshold value is 10 and the timebase is the preceding 60 seconds. 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 60-second timebase. If the counter reaches 11 within that 60 seconds, the rule would trigger.

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

Threshold values

Thresholds are the values at which potential problems may occur. In configuring a rule you can specify a threshold value that, when exceeded, triggers an action. For example, if you had specified a rule to make a RASLog entry if the CRC counter is greater than 10; when the counter hits 11, the rule is triggered and a RASLog entry is made.

Timebase

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

You can set the timebase to “day” (samples are compared once a day), “hour” (samples are compared once an hour), “minute” (samples are compared every minute), or “none” (for comparisons where the timebase is not applicable). In addition, you can specify a timebase multiplier value, allowing the timebase to be a multiple of minutes, hours, or days. The value specified for **-value** should be a positive number and less than the total number of samples in the timebase+1 period.

The following example sets the timebase for the rule “BiWeekly” to be every two weeks (14 days). The timebase portion of the command has been bolded for ease of viewing.

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

Supported timebases

The following table identifies which monitors support which timebases.

TABLE 19 Monitors and supported timebases

Monitor Name	Day	Hour	Minute	None
Domain ID change	Yes	Yes	Yes	No
Fabric logins	Yes	Yes	Yes	No
Fabric reconfigurations	Yes	Yes	Yes	No
E_Ports down	Yes	Yes	Yes	No
Segmentation changes	Yes	Yes	Yes	No
Zone changes	Yes	Yes	Yes	No
L2 Device Count	Yes	Yes	Yes	No
LSAN Device Count	Yes	Yes	Yes	No
Zone Configuration size	Yes	Yes	Yes	No

TABLE 19 Monitors and supported timebases (Continued)

Monitor Name	Day	Hour	Minute	None
FCR Count	Yes	Yes	Yes	No
Tunnel (STATE_CHG)	Yes	Yes	Yes	
Tunnel QoS (UTIL)	Yes	Yes	Yes	
QoS packet loss percentage (PKTLOSS)	Yes	Yes	Yes	
Circuit state (CIR_STATE)	Yes	Yes	Yes	
Circuit utilization percentage (CIR_UTIL)	Yes	Yes	Yes	
Circuit packet loss percentage (CIR_PKTLOSS)	Yes	Yes	Yes	
Circuit round-trip times (RTT)	Yes	Yes	Yes	
Circuit jitter (JITTER)	Yes	Yes	Yes	
Power Supply (PS_STATE)	No	No	No	No
Fan (FAN_STATE)	No	No	No	No
Blade (BLADE_STATE)	No	No	No	No
SFP (SFP_STATE)	No	No	No	No
WWN (WWN)	No	No	No	No
CRC Errors	Yes	Yes	Yes	
Invalid Transmit Words	Yes	Yes	Yes	
Sync Loss	Yes	Yes	Yes	
Link Failure	Yes	Yes	Yes	
Loss of Signal	Yes	Yes	Yes	
Protocol Errors	Yes	Yes	Yes	
Link Reset	Yes	Yes	Yes	
C3 Time outs	Yes	Yes	Yes	
State change	Yes	Yes	Yes	
SFP Current	No	No	No	Yes
SFP Receive Power	No	No	No	Yes
SFP Transmit Power	No	No	No	Yes
SFP Voltage	No	No	No	Yes
SFP Temperature	No	No	No	Yes
SFP Power On Hours	No	No	No	Yes
Flash memory percentage used (FLASH_USAGE)	No	No	No	Yes
CPU percentage used (CPU)	No	No	No	Yes
Memory percentage used (MEMORY_USAGE)	No	No	No	Yes
Ethernet management port state (ETH_MGMT_PORT_STATE)	No	No	No	Yes

TABLE 19 Monitors and supported timebases (Continued)

Monitor Name	Day	Hour	Minute	None
Temperature Sensor (TEMP)	No	No	No	Yes
DCC violations	Yes	Yes	Yes	No
HTTP violation	Yes	Yes	Yes	No
Illegal command	Yes	Yes	Yes	No
Incompatible security DB	Yes	Yes	Yes	No
Login violations	Yes	Yes	Yes	No
Invalid certifications	Yes	Yes	Yes	No
No-FCS	Yes	Yes	Yes	No
SCC violations	Yes	Yes	Yes	No
SLAP failures	Yes	Yes	Yes	No
Telnet violations	Yes	Yes	Yes	No
TS out of sync	Yes	Yes	Yes	No
Current (CURRENT)	No	No	No	No
Receive Power (RXP)	No	No	No	No
Transmit Power (TXP)	No	No	No	No
Voltage (VOLTAGE)	No	No	No	No
Temperature (TEMP)	No	No	No	No
Fabric Performance Impact (DEV_LATENCY_IMPACT)	No	No	No	Yes
Receive Bandwidth usage percentage (RX)	Yes	Yes	Yes	
Transmit Bandwidth usage percentage (TX)	Yes	Yes	Yes	
Trunk Utilization percentage (UTIL)	Yes	Yes	Yes	
Absent or faulty power supply (BAD_PWR)	No	No	No	Yes
Temperature sensors outside range (BAD_TEMP)	No	No	No	Yes
Absent or faulty fans (BAD_FAN)	No	No	No	Yes
Flash usage (FLASH_USAGE)	No	No	No	Yes
Percentage of marginal ports (MARG_PORTS)	No	No	No	Yes
Percentage of error ports (ERR_PORTS)	No	No	No	Yes
Percentage of faulty ports (FAULTY_PORTS)	No	No	No	Yes
Faulty blades (FAULTY_BLADE)	No	No	No	Yes
Faulty WWN (WWN_DOWN)	No	No	No	Yes
Core blade monitoring (DOWN_CORE)	No	No	No	Yes
HA Sync (HA_SYNC)	No	No	No	Yes
Receive throughput	Yes	Yes	Yes	

TABLE 19 Monitors and supported timebases (Continued)

Monitor Name	Day	Hour	Minute	None
Transmit Frame Count	Yes	Yes	Yes	
Receive Frame Count	Yes	Yes	Yes	
Transmit throughput	Yes	Yes	Yes	
IO Read Command Count	Yes	Yes	Yes	
IO Write Command Count	Yes	Yes	Yes	
IO Read Data	Yes	Yes	Yes	
IO Write Data	Yes	Yes	Yes	
Throughput Degradation	Yes	Yes	Yes	Yes

MAPS rules overview

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

Each rule specifies the following items:

- A group of objects to be evaluated. Refer to MAPS groups overview on page 41 for additional information.
- The element to be monitored. Refer to MAPS conditions on page 55 for additional information.
- The condition being monitored. Each rule specifies a single condition. A condition includes a timebase and a threshold. Refer to MAPS conditions on page 55 for additional information.
- The actions to take if the condition is evaluated to be true. Refer to MAPS rule actions on page 59 for additional information.

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

MAPS rule actions

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

MAPS provides the following actions for rules:

- E-mail alerts on page 61
- FICON alerts
- Port fencing and port decommissioning on page 61
- RASLog messages on page 64
- SFP marginal on page 64
- Slow Drain Device Quarantine
- SNMP traps on page 65
- Switch critical on page 66

- Switch marginal on page 66
- Port toggling
- Although it is not an action as such, MAPS allows you to define a “quiet time” for most rules in order to reduce the number of alert messages generated. Refer to Quieting a rule on page 70 for details.

For rules that are state-bound (such as those that only have conditions such 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 an error condition.

The global action settings on a switch take precedence over the actions defined in the rules. For example, if the global action settings allow RASLog alerts, but do not allow port fencing, then in the example given in the previous paragraph, if the CRC threshold is reached a RASLog message would be issued but the port would not be fenced. To enable global actions, use the **mapsconfig --actions** commands. For more details, refer to Enabling or disabling rule actions at a global level on page 60. Refer to the *Fabric OS Command Reference* for further details on using the **mapsconfig** and **mapsrule** command.

Enabling or disabling rule actions at a global level

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

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

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

The following example shows that RASLog notification (**raslog**) is not an active action on the switch, and then adds it to the list of allowed actions.

```
switch:admin> mapsconfig --show
Configured Notifications:       EMAIL,DECOM
Mail Recipient:                 admin@mycompany.com
Paused members :
PORT :
CIRCUIT :
SFP :

switch:admin> mapsconfig --actions raslog,email,decom

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

E-mail alerts

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

To configure the e-mail recipients, use the **mapsConfig --emailcfg** command. Multiple destination e-mail addresses are possible; they must be separated using a comma between each individual address and each address must be a complete e-mail address. For example, abc@12.com is a valid e-mail address; abc@12 is not. Refer to Sending alerts using e-mail on page 72 for more information.

To clear all configured email address(es), enter **mapsconfig --emailcfg -address none**. All configured email addresses will be erased.

Port fencing and port decommissioning

MAPS supports port fencing and port decommissioning for both E_Ports and F_Ports. These actions automatically take ports offline when configured thresholds in a given rule are exceeded. Port fencing immediately takes ports offline, which may cause loss of traffic. Port decommissioning takes a port offline more slowly, but without loss of traffic. Both are disabled by default. Port decommissioning and port fencing can only be configured for the port health monitoring system rules, the monitoring systems for which decommissioning is supported.

Port decommissioning cannot be configured by itself in a MAPS rule or action, as 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, MAPS behaves the same as it did in Fabric OS 7.2.x.

Port fencing and port decommissioning actions in MAPS are valid only for conditions evaluated on physical ports, E_Ports, and F_Ports, and can only be configured using the Port Health monitoring rules. Refer to the Port Health monitoring thresholds on page 134 tables for these rules.

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

Port decommissioning for E_Ports and F_Ports

For E_Ports, if port decommissioning fails, MAPS will fence the port. Switches themselves can decommission E_Ports through MAPS. In this case, when port decommissioning is triggered on an E_Port, the neighboring switches will perform a handshake so that traffic is re-routed before the port is disabled. Be aware that there are multiple reasons that the port-decommissioning operation between two E_Ports could fail; for example, if the link that fails is the last link between the two switches. To see which parameters can trigger port fencing and port decommissioning, refer to Port Health monitoring thresholds on page 134.

For F_Ports, port decommissioning will only work if BNA has been configured to perform the port decommissioning and port fencing actions, and the switch running MAPS has SNMP traps enabled. BNA can decommission F_Ports based on CRC, ITW, PE, LR, STATE_CHG, or C3TXTO criteria. MAPS notifications are integrated with BNA, which in turn must coordinate with the switch and the end device to orchestrate the port decommissioning. If BNA is not configured on a switch, MAPS will fence the F_Port.

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

Configuring port decommissioning

Port decommissioning can be configured in two ways. It can be configured along with port fencing in the MAPS actions configuration, or it can be configured as an action in a MAPS rule.

To enable port decommissioning, complete the following steps.

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

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

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

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

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

Port decommissioning and firmware downgrades

- If the port decommissioning (**decom**) configuration is in any of the logical switches, the **firmwareDownload** operation from Fabric OS 7.4.0 to Fabric OS 7.2.x version will fail with one of the following messages:
  - If the **decom** action is configured in the MAPS actions, the following error message is displayed.
    ```
    Downgrade is not allowed, because port decommmission(decom) actions are enabled
    in some logical switches.
    Please delete decom action from the FID(S) <fid1>, <fid2> …
    ```
  - If any of the MAPS rules have the **decom** action configured, the following error message is displayed.
    ```
    Downgrade is not allowed, because port decommission(decom) actions are
    configured in user defined rules in some logical switches.
     Please delete decom action from all user defined rules in the FID(S)
    <fid1>,<fid2> …
    ```
- If there are any default policy rules present with port decommissioning configured, the firmware downgrade is not blocked as in this case, the decommissioning rules are mapped to Fabric OS 7.2.x port fencing rules. That is, a default Fabric OS 7.4.0 MAPS rule with port commissioning specified will remain mapped to the same rule, but without port decommissioning as an action, when the switch is downgraded to version Fabric OS 7.2.x.
- Currently the decommission action is present for the port monitoring rules in dflt_aggressive_policy. When the switch is rebooted using version Fabric OS 7.2.x, the default rules in dflt_aggressive_policy which had port decommissioning specified will have port fencing specified.
- User-defined rules in the active policy are not checked to see if they have port decommissioning configured, as user-defined rules in the active policy are present only in memory and are erased as soon as a different policy is enabled, whether in Fabric OS 7.4.0 or any earlier version.

Enabling port fencing

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

To enable port fencing, complete the following steps.

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

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

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

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

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

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

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

RASLog messages

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

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

TABLE 20 RASLog message category for non-state based monitoring system

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 21 RASLog message category for state based monitoring system

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.

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

SFP marginal

The **SPF_MARGINAL** action sets the state of the affected small form-factor pluggable (SFP) transceiver in the MAPS dashboard to “down.” This action does not bring the SFP transceiver down, but only affects what is displayed in the dashboard. This action is valid only in the context of Advanced SFP groups.

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

Example of health status output before triggering SFP_MARGINAL

A MAPS rule can contain the **SFP_MARGINAL** action. Before the action is triggered, the health status of the ports, as shown by the **sfpshow -health** command, are not affected.

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

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

Example of health status output after triggering SFP_MARGINAL

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

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

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

SNMP traps

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 may be the most efficient notification method. You can avoid having to log in to each switch individually as you would have to do for error log notifications.

When specific events occur on a switch, SNMP generates a message (called a “trap”) that notifies a management station using SNMP. Log entries can also trigger SNMP traps if the SNMP agent is configured. When the SNMP agent is configured to a specific error message level, error messages at that level trigger SNMP traps.

An SNMP trap forwards the following information to an SNMP management station:

- Name of the element whose counter registered an event
- Area and index number of the threshold that the counter crossed

- Event type
- Value of the counter that exceeded the threshold
- State of the element that triggered the alarm
- Source of the trap

The SNMP trap only stores event information. In order to get the event notifications, you must configure the SNMP software to receive the trap information from the network device, and configure the SNMP agent on the switch to send the trap to the management station. You can configure SNMP notifications using the **snmpconfig** command or configure the switch IP in Brocade Network Advisor (refer to “Event notification” in the *Brocade Network Advisor User's Manual* or online help). For additional information on configuring the SNMP agent using **snmpconfig**, refer to the *Fabric OS Command Reference*.

SNMP MIB support

MAPS requires SNMP management information base (MIB) support on the device for management information collection. For additional information on SNMP MIB support, refer to the *Fabric OS Administrator's Guide*.

Switch critical

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

Switch marginal

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

Working with MAPS rules and actions

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

Viewing MAPS rules

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

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

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

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

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

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

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

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

```
switch:admin> mapsrule --show –all -concise
Rule Name                               |Condition                                    |Actions     |
----------------------------------------------------------------------------------------------------
defNON_E_F_PORTSCRC_0                   |NON_E_F_PORTS(CRC/MIN>0)                     |RS,SN,EM    |
defNON_E_F_PORTSCRC_2                   |NON_E_F_PORTS(CRC/MIN>2)                     |PF,SN,EM    |
defNON_E_F_PORTSCRC_10                  |NON_E_F_PORTS(CRC/MIN>10)                    |RS,SN,EM    |
defNON_E_F_PORTSCRC_20                  |NON_E_F_PORTS(CRC/MIN>20)                    |PF,SN,EM    |

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 rule

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

To create a policy rule, complete the following steps.

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

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

Example of creating a rule to generate a RASLog message

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

```
switch246:FID24:admin> mapsrule --create check_crc -monitor crc -group
critical_ports -timebase hour -op g -value 10 –action raslog -policy daily_policy
switch246:FID24:admin> mapsrule --show check_crc

Rule Data:
----------
RuleName: check_crc
Condition: critical_ports(crc/hour>10)
Actions: raslog
Policies Associated: daily_policy

switch246:FID24:admin> mapspolicy --enable daily_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.

```
switch246:FID24:admin> mapsrule --create check_crc2 -monitor crc -group MyFlow -
timebase min -op g -value 15 -action raslog -policy daily_policy2
```

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

Modifying a MAPS policy rule

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

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

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

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

Changing one parameter

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

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

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

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

Changing multiple parameters

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

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

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

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

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

Cloning a rule

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

To clone a MAPS rule, complete the following steps.

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

NOTE
If no parameters other than **--clone** *oldRuleName* **-rulename** *newRuleName* are specified, an exact copy of the original rule is created.

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

Creating an exact clone

The following example shows the existing rule ("old_rule"), creates an exact clone of that rule and names it "new_rule", and then displays the cloned rule.

```
switch:admin> mapsrule --show old_rule
Rule Data:
----------
RuleName: old_rule
Action: Raslog, Fence, SNMP
Condition: Switch(SEC_IDB/Min>0)
Associated Policies: none

switch:admin> mapsrule --clone old_rule -rulename new_rule

switch:admin> mapsrule --show new_rule
Rule Data:
----------
RuleName: new_rule
Action: Raslog, Fence, SNMP
Condition: Switch(SEC_IDB/Min>0)
Associated Policies: none
```

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. To modify the rule, use the **--config** keyword. The following example clones "myOldRule" as "myNewRule" and changes the flow that is being monitored to "flow2" and assigns it the monitor "monitor2". It then displays the rule.

```
switch:admin> mapsrule --clone myOldRule -rulename myNewRule

switch:admin> mapsrule --config -group flow2 -monitor monitor2

admin> mapsrule --show myNewRule
Rule Data:
----------
RuleName: myNewRule
Action: Raslog, Fence, Decom
Condition: Switch(SEC_IDB/Hour>10)
Associated Policies: slow_monitor2
```

Cloning a rule and changing its timebase

The following example creates a clone of "Rule1" with a timebase of an hour, and then displays the rule.

```
switch:admin> mapsrule --clone Rule1 -rulename NewRule2 -timebase hour

switch:admin> mapsrule --show NewRule2
Rule Data:
----------
RuleName: NewRule2
Action: Raslog, Fence, SNMP
Condition: Switch(SEC_IDB/Hour>0)
Associated Policies: none
```

Quieting a rule

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

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 interferes with monitoring for other faults and can rapidly fill an e-mail box. When the quiet time parameter is set, the notification is sent once, and then MAPS waits until the quiet time expires. When this happens, MAPS sends an update alert which is the same as the initial alert but includes a notification of how many times the rule has been triggered in the interim. After MAPS sends the update, it resets the quiet time timer. The following example highlights the repetition section in boldface.

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

Quiet time does not work with rules that define port fencing or port decommissioning, as there is not a good reason for those actions to be deferred. You can configure quiet time with the port fencing action, but MAPS will ignore the quiet time parameter in this case. Quiet time is not supported for monitoring TX_FCNT, RX_FCNT, TX_THPUT, RX_THPUT, IO_RD, IO_WR, IO_RD_BYTES , or IO_WR_BYTES values.

The length of a quiet time period is measured in seconds, with the minimum value determined by the timebase defined in the rule. For all time bases other than “NONE”, the minimum quiet time is 0 seconds and there is no predefined upper limit. If the rule timebase is “NONE”, then the quiet time value is different for different monitored values, as shown in the following table.

TABLE 22 Minimum quiet time values for monitoring systems that support a time base of NONE

Monitored value	Minimum quiet time (seconds)
DEV_LATENCY_IMPACT	60
DEV_NPIV_LOGINS	3600
FLASH_USAGE	60
CPU	120
MEMORY_USAGE	120
TEMP	60
PS_STATE	15
FAN_STATE	15
BLADE_STATE	15
SFP_TEMP	360
VOLTAGE	360
CURRENT	360
RXP	360
TXP	360
PWR_HRS	360
BAD_TEMP	60
BAD_PWR	60
BAD_FAN	60

The following example shows a rule that includes a 120 second quiet time period. The quiet time element has been 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
```

Rule deletion

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

Although you can use the **mapsrule --delete** *rule_name* command to delete individual instances of a user-defined rule, you must remove the rule individually from every policy that uses the rule before you can finally delete the rule itself. This could be problematic if the rule has been added to many policies. If you try to remove a rule while it is still in a policy without using the **-force** option, it will fail. 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 is deleted and then the group is deleted. Refer to Deleting groups on page 49 for information on deleting groups.

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

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

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

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

Sending alerts using e-mail

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

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

1. Configure and validate the e-mail server. Refer to Configuring e-mail server information on page 74for information on specifying the e-mail server to be used.
2. Enter the **mapsconfig --emailcfg** command to set the e-mail parameters.

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

   NOTE
   You can also send a test e-mail alert. Refer to E-mail alert testing on page 73 for additional information.

Specifying e-mail address for alerts

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

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

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

Specifying multiple e-mail addresses for alerts

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

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

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

Clearing the configured e-mail address

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

E-mail alert testing

Fabric OS allows you to 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 23 Test e-mail command parameters

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

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

Configuring e-mail server information

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

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

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

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

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **relayconfig --config -rla_ip** *relay IP address* **-rla_dname** “*relay domain name*”. The quotation marks are required.
   There is no confirmation of this action.
3. Optional: Enter **relayconfig --show**.
   This displays the configured e-mail server host address and domain name.

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

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

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

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

Viewing configured e-mail server information

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

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

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

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

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

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

Deleting e-mail server configuration

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

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

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

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

```
switch:admin> relayconfig --delete

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

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

Port Monitoring Using MAPS

Port monitoring and pausing

Pausing operations on a port does not affect flow monitoring. Flow monitoring is done at the flow level and the details of the flow passing through a particular port is transparent to MAPS.

Monitoring similar ports using the same rules

You can create groups of ports that behave in a similar manner and use this group to more easily monitor the ports using a single set of rules and thresholds. MAPS refers to these as “logical groups”.

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

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

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

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

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

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

Port monitoring using port names

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

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

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

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

Port monitoring using device WWNs

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

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

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

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

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

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

Adding a port to an existing group

If a new element, such as a host, target, or small form-factor pluggable (SFP) transceiver is added to the fabric, 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.

The following items should be kept in mind for this monitoring:

- Elements that are added manually to a group remain in the group whether they are online or offline.
- There is no validation of manual additions to a group; for example, if you add port 17 as part of an F_Port group, that port is added to the group even if it is not actually an F_Port.
- You can add a port to a predefined group, but not to the group ALL_QUARANTINED_PORTS.

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

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

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

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

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

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

Listing this group produces the following output.

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

If ports 15 and 41 go offline, the following output would result.

```
switch:admin> logicalgroup --show critical_ports
------------------------------------------------------------------
Group Name     |Predefined |Type |Member Count |Members
------------------------------------------------------------------
critical_ports |No         |Port |4            |10,25,31,41
```

NOTE
Port 41 is still considered part of the critical_ports group, even if it is offline.

Adding missing ports to a group

You can add ports to a predefined group (for example, ALL_HOST or ALL_TARGET) that may not have been included automatically.

NOTE
The same restrictions as described in Adding a port to an existing group on page 78 apply.

1. Enter **logicalgroup --show** *group_name*.
2. Enter **logicalgroup --addmember** *group_name* **-member** *member_list* to add the specified port to the named group.

For dynamic groups, 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. Dynamic groups can be created using either port names or WWNs, but you cannot use both in a single group definition. Once 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.

3. Optional: Enter **logicalgroup --show** *group_name* to confirm the addition.

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

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

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

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

Removing ports from a group

In addition to adding ports to either predefined or user-defined groups, you can remove the ports from either group type. This is useful for devices that erroneously identify themselves as both host and target.

To remove a port from a group, complete the following steps.

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

   You can specify either a single port, or specify multiple ports as either individual IDs separated by commas, or a range where the IDs are separated by a hyphen.
3. Optional: Enter **logicalGroup --show** *group_name* to confirm that the named ports are no longer part of the group.

The following example removes port 5 from the ALL_TARGET_PORTS group, and then shows that it is no longer a member of that group.

```
switch:admin> logicalgroup --delmember ALL_TARGET_PORTS -members "5"

switch:admin> logicalgroup --show ALL_TARGET_PORTS
------------------------------------------------------------------------------------
Group Name       |Predefined |Type |Member Count |Members
------------------------------------------------------------------------------------
ALL_TARGET_PORTS |Yes        |Port |5            |1,11,22,32,44
```

D_Port monitoring

In Fabric OS 7.3.0 and later, D_Ports can be monitored by MAPS using the group ALL_D_PORTS.

You can either configure a port as a D_Port using the CLI, or Fabric OS can dynamically convert a port to a D_Port. Refer to the *Flow Vision Administrator's Guide* for information on enabling the dynamic conversion. 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.

Rules based on the ALL_D_PORTS group are part of the default policies, and have error thresholds spanning multiple time windows or timebases. If any of the rules are triggered, MAPS triggers the action configured for the rule, alerts the fabric service module if appropriate, and caches the data in the dashboard.

The D_Port diagnostic output classifies the error conditions into the following states:

- Errors within operating range
- Errors outside of operating range

D_Port monitoring monitors all D_Port errors; however, the fabric service module is only notified for the following errors:

- CRC
- ITW — for the enc_out and enc_in invalid transmission words only
- LF
- LOSS_SYNC

The D_Port monitoring feature is only supported for 10 Gbps and 16 Gbps SFPs and 8Gbps LWL and ELWL ports on the following blades: CR16-4, CR16-8, FC8-32E, FC8-48E, FC16-32, FC16-48, and FC16-64.

NOTE
In versions of Fabric OS prior to 7.3, MAPS monitored D_Ports using the NON_E_F_PORTS group, but the default rules for this group did not provide the flexibility now available through the ALL_D_PORTS group.

The **mapsrule** command accepts the ALL_D_PORTS group, which can be used as shown in the following example.

```
mapsrule --create d_port_mon -group ALL_D_PORTS -monitor CRC -timebase min -op ge -
value 1 -action raslog -policy nil
```

Using the **mapsdb --show** command shows any error or rule violation during diagnostics tests on a D_Port.

```
switch:admin> mapsdb --show
1 Dashboard Information:
=======================
DB start time:                  Wed Mar 26 10:02:38 2014
Active policy:                  dflt_moderate_policy
Configured Notifications:       SW_CRITICAL,SW_MARGINAL
Fenced Ports :                  None
Decommissioned Ports :          None
Quarantined  Ports :            None

2 Switch Health Report:
=======================
Current Switch Policy Status: MARGINAL
Contributing Factors:
---------------------
*BAD_PWR (MARGINAL).

3.1 Summary Report:
===================
Category                 |Today                     |Last 7 days        |
-------------------------------------------------------------------------
Port Health              |Out of operating range    |In operating range |
BE Port Health           |No Errors                 |No Errors          |
Fru Health               |In operating range        |In operating range |
Security Violations      |No Errors                 |No Errors          |
Fabric State Changes     |Out of operating range    |In operating range |
Switch Resource          |In operating range        |In operating range |
Traffic Performance      |In operating range        |In operating range |
FCIP Health              |Not applicable            |Not applicable     |
Fabric Performance Impact|In operating range        |In operating range |

3.2 Rules Affecting Health:
===========================
Category(Rule Count)|RptCount|Rule Name           |Execution Time   |Object |Triggered Value(Units)|
----------------------------------------------------------------------------------------------------
Port Health(5)      |1       |defALL_D_PORTSCRC_1 |05/07/14 08:43:32|D_Port 20|300 Errors          |
                    |4       |defNON_E_F_PORTSLF_0|05/07/14 08:42:56|D_Port 7 |6                   |
                    |        |                    |                 |D_Port 7 |6                   |
                    |        |                    |                 |D_Port 7 |6                   |
                    |        |                    |                 |D_Port 7 |7                   |
```

You can also run the **portdporttest --show** *port_number* command to see details of an individual port. The following example shows the results for port 28.

```
switch:admin> portdporttest --show 28
D-Port Information:
===================
Port: 28
Remote WWNN: 10:00:00:05:1e:e5:e4:00
Remote port: 164
Mode: Manual
Start time: Thu Nov 7 13:43:26 2013
End time: Thu Nov 7 13:53:43 2013
Status: PASSED*
```

Refer to the *Fabric OS Command Reference* for additional information on these commands.

Back-end port monitoring

In Fabric OS 7.4.0 and later, MAPS provides RASLog, SNMP, and e-mail alerts for port counter errors on back-end ports in chassis-based switches, as well as the Brocade 5300 and Brocade 6520 switches. If you know about these errors, you can then do SerDes (Serializer/Deserializer) tuning on those ports where errors exceed the desired threshold, which will greatly improve the packet forwarding performance of these ports.

In terms of monitoring system functionality, back-end ports can be connected to ports within a fixed-port switch or to other blades within the switch chassis, and so their functionality is different from front-end ports, which connect to devices outside of the switch. The primary task of back-end ports is to route packets passing through a switch's ASICs. Switch (and consequently fabric) performance degrades when there are errors in back-end ports. MAPS error notification allows you to take corrective action earlier.

When a switch is initialized, all back-end ports are automatically brought online and stay online until the slot is powered off or the blade is removed (for chassis-based switches), or if the switch itself goes down (for fixed-port switches). This allows MAPS to continuously monitor the platform for back-end port errors. The monitoring systems are for a given switch or chassis, so all the monitoring systems are monitored only in the default switch.

MAPS monitors the port counter statistics for back-end ports through the group ALL_BE_PORTS, which identifies each port using a slot-port combination such as 3/3/1. In case of fixed-port switches, the slot number is 0. WWN IDs and Port Names are not supported. The predefined groups for front-end ports do not apply to back-end ports. History data for back-end ports is collected for a period of seven days and is displayed in the “Backend port History Data” section of the MAPS dashboard.

The following example is typical of a RASLog message generated for a back-end port. In this example, the rule set the threshold at more than 35 CRC errors in a minute (CRC/min>35).

```
2014/08/14-19:58:01, [MAPS-1003], 6313, SLOT 5 | FID 128, WARNING, sw0, BE
Port 6/8, Condition=ALL_BE_PORTS(CRC/min>35), Current Value:[CRC,96 CRCs],
RuleName=test_port_rule_90, Dashboard Category=N/A.
```

For more information on Back-end health monitoring, refer to Back-end Health on page 34 and Back-end port monitoring thresholds on page 132.

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 shown in bold for highlighting. Be aware that the column headings in the example have been edited slightly so as 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 | 07/21/14 01:30:60 |Port 6/8   | 50 CRCs

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

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

Monitoring Flow Vision Flows with MAPS

Viewing Flow Vision Flow Monitor data with MAPS

The Monitoring and Alerting Policy Suite (MAPS) can monitor flows created using the Flow Monitor feature of Flow Vision. Flows created by the Flow Vision Flow Generator and Flow Mirror features cannot be monitored using MAPS.

For details on flows and Flow Vision, refer to the Flow Monitor section of the *Flow Vision Administrator's Guide*.

On Brocade Analytics Monitoring Platform, the rules and flows can be created on a single platform. The AMT learns all the flows in the fabric, making it easy to monitor and configure flows in a large fabric. There is no need to login to multiple switches to configure the flows.

To monitor flows using MAPS, complete the following steps.

1. Create the flow in Flow Vision using the **flow --create** command.
2. Import the flow into MAPS using the **mapsconfig --import** command.
3. Enter **logicalgroup --show** to confirm that the flow was correctly imported into MAPS. The imported flow name indicates the groups that can be monitored.
4. Define a MAPS rule using the **mapsrule --create** command (for the supported timebases).

   Refer to MAPS rules overview on page 59 for information on creating and using rules.
5. Enter **mapspolicy --enablepolicy** *policy_name* to activate the flow.

The following example illustrates the flow-monitoring steps. The first command line creates the flow, the second command line imports it, and the third command line displays the members of the logical groups, with the imported flow in bold (only for the example). The fourth command line creates a rule for the group and the fifth command line enables the flow with the new rule active.

```
switch246:FID128:admin> flow --create myflow_22 -feature monitor -egrport 21 -srcdev
0x010200 -dstdev 0x011500

switch246:FID128:admin> mapsconfig --import myflow_22

switch246:FID128:admin> logicalgroup --show
-------------------------------------------------------------------
Group Name            |Predefined|Type |Member Count|Members
-------------------------------------------------------------------
ALL_PORTS             |Yes       |Port |8           |3/4,3/6-15
NON_E_F_PORTS         |Yes       |Port |2           |3/4,3/6
ALL_E_PORTS           |Yes       |Port |0           |
ALL_F_PORTS           |Yes       |Port |5           |8/0-1,8/4
ALL_OTHER_F_PORTS     |Yes       |Port |1           |8/0
ALL_HOST_PORTS        |Yes       |Port |1           |8/8
ALL_TARGET_PORTS      |Yes       |Port |0           |
ALL_QUARANTINED_PORTS |Yes       |Port |2           |8/0,8/4
ALL_2K_QSFP           |Yes       |Sfp  |4           |8/28-31
ALL_100M_16GSWL_QSFP  |Yes       |Sfp  |0           |
myflow_22             |No        |Port |3           |Monitored Flow

switch246:FID128:admin> mapsrule --create myRule_22 -group myflow22 -monitor TX_FCNT
-timebase hour -op g -value 22 -action RASLOG -policy myPolicy

switch246:FID128:admin> mapspolicy --enable policy myPolicy22
```

Examples of using MAPS to monitor traffic performance

The following examples illustrate how to use MAPS to monitor traffic performance.

Monitoring end-to-end performance

In the following example, MAPS uses the flow “E2E_flow” to monitor the percentage of frames passing through port 5 between two devices that exceed the configured RX and TX threshold values. To achieve this, it defines a flow using the **-feature monitor** parameter for a particular Source ID, Destination ID, and port.

```
switch246:admin> flow --create E2E_flow -feature monitor -ingrport 5 -scrdev
0x010200 -dstdev 0x020300

switch246:admin> mapsconfig --import E2E_flow

switch246:admin> mapsrule --create E2E_rule -monitor TX_THPUT -group E2E_flow -
timebase min -op g -value 10 -action rasLog -policy flowpolicy

switch246:admin> mapspolicy --enable flowpolicy
```

NOTE
The group name needs to match the imported flow name. In this case “E2E_flow”.

Monitoring frames for a specified set of criteria

In the following example, MAPS uses the flow “abtsflow” to watch for frames in a flow going through port 128 that contain SCSI ABORT sequence markers.

```
switch246:admin> flow --create abtsflow -feature mon -ingrport 128 -frametype abts
```

```
switch246:admin> mapsconfig --import abtsflow
```

You can then define rules for this flow (group), and then re-enable the policy so they take effect. THe following example creates only one rule, “abts_rule”.

```
switch246:admin> mapsrule --create abts_rule -monitor txfcnt -group abtsflow -
timebase min -op ge -value 10 -action raslog -policy flowpolicy
switch246:admin> mapspolicy --enable flowpolicy
```

NOTE
Any new rule you create will not take effect until you enable the policy associated with it.

Examples of monitoring flows at the sub-flow level

The following examples illustrate some of the ways you might want to monitor flows at the sub-flow level.

Excessive throughput notification

To be notified of all the Source ID-Destination ID device pairs for which the RX-TX throughput is greater than a required 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
```

Possible bottleneck notification

To determine which Source ID-Destination ID device pairs have an “abort frame” count that exceeds a specified limit (and consequently cause network bottlenecking), you would import a learning flow that monitors the SCSI Abort notifications with both the Source ID and Destination ID specified as “*” and define a rule to provide the notification, as shown in the following example.

```
switch:admin> flow --create frmcntflow -feature monitor -egrp 234 -srcdev "*" -dstdev
"*" -frametype abts
```

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

```
switch:admin> mapsrule --create frmcntflow_framecnt_1000 -group frmcntflow -timebase
hour -m TX_FCNT -op ge -v 1000 -a RASLOG,EMAIL
```

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.

MAPS dashboard display options

The **mapsdb** command allows you to configure the period of time used for the MAPS dashboard. You can configure the dashboard to display data gathered since midnight, for any 60-minute period since midnight, or for the last seven days on which errors were recorded.

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

NOTE
If the MAPS license is not active, then only the results for unlicensed features will be displayed in the MAPS dashboard. Refer to MAPS commands that do not need a Fabric Vision license on page 22 for a list of these features.

MAPS dashboard sections

The MAPS dashboard output is divided into three main sections: high-level dashboard information, general switch health information, and categorized switch health information. A history section is displayed if you enter **mapsdb --show all**.

Dashboard high-level information section

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

The following output extract shows that the dashboard was started at 1:02 PM on February 27, 2015, the active policy is “dflt_base_policy”, and that there are no fenced, decommissioned, or quarantined ports.

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

1 Dashboard Information:
=======================
DB start time:                Fri Feb 27 13:02:19 2015
Active policy:                dflt_base_policy
```

```
Configured Notifications:        None
Fenced Ports :                   None
Decommissioned Ports :           None
Quarantined Ports :              None
     (output truncated)
```

Switch Health Report section

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

The following output extract shows a sample Switch Health Report section; revealing that the switch status is CRITICAL due to problems with a power supply and a fan.

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

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

     (output truncated)
```

Refer to Switch Policy Status on page 40 for more details on switch policies.

Summary Report section

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

The following categories are monitored by MAPS:

- Port Health on page 32
- Back-end Health on page 34
- FRU Health on page 34
- Security Violations on page 35
- Fabric State Changes on page 35
- Switch Resource on page 36
- Traffic Performance on page 37
- FCIP Health on page 38
- Fabric Performance Impact on page 39

The following output extract shows a sample Summary Report section.

```
3.1 Summary Report:
===================
Category                 |Today                     |Last 7 days               |
--------------------------------------------------------------------------------
Port Health              |No Errors                 |No Errors                 |
BE Port Health           |No Errors                 |No Errors                 |
Fru Health               |In operating range        |In operating range        |
Security Violations      |No Errors                 |No Errors                 |
Fabric State Changes     |No Errors                 |No Errors                 |
Switch Resource          |In operating range        |In operating range        |
Traffic Performance      |In operating range        |In operating range        |
FCIP Health              |Not applicable            |Not applicable            |
Fabric Performance Impact|In operating range        |In operating range        |
     (output truncated)
```

When a category contains an “out-of-range” error, the dashboard displays a table showing the rules triggered in that category since the previous 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 number of times that a rule was triggered in the hour that it was triggered
- The entities (ports, circuits, and others) that triggered the rule
- The values set for these entities when the rule was triggered

For each category, the dashboard stores the five most recent distinct rule violations that occurred in each hour since the previous midnight. For each rule violation the dashboard stores the five most recent entities on which the rules were triggered. Consequently, while a rule might be triggered multiple times within a given hour, only the timestamp of the latest violation is stored, along with the last five entities on which the rule was triggered. However, each violation of a rule individually is reflected in the rule count for that category and the repeat count for that rule in that hour.

For example, if the same rule was triggered 12 times in one hour, the repeat count value (shown as `RptCnt` 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 can include 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, showing that there are six errors affecting four rules. The column headings in the example have been edited slightly so as to allow the example to display clearly.

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

Category(Rule Cnt)|RptCnt|Rule Name                 |Execution Time   |Object  |Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------
Port Health(2)    |1     |defALL_OTHER_F_PORTSCRC_40|07/09/13 17:18:18|Port 1  |876 CRCs              |
                  |1     |defALL_OTHER_F_PORTSCRC_21|07/09/13 17:18:18|Port 1  |876 CRCs              |
Fru Health(2)     |2     |defALL_FANFAN_STATE_FAULTY|07/09/13 19:15:17|Fan 2   |FAULTY                |
                  |      |                          |                 |Fan 1   |FAULTY                |
FCIP Health(2)    |1     |low_tunnel_mon            |11/20/13 06:19:6 |Tunnel  |25                    |

     (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 on page 98.

The following output extract shows a sample History Data section.

```
(output truncated)
4.1 Front end port History Data:
=================================
Stats(Units)        Current   07/21/14  07/14/14    --/--/--    --/--/--   --/--/--   --/--/--
                    Port(val) Port(val) Port(val)
----------------------------------------------------------------------------------------------
CRC(CRCs)           1/13(20)   -          -          -          -          -          -
ITW(ITWs)           -          1/13(612)  -          -          -          -          -
```

```
LOSS_SYNC(SyncLoss) -          -          -          -          -          -          -
LF                  -          -          -          -          -          -          -
LOSS_SIGNAL(LOS)    -          -          -          -          -          -          -
PE(Errors)          -          -          -          -          -          -          -
STATE_CHG           -          -          -          -          -          -          -
C3TXTO(Timeouts)    -          -          -          -          -          -          -
RX(%)               -          -          -          -          -          -          -
TX(%)               -          -          -          -          -          -          -
UTIL(%)             -          -          -          -          -          -          -
BN_SECS(Seconds)    -          -          -          -          -          -          -

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

Notes on dashboard data

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

- 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 dashboard state conditions may be displayed:
  - No Errors: Displayed if there are no errors for the switch ports, security, fabric, or FCIP health; for example, if no port has had an error since midnight.
  - In operating range: Displayed if there are no errors for any of the other categories, 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.
- CIR_UTIL errors (RX, TX, UTIL) are not displayed in the History Data section unless other errors are recorded for that day.
- RX, TX, and UTIL-related rules are monitored under both Traffic Performance and FPI categories.
- The “Rule Count” value is the absolute number of different violations in that category since the previous midnight. The “Repeat Count” is the number of times a rule has been violated in the hour, for example, between 10:00:00 and 10:59:59.
- By default, only the last five violations are displayed for each category. However, entering **mapsdb --show all** causes the dashboard to display all the rule violations currently stored along with additional historical data.

Viewing the MAPS dashboard

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

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

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

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

```
switch:admin> mapsdb --show all
1 Dashboard Information:
=======================
DB start time:                  Mon Mar  2 06:02:07 2015
Active policy:                  my_moderate_policy
Configured Notifications:       RASLOG,SNMP,EMAIL,SDDQ
Fenced Ports :                  None
Decommissioned Ports :          None
Quarantined Ports :             7/28

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


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


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

4 History Data:
===============
Stats(Units)        Current     03/03/15    03/02/15    --/--/--    --/--/--    --/--/--    --/--/--
                    Port(val)   Port(val)   Port(val)
----------------------------------------------------------------------------------------------------
CRC(CRCs)           -           -           -           -           -           -           -
ITW(ITWs)           -           -           -           -           -           -           -
LOSS_SYNC(SyncLoss) -           -           -           -           -           -           -
LF                  -           -           7/13(65)    -           -           -           -
                    -           -           7/14(1)     -           -           -           -
                    -           -           7/15(1)     -           -           -           -
LOSS_SIGNAL(LOS)    7/4(51)     7/4(52)     7/4(44)     -           -           -           -
                    7/5(51)     7/5(52)     7/5(44)     -           -           -           -
                    7/6(51)     7/6(52)     7/6(44)     -           -           -           -
                    7/7(51)     7/7(52)     7/7(44)     -           -           -           -
                    -           -           7/14(1)     -           -           -           -
PE(Errors)          -           -           -           -           -           -           -
STATE_CHG           -           -           7/14(2)     -           -           -           -
                    -           -           7/15(2)     -           -           -           -
                    -           -           7/13(1)     -           -           -           -
LR                  -           -           7/13(55)    -           -           -           -
                    -           -           7/14(3)     -           -           -           -
                    -           -           7/15(3)     -           -           -           -
C3TXTO(Timeouts)    -           -           -           -           -           -           -
RX(%)               2/11(38.46) 7/31(5.87)  2/11(7.39)  -           -           -           -
                    7/31(37.38) -           7/31(6.82)  -           -           -           -
                    7/0(29.11)  -           7/0(4.81)   -           -           -           -
                    7/8(2.79)   -           -           -           -           -           -
TX(%)               2/11(38.46) 2/11(6.20)  2/11(7.68)  -           -           -           -
                    7/31(37.39) 7/0(4.25)   7/31(6.54)  -           -           -           -
                    7/0(29.10)  7/8(1.89)   7/0(5.00)   -           -           -           -
                    7/8(16.74)  -           7/8(4.02)   -           -           -           -
UTIL(%)             2/11(38.46) 2/11(3.38)  2/11(7.53)  -           -           -           -
                    7/31(37.39) 7/31(3.20)  7/31(6.68)  -           -           -           -
                    7/0(29.11)  7/0(2.33)   7/0(4.90)   -           -           -           -
                    7/8(9.77)   7/8(1.01)   7/8(2.21)   -           -           -           -
BN_SECS(Seconds)    -           -           -           -           -           -           -
```

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

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

Viewing a summary switch status report

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

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

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

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

```
switch:admin> mapsdb --show

1 Dashboard Information:
=======================
DB start time:                  Tue Mar  3 22:27:52 2015
Active policy:                  dflt_aggressive_policy
Configured Notifications:       RASLOG
Fenced Ports :                  None
Decommissioned Ports :          None
Quarantined Ports :             None

2 Switch Health Report:
=======================
Current Switch Policy Status: MARGINAL
Contributing Factors:
---------------------
*BAD_PWR (MARGINAL).
*BAD_FAN (MARGINAL).


3.1 Summary Report:
===================
Category                 |Today                     |Last 7 days               |
--------------------------------------------------------------------------------
Port Health              |No Errors                 |No Errors                 |
BE Port Health           |No Errors                 |No Errors                 |
Fru Health               |Out of operating range    |In operating range        |
Security Violations      |No Errors                 |No Errors                 |
Fabric State Changes     |No Errors                 |No Errors                 |
Switch Resource          |In operating range        |In operating range        |
Traffic Performance      |In operating range        |In operating range        |
FCIP Health              |Not applicable            |Not applicable            |
Fabric Performance Impact|In operating range        |In operating range        |


3.2 Rules Affecting Health:
===========================
Category(Rule Count)|RptCnt|Rule Name               |Execution Time   |Object         |Triggered Value|
                                                                                        (Units)       |
-------------------------------------------------------------------------------------------------------
Fru Health(1)       |1     |defALL_PSPS_STATE_FAULTY|03/04/15 00:03:09|Power Supply 2 |FAULTY         |
```

Sub-flow rule violation summaries

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

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

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

The following extract provides an illustration of violations of the “thruputflow_thput_10” rule. The output has been split at \ to allow the example to display clearly.

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

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

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

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

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

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

Viewing a detailed switch status report

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

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

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

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

```
1 Dashboard Information:
=======================
DB start time:                  Wed May 14 23:40:58 2014
Active policy:                  dflt_aggressive_policy
Configured Notifications:       SW_CRITICAL,SW_MARGINAL
Fenced Ports :                  None
Decommissioned Ports :          None
Quarantined  Ports :            None

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

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               |Out of operating range    |In operating range  |
Security Violations      |No Errors                 |No Errors           |
Fabric State Changes     |No Errors                 |No Errors           |
Switch Resource          |In operating range        |In operating range  |
Traffic Performance      |In operating range        |In operating range  |
FCIP Health              |Out of operating range    |In operating range  |
Fabric Performance Impact|Out of operating range    |In operating range  |

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

Category       |Repeat|Rule Name                 |Execution        |Object  |Triggered Value|
(Rule Count)   |Count |                          |Time             |        |(Units)        |
--------------------------------------------------------------------------------------------
Port Health(2) |1     |defALL_OTHER_F_PORTSCRC_40|07/09/13 17:18:18|Port 1  |876 CRCs      |
               |1     |defALL_OTHER_F_PORTSCRC_21|07/09/13 17:18:18|Port 1  |876 CRCs      |
Fru Health(2)  |2     |defALL_FANFAN_STATE_FAULTY|07/09/13 19:15:17|Fan 2   |FAULTY        |
               |      |                          |                 |Fan 1   |FAULTY        |
FCIP Health(2) |1     |low_tunnel_mon            |11/20/13 06:19:6 |Tunnel  |25            |

4.1 Front-end port History Data:
=================================
Stats(Units)        Current   07/21/13   07/14/13   --/--/--   --/--/--   --/--/--   --/--/--
                    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   07/21/13   07/14/13   --/--/--   --/--/--   --/--/--   --/--/--
                    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)         -          -          -          -
```

Viewing historical data

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

NOTE
The output of the **mapsdb --show history** command differs depending on the platform on which you run it. On fixed-port switches, ports are shown in port index format; on chassis-based platforms, ports are shown in slot/port format.

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

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

The following example displays all stored historical port performance data.

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

Front-end port History Data:
=================================
Stats(Units)        Current   07/21/13   07/14/13   --/--/--   --/--/--   --/--/--   --/--/--
                    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)    -          -          -          -          -          -          -

Back-end port History Data:
=================================

Stats(Units)        Current   07/21/13   07/14/13   --/--/--   --/--/--   --/--/--   --/--/--
                    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)         -          -          -          -
```

Viewing data for a specific time window

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

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

1. Connect to the switch and log in using an account with admin permissions.
2. Specify either the day or the hour of the current day you want to view:

   - To specify the day, enter **mapsdb --show details -day** *mm/dd/yyyy*.
   - To specify the hour, enter **mapsdb --show details -hr** *hh*.

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

```
switch:admin> mapsdb --show details -day 11/29/2014

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

DB start time:                  Sat November 29 23:40:58 2014
Active policy:                  dflt_aggressive_policy
Configured Notifications:       SW_CRITICAL,SW_MARGINAL
Fenced Ports :                  None
Decommissioned Ports :          None
Quarantined  Ports :            None

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

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               |Out of operating range    |In operating range        |
Security Violations      |No Errors                 |No Errors                 |
Fabric State Changes     |No Errors                 |No Errors                 |
Switch Resource          |In operating range        |In operating range        |
Traffic Performance      |In operating range        |In operating range        |
FCIP Health              |Out of operating range    |In operating range        |
Fabric Performance Impact|Out of operating range    |In operating range        |

3.2 Rules Affecting Health:
===========================
Category(RuleCnt)|RptCnt|Rule Name                 |Execution Time   |Object   |Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------
Port Health(2)   |1     |defALL_OTHER_F_PORTSCRC_40|11/29/14 17:18:18|F_Port 1 |876 CRCs              |
                 |1     |defALL_OTHER_F_PORTSCRC_21|11/29/14 17:18:18|F_Port 1 |876 CRCs              |
Fru Health(2)    |2     |defALL_FANFAN_STATE_FAULTY|11/29/14 19:15:17|Fan 2    |FAULTY                |
                 |      |                          |                 |Fan 1    |FAULTY                |
FCIP Health(2)   |1     |low_tunnel_mon            |11/20/13 06:19:6 |VE_Port 2|25                    |

4 History Data:
===============
Stats(Units)        Current   11/29/14   --/--/--   --/--/--   --/--/--   --/--/--   --/--/--
                    Port(val)
---------------------------------------------------------------------------------------------
CRC(CRCs)           0(>9999)   0(>9999)   -          -          -          -          -
                    1(876)     1(876)     -          -          -          -          -
ITW(ITWs)           -          -          -          -          -          -          -
LOSS_SYNC(SyncLoss) -          -          -          -          -          -          -
LF                  -          -          -          -          -          -          -
LOSS_SIGNAL(LOS)    -          -          -          -          -          -          -
PE(Errors)          -          -          -          -          -          -          -
STATE_CHG           -          -          -          -          -          -          -
LR                  -          -          -          -          -          -          -
C3TXTO(Timeouts)    -          -          -          -          -          -          -
RX(%)               -          -          -          -          -          -          -
TX(%)               -          -          -          -          -          -          -
UTIL(%)             -          -          -          -          -          -          -
BN_SECS(Seconds)    -          -          -          -          -          -          -
```

```
5 History Data for Back-end ports:
=================================
Stats(Units)        Current   11/29/14   --/--/--   --/--/--   --/--/--   --/--/--   --/--/--
---------------------------------------------------------------------------------------------
CRC(CRCs)             -          -          -          -          -          -          -
ITW(ITWs)             -          -          -          -          -          -          -
LR(Errors)            -          -          -          -          -          -          -
BAD_OS(Errors)        -          -          -          -          -          -          -
FRM_LONG(Errors)      -          -          -          -          -          -          -
FRM_TRUNC(Errors)     -          -          -          -          -          -          -
```

The following example displays historical port performance data for the hour from 2 PM to 4 PM on a chassis-based platform. Note that the History Data section has been trimmed by two days so that the output will display correctly here. Normally there would be two more days of data.

```
switch:admin> mapsdb --show details -hr 2

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

DB start time:            Tue Mar 10 02:00:00 2015
Active policy:            dflt_conservative_policy
Configured Notifications:    RASLOG
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    |No Errors                 |
BE Port Health           |No Errors                 |No Errors                 |
Fru Health               |In operating range        |In operating range        |
Security Violations      |No Errors                 |No Errors                 |
Fabric State Changes     |No Errors                 |No Errors                 |
Switch Resource          |In operating range        |In operating range        |
Traffic Performance      |In operating range        |In operating range        |
FCIP Health              |Not applicable            |Not applicable            |
Fabric Performance Impact|In operating range        |In operating range        |


3.2 Rules Affecting Health:
===========================
Category(Rule Count)|RptCnt|Rule Name            |Execution Time   |Object     |Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------
Port Health(28)     |14    |defALL_E_PORTSITW_80 |03/11/15 02:12:55| 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         |
                    |14    |defALL_E_PORTSITW_41 |03/11/15 02:12:55| 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         |

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

Stats(Units)        Current        03/10/15       03/09/15       03/08/15       03/07/15
                    Port(val)      Port(val)      Port(val)      Port(val)      Port(val)
--------------------------------------------------------------------------------------------
CRC(CRCs)           -              13(169)        41(9)          41(13)         41(8)
                    -              41(9)          -              -              -
ITW(ITWs)           12(1774522342) 12(3464366760) 12(2347926598) 12(1925200209) 13(1955576182)
                    -              13(735674759)  13(1895555442) 13(1034845305) 12(1013464742)
                    -               5(4736)       -              -              -
LOSS_SYNC(SyncLoss) -              13(2)          -              -              -
LF                  5(2)            5(4)          20(3)          -              -
                    20(2)          20(4)          21(3)          -              -
                    21(2)          21(4)           5(2)          -              -
LOSS_SIGNAL(LOS)    5(2)           13(75)          5(4)          -              -
                    20(2)           5(4)          20(4)          -              -
                    21(2)          20(4)          21(4)          -              -
                    -              21(4)          -              -              -
PE(Errors)          -              -              -              -              -
```

```
STATE_CHG           5(4)           13(158)          5(8)          -              -
                    20(4)           5(8)           20(8)          -              -
                    21(4)
```

Clearing MAPS dashboard data

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

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

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

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

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

The following example clears only the dashboard summary data.

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

Additional MAPS Features

Fabric performance monitoring using MAPS

MAPS allows you to monitor your fabrics for performance impacts, including timeouts, latency, and throughput.

There are many distinct elements and layers in a fabric (applications, servers, switches, targets, LUNs, and so on) and consequently multiple places that could possibly be the cause of fabric performance impacts (bottlenecks). As each application’s behavior is unique, the impact of a bottleneck on one individual application might be different from its impact on another application. Each MAPS event needs to be viewed in conjunction with other server or application events to determine the actual root cause of the problem.

The Brocade blades, chassis, and fixed-port switches are also continuously monitored for thermal safety. For more information, refer to “System temperature monitoring” in *Fabric OS Administrator's Guide*.

MAPS latency monitoring

MAPS latency detection is based on data retrieved from the port on the switch (which is just one element in the fabric) and uses this to determine the potential impact to other flows using the fabric. MAPS monitors the fabric impact state of individual F_Ports (but not F_Port trunks) on both individual switches and switches operating in Access Gateway mode. On an Access Gateway set up with F_Port trunks, fabric performance monitoring is done only on those F_Ports actually present on the Access Gateway.

MAPS monitors the current latency on F_Ports over different time windows to determine the impact of latency on the fabric. If it determines the latencies on these ports are severe enough to significantly impact fabric performance, the state of that port is changed to IO_PERF_IMPACT, 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 counter, while IO_FRAME_LOSS is calculated using transient queue latency only.

The following example shows first the MAPS dashboard displaying the IO_PERF_IMPACT report, and then the IO_LATENCY_CLEAR report. The dashboard has been edited to show only Section 3, and the significant text is highlighted in bold. The slash character (\) in the following examples indicates a break inserted because the output is too long to display here as a single line.

```
switch:admin> mapsdb --show
      (output truncated)
3.1 Summary Report:
===================
Category                 |Today                     |Last 7 days               |
--------------------------------------------------------------------------------
Port Health              |Out of operating range    |No Errors                 |
BE Port Health           |No Errors                 |No Errors                 |
Fru Health               |In operating range        |In operating range        |
Security Violations      |No Errors                 |No Errors                 |
Fabric State Changes     |In operating range        |No Errors                 |
Switch Resource          |In operating range        |In operating range        |
Traffic Performance      |In operating range        |In operating range        |
FCIP Health              |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_F_PORTS_IO_PERF_IMPACT|\

\|Execution Time   |Object     |Triggered Value(Units)|
\-----------------------------------------------------
\|08/19/14 22:48:01|F_Port 8/8 |IO_PERF_IMPACT        |
(output truncated)
```

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

```
switch:admin> mapsdb --show
     (output truncated)
3.1 Summary Report:
===================
Category                 |Today                     |Last 7 days               |
--------------------------------------------------------------------------------
Port Health              |Out of operating range    |No Errors                 |
BE Port Health           |No Errors                 |No Errors                 |
Fru Health               |In operating range        |In operating range        |
Security Violations      |No Errors                 |No Errors                 |
Fabric State Changes     |In operating range        |No Errors                 |
Switch Resource          |In operating range        |In operating range        |
Traffic Performance      |In operating range        |In operating range        |
FCIP Health              |Out of operating range    |No Errors                 |
Fabric Performance Impact|In operating range           |In operating range        |

3.2 Rules Affecting Health:
===========================
Category(Rule Count)        |Repeat Count|Rule Name                      |\
--------------------------------------------------------------------------\
Fabric Performance impact(1)|1           |defALL_F_PORTS_IO_LATENCY_CLEAR|\

\|Execution Time   |Object     |Triggered Value(Units)|
\-----------------------------------------------------
\|08/19/14 23:48:01|F_Port 8/8 |IO_LATENCY_CLEAR      |
(output truncated)
```

Frame timeout latency monitoring

MAPS monitors for Class 3 frame timeout errors (C3TXTO) on individual ports and when a timeout is detected on a port, MAPS reports them by setting the port state to IO_FRAME_LOSS and posting a

RASLog message containing the number of frames that have timed out. This state is also reported on the MAPS dashboard.

The following example displays a typical RASlog entry for this condition.

```
 2014/11/30-20:15:59, [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.
```

Transient queue latency counter monitoring

MAPS monitors port transient queue latency (TXQ) for every port to identify ports which may 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, as increased latency on a switch or port can propagate through a switch to the network as a whole. The cumulative effect of many individual device latencies on a port can degrade the performance of the entire port. While the latency for each device may not be a problem, the presence of too many flows with significant latency passing through a port may become a problem.

Brocade Adaptive Networking uses virtual circuits (VCs) to facilitate fabric-wide resource monitoring. The VC architecture provides a flexible way to apply Quality of Service (QoS) monitoring for applications in virtual server environments. If the latency for a VC queue is high, the performance of the traffic flows assigned to that queue will be degraded. While latency is calculated individually for each VC in a port at a rate of once per second, transient latency is monitored only at the VC level, not at the port level. The latency of the VC having the greatest latency time is what is used to determine when an action is triggered. When the latency value crosses the threshold specified in a rule, the configured actions are taken for the virtual circuit for the given port. Possible actions include RASLog, SNMP, or e-mail notifications, as well as port toggling and slow drain device quarantining. Refer to Port toggling support on page 109 and Slow Drain Device quarantining on page 111 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. Now, the TXQ latency calculation is based on the actual latency time seen by a frame, that is, the amount of time it takes for a frame to move out of the switch after it arrives at the time when the sample is taken, and is applicable to all ports.

For determining transient queue latency, MAPS has two predefined threshold states: IO_PERF_IMPACT and IO_FRAME_LOSS. The IO_PERF_IMPACT state is set for a port when latency is between the low threshold and high threshold values, the IO_FRAME_LOSS state is set for a port when latency is greater than the high threshold value.

The following example displays typical RASLogs created when IO_FRAME_LOSS and IO_PERF_IMPACT states are set. The specific issue has been highlighted in bold for the example.

```
 2015/03/19-21:18:00, [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.
2015/03/19-21:18:50, [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 slash character (\) indicates a break inserted because the output is too long to display here as a single line.

```
switch:admin> mapsdb --show
     (output truncated)
3.1 Summary Report:
===================
Category                 |Today                     |Last 7 days              |
--------------------------------------------------------------------------------
Port Health              |Out of operating range    |No Errors                |
BE Port Health           |No Errors                 |No Errors                |
Fru Health               |In operating range        |In operating range       |
Security Violations      |No Errors                 |No Errors                |
Fabric State Changes     |In operating range        |No Errors                |
Switch Resource          |In operating range        |In operating range       |
Traffic Performance      |In operating range        |In operating range       |
FCIP Health              |Out of operating range    |No Errors                |
Fabric Performance Impact|Out of operating range    |In operating range       |

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

Category(Rule Count)        |Repeat Count|Rule Name                     |\
-----------------------------------------------------------------------------\
Fabric Performance impact(1)|1           |defALL_ALL_PORTS_IO_LATENCY_CLEAR|\

\ Execution Time   |Object   |Triggered Value(Units)|
\ ------------------------------------
\ 03/19/15 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 highlighted in boldface are the problem ports.

```
switch:admin> portstatsshow 1
     (output truncated)
stat_mc_tx              0           Multicast frames transmitted
tim_rdy_pri             0           Time R_RDY high priority
tim_txcrd_z             0           Time TX Credit Zero (2.5Us ticks)
tim_txcrd_z_vc  0- 3:   0           0           0           0
tim_txcrd_z_vc  4- 7:   0           0           0           0
tim_txcrd_z_vc  8-11:   0           0           0           0
tim_txcrd_z_vc 12-15:   0           0           0           0
tim_latency_vc  0- 3:   1           1           1           1
tim_latency_vc  4- 7:   1           1           1           1
tim_latency_vc  8-11:   1           1           1           1
tim_latency_vc 12-15:   1           1           1           1
fec_cor_detected        0           Count of blocks that were corrected by FEC
(output truncated)
```

Buffer credit zero counter monitoring

Buffer credit zero counter increments are indirect indication 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 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.

```
2014/11/30-21:10:00, [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.

Latency state clearing

When the detected frame timeout backpressure due to either the buffer cred zero counter value or the queue latency condition is cleared, and 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.

```
 2014/11/30-21:11:00, [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.
```

MAPS and bottleneck detection

Fabric OS bottleneck detection supports the legacy bottleneck monitoring feature responsible for detecting persistent bottlenecks and providing notifications as an alternative to Fabric Performance Impact monitoring.

If you do not use Fabric Performance Impact monitoring in MAPS, you can continue to use the Fabric OS legacy bottleneck monitoring features as no existing bottleneck detection logic or behaviors have been removed.

For the legacy bottleneck monitoring features to produce notifications, the sub-second bottleneck monitoring parameters must be correctly configured and the bottlenecks must be seen persistently on the ports. All bottleneck configurations must be made using **bottleneckmon** commands.

Refer to Bottleneck detection with the MAPS dashboard on page 107 for additional details. Refer to the “Bottleneck Detection” chapter in the *Fabric OS Administrator's Guide* for specific **bottleneckmon** command details and bottleneck monitoring parameters.

Bottleneck detection with the MAPS dashboard

Bottleneck monitoring based on the Fabric OS bottleneck daemon is integrated with the MAPS dashboard, enabling you to easily see which ports are impacted by either persistent or transient bottlenecks.

The MAPS dashboard displays the following latency events:

- Latency bottlenecks on any port
- Timeouts occurring on any 16 Gbps-capable Fibre Channel platform port
- A stuck Virtual Channel on any port
- Congestion bottleneck events (backpressure) on individual F_Ports (but not F_Port trunks)

The MAPS dashboard identifies the ports on which bottlenecks are seen and 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, and includes bottleneck events detected by the bottleneck daemon. However, even if the bottleneck daemon does not log a bottleneck event (due to lack of persistence), the data shown in the “History Data” section displays entries both for those ports that have bottlenecks detected by the daemon and for those ports that have cred_zero counters that are not zero. If the cred_zero counter increases for a port but no bottleneck time is recorded, this indicates a potential transient bottleneck on the port.

In the following example, the last three lines list bottlenecks, with the final bottleneck caused by a timeout rather than a numeric value. Be aware that the column headings in the example have been edited slightly so as to allow the example to display clearly.

```
3. Rules Affecting Health:
===========================
Category(RuleCnt)|RptCnt|Rule Name                    |Execution Time  |Object   |Triggered Value
                                                                                  (Units)             |
-------------------------------------------------------------------------------------------------------
Port Health(12)  |1     |defALL_OTHER_F_PORTSLR_10    |08/21/02 0:30:06|D_Port 23|11                  |
                  1     |defALL_OTHER_F_PORTSLR_5     |08/21/02 0:29:54|D_Port 23|7                   |
                  1     |defALL_OTHER_F_PORTSC3TXTO_3 |08/21/02 0:29:36|D_Port 23|57                  |
                  1     |defALL_OTHER_F_PORTSC3TXTO_10|08/21/02 0:29:36|D_Port 23|57                  |
                  6     |Bottleneck_stuckvc           |08/21/02 0:30:24|D_Port 23|STUCKVC             |
                  1     |Bottleneck_latency           |08/21/02 0:30:20|D_Port 23|60                  |
                  1     |Bottleneck_timeout           |08/21/02 0:30:27|D_Port 23|TIMEOUT             |
     (output truncated)
```

When a latency rule is triggered, the instance is listed as part of the Traffic Performance category. In the both the “Front-end port History Data” section and the “Back-end port History Data” sections, the five ports with the longest total backpressure times since the previous midnight are shown, as shown in the following example. Be aware that the headings in the example have been edited slightly so as to allow the example to display clearly.

```
3. Rules Affecting Health:
=================================
Category(RuleCnt)    |RptCnt|Rule Name                      |Execution Time  |Object   |Trig
Val(Units)|
Fabric Perf Impact(5)|2     |defALL_PORTS_IO_PERF_IMPACT    |08/21/02 0:30:6 |F_Port 13|IO_PERF_IMPACT
                                                             08/21/02 10:30:6|F_Port 22|IO_PERF_IMPACT
                      3     |defALL_PORTS_IO_FRAME_LOSS     |08/21/02 0:30:6 |F_Port 3 |
IO_FRAME_LOSS
                                                             08/21/02 10:30:6|F_Port 2 |IO_FRAME_LOSS
                                                             08/21/02 10:30:6|F_Port 4 |IO_FRAME_LOSS
4.1 Front-end port History Data:
=================================
Stats(Units)        Current   07/21/13   07/14/13   --/--/--   --/--/--   --/--/--   --/--/--
                    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   07/21/13   07/14/13   --/--/--   --/--/--   --/--/--   --/--/--
                    Port(val) Port(val)  Port(val)
---------------------------------------------------------------------------------------------
CRC(CRCs)           2/1/0(15)     -          -          -          -          -          -
LOSS_SYNC(SyncLoss) 2/1/0(1)  3/3/1(2)    3/3/1(2)         -          -          -          -
```

NOTE
The MAPS dashboard will continue to log events whether RASLogs are set to on or off in the bottleneck configuration.

Changes to bottleneck rules from previous versions of Fabric OS

The following bottleneck detection rule changes apply to Fabric OS 7.4.0:

- Transient queue latency is used to determine the Fabric Performance Impact states IO_PERF_IMPACT and IO_FRAME_LOSS.
- The bottleneck detection rules are now applicable to all ports.
- The IO_LATENCY_CLEAR state is monitored. Because of this, the frequency of alerts will change. Refer to Transient queue latency counter monitoring on page 105 and MAPS latency monitoring on page 103 for more details on this state.
- The actions SDDQ, TOGGLE and QT (quiet time) are new. Refer to the **mapsrule** command in MAPS commands altered in this release and the *Fabric OS Command Reference* for more information.

Port toggling support

MAPS supports port toggling, which allows Fabric OS to recover a port that has been bottlenecked by a target device. While there are many reasons why the target device could be bottlenecked, one of the most common is a temporary glitch in an adapter or its software.

Port toggling in MAPS temporarily disables a port and then re-enables it, allowing the port to reset and recover from the issue. If the port does not recover, Fabric OS suspends the port, forcing the port traffic to switch over to a different path if one is available.

To enable recovering ports using port toggling, MAPS assumes that there is a redundant path to the target device. It does not check to see if there is one, nor can it check to see if traffic to or from the target device has been switched over to a redundant path. MAPS also assumes that while the port is being toggled, the operational state of the port will not be changed by any other mechanism, such as an administrator disabling or moving the port, or a port fencing operation.

NOTE
Port toggling cannot be used in conjunction with automatic VC quarantining, as this may 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_crc_rule -group ALL_PORTS -monitor CRC -
timebase MIN -op g -value 0 -action TOGGLE -tt 180
```

When a port has been toggled by a MAPS rule, TOGGLE appears 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:                 Thu Aug 21 14:58:06 2014
Active policy:                 nil5
Configured Notifications:      RASLOG,TOGGLE <== Emphasis added
Fenced Ports :                 None
Decommissioned Ports :         None
Quarantined Ports :            None
       (output truncated)
```

When MAPS toggles a port, the **switchshow** command lists the reason for the port being disabled as “Transient”.

The following example displays a sample output for **switchshow** when a port has been toggled. In this example, Port 65 is listed as “Disabled (Transient)”.

```
switch:admin> switchshow
444  8 28   ------   cu   8G    No_Sync   FC
445  8 29   ------   cu   8G    No_Sync   FC
446  8 30   ------   cu   8G    No_Sync   FC
447  8 31   ------   cu   8G    No_Sync   FC
 64  9  0   014000   id   N2    Online    FC  F-Port  10:00:00:00:00:01:00:01
 65  9  1   014100   id   N4    In_Sync   FC  Disabled (Transient) <== Emphasis added
 66  9  2   014200   --   N8    No_Module FC
 67  9  3   014300   --   N8    No_Module FC
 68  9  4   014400   --   N8    No_Module FC
```

Enabling MAPS Fabric Performance Impact monitoring

MAPS Fabric Performance Impact (FPI) is automatically enabled for new Brocade switches already running Fabric OS 7.3.0 or later, or if the legacy bottleneck monitoring feature was not enabled before the switch firmware was upgraded to Fabric OS 7.3.x or 7.4.0.

NOTE
If you want to use FPI monitoring, the legacy bottleneck monitoring feature cannot be enabled.

If FPI has been disabled for some reason, complete the following steps to re-enable it.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **mapsconfig --enableFPImon**. This will return one of two messages:
   - If the legacy bottleneck monitoring feature *is not* enabled, you will see the following message:
     `Error: MAPS Fabric Performance Impact monitoring is enabled`
   - If the legacy bottleneck monitoring feature *is* enabled, you will see the following message:
     `Operation failed. Bottleneck monitoring is enabled. Please disable bottleneck monitoring and retry.`

     Enter **bottleneckmon --disable** to disable the legacy bottleneck monitoring feature, and then enter **mapsconfig --enableFPImon** again to enable MAPS Fabric Performance Impact monitoring.

For more information on the relationship of MAPS FPI monitoring and the legacy bottleneck monitoring feature, refer to MAPS Fabric Performance Impact monitoring and legacy bottleneck monitoring on page 111.

MAPS Fabric Performance Impact monitoring and legacy bottleneck monitoring

The following conditions apply to MAPS Fabric Performance Impact (FPI) monitoring and legacy bottleneck monitoring:

- MAPS FPI monitoring and the legacy bottleneck monitoring feature are mutually exclusive. If the legacy bottleneck monitoring feature is not enabled on a switch, MAPS will automatically start monitoring for impacts on fabric performance when the switch is restarted after upgrading to Fabric OS 7.3.0. However, if the legacy bottleneck monitoring feature was enabled before the upgrade, MAPS will not monitor for impacts on fabric performance. To use MAPS FPI monitoring, you will need to both explicitly disable the legacy bottleneck monitoring feature and enable MAPS FPI monitoring.
- Once the MAPS FPI monitoring feature is enabled, the legacy bottleneck monitoring feature cannot be enabled.
- You cannot migrate the legacy bottleneck monitoring configurations to MAPS FPI monitoring. Once MAPS FPI monitoring is enabled, all monitoring will be done using the predefined MAPS thresholds.
- Legacy bottleneck monitoring must be disabled on every logical switch before you can proceed with configuring the netmon application using the **mapsconfig** command.
- Disabling legacy bottleneck monitoring only disables the latency and congestion features, it does not disable stuck VC monitoring. This means that a RASLog message stating “Severe latency detected” is still posted if that condition occurs.

Disabling MAPS Fabric Performance Impact monitoring

MAPS allows you to disable MAPS Fabric Performance Impact monitoring in order to enable legacy bottleneck monitoring.

To disable MAPS Fabric Performance Impact monitoring, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **mapsconfig --disableFPImon**.
   MAPS Fabric Performance Impact monitoring is automatically disabled.

NOTE
Legacy bottleneck monitoring is not automatically enabled if you disable MAPS Fabric Performance Impact monitoring.

Slow Drain Device quarantining

In a fabric, many flows share the same link or virtual circuit (VC). However, the credits used to send traffic or packets across the link are common to all the flows using the same link. Because of this, a slow-draining device may slow down the return of credits and have a negative effect on the healthy flows through the link.

The following figure illustrates how this functions. In it, the inability of the slow-draining device (1) to clear frames quickly enough is causing backpressure not just to the edge device (2), but also to the core (3) and the appliance (4) that is trying to reach the slow-draining device.

FIGURE 1 Slow-drain data flow

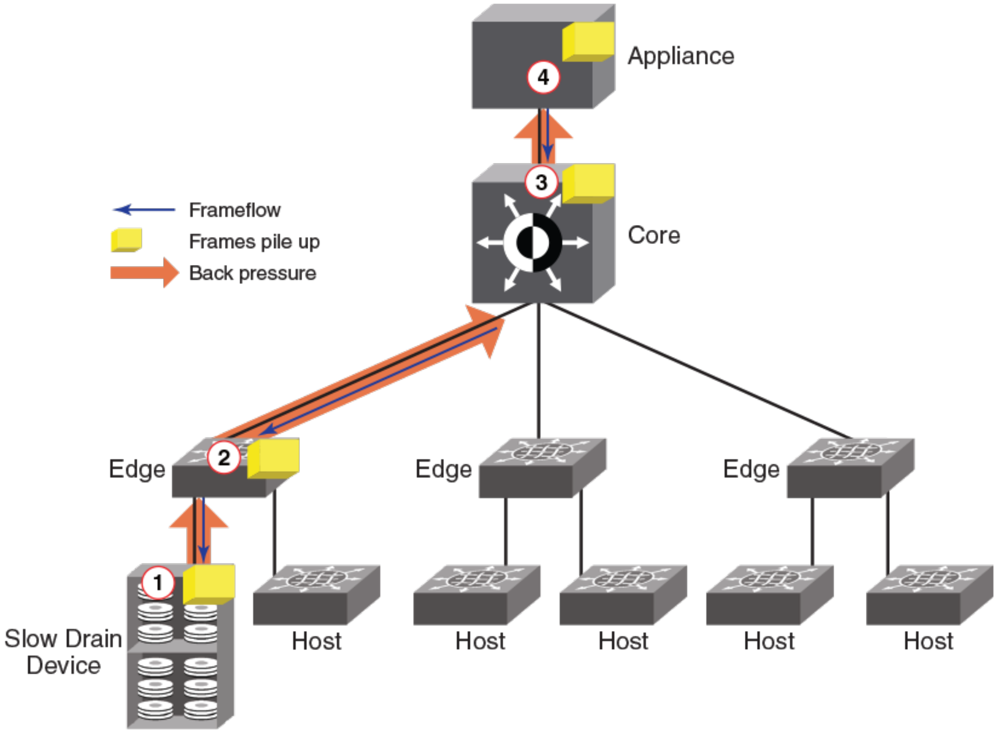

To remedy this, Brocade created Slow Drain Device Quarantine (SDDQ), a feature which —in conjunction with Quality of Service (QoS) monitoring— allows MAPS to identify a slow-draining device and quarantine it by automatically moving all traffic destined to the F_Port that is connected to the slow-draining device to a low-priority VC so that the traffic in the original VC does not experience backpressure.

MAPS accomplishes the isolation of the slow-draining device by moving the slow-draining port to the MAPS ALL_QUARANTINED_PORTS group, and then notifying the Name Server to use a low-priority virtual circuit for all flows directed to the slow-draining port. Once the traffic targeted to the port is isolated, MAPS moves the port to the quarantine group. After this is done, MAPS continues to monitor the device latency on the port. If MAPS clears the latency impact state by changing the status to IO_LATENCY_CLEAR, it creates a RASLog or other notice for that change, but the port is kept in the quarantine group. If any quarantined ports go offline or become disabled, the ports remain in the quarantined group. When the quarantined ports come online, MAPS automatically notifies the Name Server that the port is still quarantined.

To restore the traffic targeted to a quarantined port to normal QoS priority, the **sddquarantine --clear** command is the only way to move the port out of the quarantine group. Refer to Clearing quarantined ports on page 117 for more details on this procedure.

Slow Drain Device Quarantine licensing

For Slow Drain Device Quarantining (SDDQ) to take effect, the Fabric Vision license must be installed on the switch where the slow draining device is detected, as well as on the switch where the quarantine action is to occur. Intermediate switches do not need the Fabric Vision license for this feature to work, but they must have QoS enabled on all ISLs.

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

If the Fabric Vision license is not installed on a switch which has slow-draining devices (either remote or local) attached, and no quarantine action has been applied previously to those devices, simply enabling the Fabric Vision license on that switch will not immediately trigger the quarantining action. However, the quarantine action will be applied to all zoned devices coming online after the Fabric Vision license has been enabled. In order to have the same behavior across all Fabric Vision-licensed switches, Brocade recommends that you toggle the switch, devices, or ports in the zone involving the slow-draining device after you enable the Fabric Vision license.

When MAPS is installed and the license is active, SDDQ is available to all switches in the fabric that can receive the SDDQ registered state change notification (RSCN). It does not matter if the quarantine action is explicitly enabled on that switch or not. However, those switches that do not have the SDDQ action enabled will not be able to flag or isolate any slow-draining devices.

Notes on Slow Drain Device Quarantining

The following items should be kept in mind when using Slow Drain Device Quarantining (SDDQ):

- SDDQ is disabled by default on installation.
- SDDQ is supported:
  - For N_Port ID Virtualization (NPIV) devices connected to F_Ports, but not for ports connected to Access Gateway devices.
  - On all Fabric OS platforms that are running Fabric OS 7.4.0, including FICON environments.
- SDDQ is not supported:
  - On switches running Fabric OS versions earlier than 7.4.0.
  - On switches running in Access Gateway mode.
  - On F_Port trunks.
  - In Fibre Channel Routing (FCR) backbone fabrics, but can be supported on devices within an FCR edge fabric. However, the edge fabric will not apply this quarantine feature to the flows for any FCR-imported devices.
- The SDDQ action is supported by the default rules defALL_IO_PERF_IMPACT and defALL_IO_FRAME_LOSS.
- While MAPS does allow you to enable and disable SDDQ for individual logical switches, Brocade recommends that you enable SDDQ for the entire fabric. The reason for this is that it is difficult to predict when a device may become slow-draining.
- You cannot add or delete members from the ALL_QUARANTINED_PORTS group. No default rules are defined for the group, but you can create custom rules to monitor the ports in this group.
- When a device is marked as being slow-draining, only the flows destined to it are shifted to the low-priority Virtual Circuit (VC). Flows in the reverse direction are not affected.
- QoS zoning rules may also route flows to low-priority virtual channels, which may further slow these flows. In addition, if a device is already in a QoS zone and the device has been identified as slow-draining, then any flows to that device will be assigned a low priority independent of the QoS zone priority until you explicitly restore the device to its original VC. Alternatively, if a QoS zone is newly configured with the slow-draining device and pushed across the fabric, the flow will not shift to the new priority because it is still marked as a slow-draining flow. It will be kept in the low-priority VC until you remove it from that VC.

- If device latency is due to Class 3 timeouts (C3TXTO) and the active C3TXTO rule has port fencing as an action, then the port fencing (disabling) may be performed first and quarantining will not occur.
- If there are switches in the fabric that do not have QoS feature support, then end-to-end slow drain flow isolation may not be possible.
- The maximum number of devices that can be isolated per unit (chassis or fixed-port switch) is 32. The default value is 8. This is controlled by the chassis-wide “Chassis SDDQ limit” user-configurable key, which is set using the **configurechassis** command. Refer to the *Fabric OS Command Reference* for more information on this command.
- To prevent SDDQ from consuming an excessive amount of system resources in FICON environments and large-scale fabrics, the SDDQ action is blocked in the following scenarios to prevent spurious IOCTLs (input/output control actions) due to quarantine action on hundreds of source ports:
  - If the total number of ports zoned to the slow drain port is set to 32.
  - If the defzone is labeled as “all access”, and there is no user-defined zoning configuration.

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

Enabling Slow Drain Device Quarantining

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

To enable SDDQ, complete the following steps.

1. Connect to the device and log in using an account with admin permissions.
2. Use the **mapsconfig --actions** command, and include SDDQ as one of the actions.

The following example enables SDDQ as an available action.

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

Disabling Slow Drain Device Quarantining

To disable the Slow Drain Device Quarantine feature, exclude SDDQ from the configured list of MAPS actions, as shown in the following procedure.

1. Connect to the device and log in using an account with admin permissions.
2. Use the **mapsconfig --actions** command, and do not include SDDQ as one of the actions.

The following example removes Slow Drain Device Quarantining from the list of available actions.

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

Confirming the slow-draining status of a device

You can use the **nsshow**, **nscamshow**, and **nodefind** commands to verify that a device is slow-draining.

The following example shows the output of the **nsshow** command where there is a slow-draining device. The identifying line has been called out in this example. If there not a slow-draining device, the line would not appear.

```
switch:admin> nsshow
{
Type Pid    COS     PortName                NodeName                 TTL(sec)
N    010000;    2,3;20:00:00:05:1e:92:e8:00;20:00:00:05:1e:92:e8:00; na
    Fabric Port Name: 20:00:00:05:1e:92:e8:00
    Permanent Port Name: 20:00:00:05:1e:92:e8:00
    Port Index: 0
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
    Port Properties: SIM Port
N    014a00;    2,3;30:0a:00:05:1e:84:b5:c3;10:00:00:05:1e:84:b5:c3; na
    FC4s: FCP
    NodeSymb: [42] "dsim:fdmi_host:sw_5300edsim[172.26.26.188]"
    Fabric Port Name: 20:4a:00:05:1e:92:e8:00
    Permanent Port Name: 30:0a:00:05:1e:84:b5:c3
    Port Index: 74
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
N    015000;      3;10:00:00:00:00:01:00:01;10:00:00:00:00:00:01:01; na
    Fabric Port Name: 20:50:00:05:1e:92:e8:00
    Permanent Port Name: 10:00:00:00:00:01:00:01
    Port Index: 80
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
    Slow Drain Device: Yes     <== Slow-draining device identified
N    015100;      3;10:00:00:00:00:02:00:01;10:00:00:00:00:00:02:01; na
    Fabric Port Name: 20:51:00:05:1e:92:e8:00
    Permanent Port Name: 10:00:00:00:00:02:00:01
    Port Index: 81
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
The Local Name Server has 4 entries }
```

The following example shows the output of the **nscamshow** command where there is a slow-draining device. The identifying line has been called out in this example. If there was not a slow-draining device, the line would not appear.

```
switch:admin> nscamshow
nscam show for remote switches:
Switch entry for 2
  state  rev    owner  cap_available
  known   v740  0xfffc01 1
  Device list: count 7
    Type Pid    COS     PortName                NodeName
    N    020a00;      3;10:00:00:00:00:04:00:01;10:00:00:00:00:00:04:01;
        Fabric Port Name: 20:0a:00:05:1e:84:98:c7
        Permanent Port Name: 10:00:00:00:00:04:00:01
        Port Index: 10
        Share Area: No
        Device Shared in Other AD: No
        Redirect: No
        Partial: No
    N    020b00;      3;10:00:00:00:00:0a:00:01;10:00:00:00:00:00:0a:01;
        Fabric Port Name: 20:0b:00:05:1e:84:98:c7
        Permanent Port Name: 10:00:00:00:00:0a:00:01
        Port Index: 11
        Share Area: No
        Device Shared in Other AD: No
        Redirect: No
        Partial: No
        Slow Drain Device: Yes     <== Slow-draining device identified
    N    021000;      3;10:00:00:00:00:05:00:01;10:00:00:00:00:00:05:01;
        Fabric Port Name: 20:10:00:05:1e:84:98:c7
        Permanent Port Name: 10:00:00:00:00:05:00:01
        Port Index: 16
        Share Area: No
        Device Shared in Other AD: No
        Redirect: No
        Partial: No
```

The following example shows the output of the **nodefind** command where there is a slow-draining device. The identifying line has been called out in this example. If there was not a slow-draining device, the line would not appear.

```
switch:admin> nodefind 015000
Local:
 Type Pid COS PortName NodeName SCR
 N 015000; 3;10:00:00:00:00:01:00:01;10:00:00:00:00:00:01:01; 0x00000003
    Fabric Port Name: 20:50:00:05:1e:92:e8:00
    Permanent Port Name: 10:00:00:00:00:01:00:01
    Device type: Physical Unknown(initiator/target)
    Port Index: 80
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
    Slow Drain Device: Yes     <== Slow-draining device identified
    Aliases:
```

Displaying quarantined ports

MAPS allows you to display a list of ports which have been quarantined in the ALL_QUARANTINED_GROUP.

To display a list of ports which have been quarantined by MAPS, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **sddquarantine --show** to view a list of the quarantined ports.

The following example shows the offline quarantined local ports and the online quarantined device information across the fabric.

```
switch:admin> sddquarantine --show
----------------------------------------------------------
Ports marked as Slow Drain Quarantined in the Local Switch:   2/1, 1/3
----------------------------------------------------------
Ports marked as Slow Drain Quarantined but not enforced:   1/3
----------------------------------------------------------

Online Quarantined Devices across the fabric
----------------------------------------------------------
 Port Index |   PID   |            PWWN
----------------------------------------------------------
      17    | 051100 | 30:10:00:05:33:ac:c6:13
      17    | 051101 | 30:10:01:05:33:ac:c6:13
----------------------------------------------------------
```

Clearing quarantined ports

MAPS allows you to remove ports which have been quarantined from the quarantine group (ALL_QUARANTINED_GROUP).

To remove ports which have been quarantined by MAPS from the quarantine group, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **sddquarantine --show** to view a list of the quarantined ports.
3. Enter **sddquarantine --clear** followed by either:
   - The slot/port ID that is to be removed from the quarantine group
   - The keyword **all** to clear all quarantined ports

   If the port is offline, the port is removed from the ALL_QUARANTINED_GROUP. There is no confirmation of this action.

   A port is not allowed to be cleared from quarantine while the latency or frame loss condition that caused it to be quarantined persists, as the port will not be flagged again until the condition is cleared and then retriggered. However, you can override this by using the **-force** keyword as part of the **sddquarantine** command.

NOTE
The "Ports marked as Slow Drain Quarantined but not enforced" line only appears in the output when devices could not be quarantined because the chassis SDDQ limit of 32 quarantined devices has been reached.

The first part of the following example uses the **sddquarantine --show** command to display the offline quarantined local ports and the online quarantined device information across the fabric. The second part shows using the **sddquarantine --clear** command to remove the ports from quarantine. Notice that port 3 was not able to be cleared using the **all** option as it was still in either the IO_FRAME_LOSS or the IO_PERF_IMPACT condition, but this port was able to be cleared using the **-force** option. These two options can also be used together.

```
switch:admin> sddquarantine --show
----------------------------------------------------------
Ports marked as Slow Drain Quarantined in the Local Switch:   2/1, 1/3
----------------------------------------------------------
Ports marked as Slow Drain Quarantined but not enforced:   1/3
----------------------------------------------------------

Online Quarantined Devices across the fabric
----------------------------------------------------------
 Port Index |   PID   |            PWWN
----------------------------------------------------------
      17    | 051100 | 30:10:00:05:33:ac:c6:13
      17    | 051101 | 30:10:01:05:33:ac:c6:13
----------------------------------------------------------

switch:admin> sddquarantine --clear 2/1, 1/3
Initiated clearing port from quarantined state

switch:admin> sddquarantine --clear 1/3
Clear failed since port is still in latency state

switch:admin> sddquarantine --clear all
The clear action was not initiated for the following port(s). Try with individual
ports
                3
Initiated quarantine action on other ports

switch:admin> sddquarantine --clear 1/3 -force
Initiated clearing port from quarantined state
```

Slow Drain Device quarantining and FICON

Brocade does not recommend that you enable the Slow Drain Device Quarantine (SDDQ) feature in a FICON environment because a FICON environment typically has a single zone for all devices. However, if you are using 1-to-1 zoning, SDDQ may help with slow-draining device issues.

When SDDQ is enabled, all traffic to a slow-draining device is moved to the lowest-priority virtual circuit. In a single-zone environment this impacts all traffic in the zone.

In a FICON environment, the maximum number of devices that can be isolated per unit (chassis or fixed-port switch) is 32, and the default number of devices is 10. The first time a port is detected as a slow drain, the Slow Drain Device quarantine will be applied to it. The maximum number of ports that can be zoned with a slow-draining port is 32. If the 32-port zone limit is exceeded, the quarantine action is not taken, and any new zoned device or (33rd) port coming online after this limit is reached will not be quarantined, even if it is slow-draining. However, MAPS will add the problem port to the quarantine list. This means that once the number of zoned ports is back below the limit, the next time the listed port comes online the listed port may be quarantined.

If there are quarantined slow-draining ports in a fabric and their zone changes, these ports will remain quarantined even if this violates the limit on the number of ports that can be zoned together. Once the limit is reached, any reduction in the number of ports in the zone will not affect the quarantine decision until the next time MAPS detects the slow drain condition. Be aware that if the slow drain conditions persist, you may not get another alert, as MAPS has already raised an alert. If a slow-draining device port avoids quarantine due to a limit on the number of ports in the zone, any zoned device connecting to that port that comes online after the limit check will not be quarantined.

The **sddquarantine --show** command output displays the list of successfully quarantined ports as well as those that were identified as problematic but were not quarantined, so that you can take any necessary measures to recover these ports. For example, you might choose to disable the ports, or (after correcting what is causing the slow drain) restore the ports. The restoration action for quarantined ports does not have any constraints and is not affected by any zone restrictions.

NOTE
To avoid excessive consumption of system resources, Slow Drain Device quarantining based on IOCTL errors is not permitted.

Scalability limit monitoring

MAPS monitors fabric level changes such as the count of logged-in devices in a pure Layer 2 (L2) fabric, the count of LSAN devices, zone configuration size and the number of Fiber Channel Router configurations. These fabric-level monitoring systems all have scalability limits, which MAPS can monitor and send alerts using RASLog entries, SNMP messages, or e-mail. The monitoring results are captured in the MAPS dashboard under the “Fabric State Changes” category. While 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 imported Logical SAN (LSAN) device connections (this includes both edge fabric and Backbone fabric device connections).

NOTE
MAPS does not monitor the individual device counts for edge and Backbone fabrics.

When a rule is triggered, the corresponding RASLogs appear in the summary section of the dashboard. The following example shows two rules (LSan_Dev_Count and L2_Dev_Count) have been triggered. The column headings in the example have been edited slightly so as to allow the example to display clearly.

```
3.1 Summary Report:
===================
Category                 |Today                     |Last 7 days               |
--------------------------------------------------------------------------------
Port Health              |No Errors                 |No Errors                 |
BE Port Health           |No Errors                 |No Errors                 |
Fru Health               |In operating range        |In operating range        |
Security Violations      |No Errors                 |No Errors                 |
Fabric State Changes     |No Errors                 |No Errors                 |
Switch Resource          |In operating range        |In operating range        |
Traffic Performance      |In operating range        |In operating range        |
FCIP Health              |Not applicable            |Not applicable            |
Fabric Performance Impact|In operating range        |In operating range        |

3.2 Rules Affecting Health:
===========================
Category(Rule Count)   |RptCnt|Rule Name     |Execution Time  |Object   |Triggered Value (Units)|
-------------------------------------------------------------------------------------------------
Fabric State Changes(2)|1     |LSan_Dev_Count|08/21/02 00:30:6|D_Port 23|12 %
|
                       |1     |L2_Dev_Count  |08/21/13 01:04:6|D_Port 23|12 %                   |
```

For more detailed information on scalability limits, refer to *Brocade SAN Scalability Guidelines: Brocade Fabric OS v7.X*.

Layer 2 fabric device connection monitoring

A pure Layer 2 fabric is a collection of Fibre Channel switches and devices and switches that doesn't participate in a metaSAN. In such a fabric, rules for device counts are calculated as a percentage of the total number of devices. For example, a Layer 2 fabric with 5500 devices logged in is using 92 percent of the maximum limit of 6000 devices for a Layer 2 fabric. So if user have configured a rule to trigger an alert at 90 percent or greater, then MAPS triggers the action configured for that rule and sends the data to the dashboard.

Imported 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 imported from all edge fabrics) in a metaSAN can be monitored for a scalability limit.

NOTE
MAPS rules for monitoring imported LSAN device connections in a metaSAN can be configured only on switches that are a part of the Backbone fabric.

Device counts in this framework are calculated as a percentage of the total number of LSAN devices in a metaSAN (including imported devices from all edge fabric). For example: if a fabric has four switches in the Backbone fabric and four switches each in four edge fabrics, the total number of LSAN devices in this metaSAN (including imported devices from all edge fabrics) is 1200. Given a maximum of 10000 devices, this is 12 percent. If you have configured a rule to trigger at 10 percent or greater, then MAPS triggers the action configured for the rule, but only on those switches that are part of the Backbone fabric, and caches the data in the dashboard.

Backbone fabric Fibre Channel router count monitoring

In a Backbone fabric, there can be maximum number of 12 Fibre Channel routers (FCRs). MAPS rules can be configured to monitor the number of Fibre Channel routers in the Backbone fabric as an absolute value. If the number of Fibre Channel routers reaches the configured threshold, MAPS triggers the action configured for the rule and caches the data in the dashboard. Refer to Default rules for scalability limit monitoring on page 123 for these values.

The following example shows a typical RASLog entry for exceeding the threshold for the number of Fibre Channel routers in the Backbone fabric:

```
2014/05/27-17:02:00, [MAPS-1003], 14816, SLOT 4 | FID 20, WARNING, switch_20, Switch,
Condition=SWITCH(BB_FCR_CNT>12), Current Value:[BB_FCR_CNT,13], RuleName=
defSWITCHBB_FCR_CNT_12, Dashboard Category=Fabric State Changes.
```

The following example shows a typical MAPS dashboard entry for exceeding the threshold for the number of Fibre Channel routers in the Backbone fabric. Be aware that the column headings in section 3.2 of the example have been edited slightly so as to allow the example to display clearly.

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

3.2 Rules Affecting Health:
===========================
Category(RuleCount) |RptCount|Rule Name             |Execution Time   |Object|Triggered Value(Units)|
----------------------------------------------------------------------------------------------------
Fabric State Changes|1       |defSWITCHBB_FCR_CNT_12|05/27/14 17:02:00|Switch|13                    |
(1)                 |        |                      |                 |      |                      |
```

Zone configuration size monitoring

In Fabric OS 7.3.0 and later, MAPS can monitor zone configuration size. Based on the platform, a switch supports either a maximum zone configuration size of 1 MB or 2 MB. The monitoring value is calculated as a percentage of the zone configuration space used. If the configuration size reaches the configured threshold limit, MAPS triggers the action configured for the rule and caches the data in the dashboard. Refer to Default rules for scalability limit monitoring on page 123 for these limit values.

NOTE
MAPS zone configuration size monitoring is only for the default switch, as the total memory size is for the chassis as a whole. The maximum available zone configuration limit is determined at the chassis level and shared by all logical switches.

Monitoring NPIV logins to F_Ports

MAPS can monitor the number of N_Port ID Virtualization (NPIV) logins to individual F_Ports and generates RASLog, SNMP, or e-mail alerts if the login threshold for a port is reached.

NPIV is a Fibre Channel mechanism that allows host virtual machines to obtain a virtual world wide node (WWN) name. This allows multiple virtual machines to share the same physical Fibre Channel host bus adapter (HBA). When a switch gets a new connection request as part of the FLOGI, a unique N_Port ID is assigned to the device. This means that for each device there is a one-to-one mapping between the virtual world wide node name and the N_Port ID.

Monitoring NPIV logins to F_Ports is important because Access Gateway leverages NPIV to present open system Fibre Channel host connections as logical devices to SAN fabrics. This reduces switch-to-switch interoperability problems by allowing servers to seamlessly connect to SAN fabrics. However, as there is a limit to the number of logins that can be made to a switch, it is important to monitor this number and be able to alert administrators when the limit is approached.

NOTE
NPIV monitoring is not High Availability (HA)-capable. As a consequence, if there is a reboot or an HA failover, existing NPIV logins are not preserved, and new ones are assigned on a first-come, first-served basis.

MAPS supports monitoring all F_Ports in a switch for the number of NPIV logins as part of scalability limit monitoring. The value monitored is calculated as a percentage of the number of devices logged in relative to the maximum number of logins configured for that port. When the number of devices logged in to the switch reaches the rule threshold, MAPS posts a RASLog, SNMP, or e-mail message, allowing you to block any further logins. This threshold monitoring helps keep the switch from being overloaded with connection requests. For each port, MAPS allows you to configure the maximum number of logins value that is used in the calculation.

When a rule is triggered, the triggered rule name appears in the dashboard summary section as a Port Health item. The following example shows the relevant portion of the **mapsdb --show** command output for a switch configured with a maximum NPIV login limit of 100 for all ports. Be aware that the column headings in the example have been edited slightly so as to allow the example to display clearly.

```
     (output truncated)
Category                 Today               Last 7 Calendar Days
==============================================================================
Port Health         :  Out of range             In range
BE Port Health      :  In range                 In range
FRU Health          :  In range                 In range
FCIP Health         :  In range                 In range
Security Violations :  No Errors                   In range
Fabric Health       :  In range                    In range
Switch Resources    :  In range                    In range
Traffic Performance :  In range                    In range

Rules Affecting Health:
=================================
Category(RuleCount)|RptCnt|Rule Name                   |Execution Time    |Object     |Triggered
Value|
                                                                                        (Units)
|
-------------------------------------------------------------------------------------------------------
|
Port Health(2)     |1     |defALL_F_PORTSDEV_NPIV_LOGIN |01/14/15 12:59:58 |F-Port 7/11| 65.00 %
|
                   |1     |defALL_OTHER_F_PORTSLOSS_SYNC|01/14/15 12:52:34 |F-Port 7/1 | 1 SyncLoss
|
     (output truncated)
```

MAPS includes the following default rules for monitoring NPIV logins to F_Ports in the default policies. The counter monitored is the absolute value in percentage and the timebase is “NONE”.

TABLE 24 Default rules for monitoring NPIV logins to F_Ports

Policy	Rule Name	Condition	Actions
dflt_aggressive_policy	defALL_NPIV_LOGIN_PER_60	ALL_F_PORTS(DEV_NPIV_LOGINS/NONE > 60)	SNMP, RASLOG, EMAIL
dflt_moderate_policy	defALL_NPIV_LOGIN_PER_75	ALL_F_PORTS(DEV_NPIV_LOGINS/NONE > 75)	SNMP, RASLOG, EMAIL
dflt_conservative_policy	defALL_NPIV_LOGIN_PER_90	ALL_F_PORTS(DEV_NPIV_LOGINS/NONE > 90)	SNMP, RASLOG, EMAIL

Scalability limit monitoring assumptions and dependencies

The following assumptions and dependencies should be kept in mind when considering scalability limit monitoring.

- All the scalability limits are soft limits, not hard limits; the monitored value can be greater than 100 percent.
- The Backbone fabric can also have Layer 2 switches; these switches are not considered as part of any of the scalability limit metrics.
- The number of device connections in an edge fabric or Backbone fabric also have scalability limits themselves, and these cannot be monitored using MAPS.
- Scalability limit monitoring (using L2_DEVCNT_PER) occurs only at midnight. Therefor, if a switch is moved from being a part of the Layer 2 fabric to being a part of the edge fabric, the device count metrics (how many devices in the fabric) will not change until the next midnight.
- The "LSAN-imported device" metric is only monitored in switches that are a part of a Backbone fabric.
- Scalability limits that are determined internally by a device cannot be monitored by MAPS.

Default rules for scalability limit monitoring

The following table lists the scalability monitoring default rules in each of the default policies, and shows the actions and condition for each rule.

TABLE 25 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_12	L2_DEVCNT_PER greater than 90 LSAN_DEVCNT_PER greater than 90 ZONE_CFGSZ_PER greater than 90 BB_FCR_CNT greater than 12	RASLOG, SNMP, EMAIL
dflt_moderate_policy	defSWITCHL2_DEVCNT_PER_75 defSWITCHLSAN_DEVCNT_PER_75 defSWITCHZONE_CFGSZ_PER_80 defSWITCHBB_FCR_CNT_12	L2_DEVCNT_PER greater than 75 LSAN_DEVCNT_PER greater than 75 ZONE_CFGSZ_PER greater than 80 BB_FCR_CNT greater than 12	RASLOG, SNMP, EMAIL

TABLE 25 Scalability monitoring default rules (Continued)

Policy name	Rule name	Rule condition	Rule action
dflt_aggressive_policy	defSWITCHL2_DEVCNT_PER_60 defSWITCHLSAN_DEVCNT_PER_60 defSWITCHZONE_CFGSZ_PER_70 defSWITCHBB_FCR_CNT_12	L2_DEVCNT_PER greater than 60 LSAN_DEVCNT_PER greater than 60 ZONE_CFGSZ_PER greater than 70 BB_FCR_CNT greater than 12	RASLOG, SNMP, EMAIL

Examples of scalability limit rules

The following examples show the patterns for creating device counts, Fibre Channel router counts, and zone configuration usage rules for MAPS.

Rule for device counts in a Layer 2 fabric

In the following example, when the total device count in all switches that are part of the Layer 2 fabric rises above 90 percent of the total permissible count in the fabric, MAPS reports the threshold violation using a RASLog message.

```
switch123:FID128:admin> mapsrule --create L2_Dev_Count -group SWITCH -monitor
L2_DEVCNT_PER -timebase none -op ge -value 90 -action RASLOG -policy
scalability_policy
```

Rule for LSAN device counts

In the following example, when the total device count in all switches that are part of the metaSAN (edge plus Backbone) fabric rises above 90 percent of the total permissible count in the fabric, MAPS reports the threshold violation using a RASLog message on that platform.

```
switch123:FID128:admin> mapsrule --create LSan_Dev_Count -group SWITCH -monitor
LSAN_DEVCNT_PER -timebase none -op ge -value 90 -action RASLOG -policy
scalability_policy
```

Rule for Fibre Channel router count in Backbone fabric

In the following example, when the maximum limit of 12 Fibre Channel routers in the Backbone fabric is reached, MAPS reports the threshold violation using a RASLog message.

```
switch123:admin> mapsrule --create FCRCount -group SWITCH -monitor BB_FCR_CNT -
timebase none -op ge -value 12 -action RASLOG -policy scalability_policy
```

Rule for zone configuration size

In the following example, when the zone configuration size limit reaches 90 percent of the total size, MAPS reports the threshold violation using a RASLog message.

```
switch123:admin> mapsrule --create ZoneConfigSize -group SWITCH -monitor
ZONE_CFGSZ_PER -timebase none -op ge -value 90 -action RASLOG -policy
scalability_policy
```

For ZONE_CFGSZ_PER policy, the default time base is "none" and the system performs a daily check. The policy does not support other time base.

MAPS Service Availability Module

The MAPS Service Availability Module (MAPSSAM) reports display CPU, RAM, and flash memory usage, and the port status for every physical and virtual Fibre Channel port on the switch.

There are three MAPSSAM commands: **mapsSam --show**, **mapsSam --clear**, and **mapsSam --help**. Only **mapsSam --show** has additional parameters. These parameters are listed in the following table and illustrated in the following examples.

TABLE 26 MAPSSAM command options

Option	Details
--show (default option)	For each physical and virtual Fibre Channel port on a switch, this displays the total up, down, and offline time (as percentages), and the number of times the port has been down. This enables you to see if a particular port is failing more often than the others. **NOTE** The report does not distinguish why a port is recorded as down, it only reports how long the port has been down.
--show cpu	Displays the CPU usage as a percentage.
--show memory	Displays the general RAM memory usage as a percentage, along with total, used, and free memory values.
--show flash	Displays the flash memory usage as a percentage.

The following examples show the from using the various **mapsSam --show** parameters.

Using only “--show”

In this form, the report lists the following information for each port:

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

NOTE
The MAPSSAM report does not include the health status of gigabyte Ethernet (GbE) ports.

- Total up time — Percentage of time the port was up.
- Total down time — Percentage of time the port was faulty.
- Down occurrence — Number of times the port was faulty.
- Total Offline time — Percentage of time the port was offline.
- Number of ports

All percentages are based on the total time the switch was up or down since the switch was rebooted, MAPS was activated, or the **mapsSam --clear** command was last run.

The following example shows typical output for **mapsSam --show**.

```
switch0:FID128: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%
```

FCIP QoS monitoring using MAPS

FCIP Quality of Service (QoS) monitoring uses policies based on a combination of data characteristics and delivery requirements to appropriately prioritize data traffic. For example, while ordinary data traffic is tolerant of delays and dropped packets, real-time voice and video data are not. MAPS QoS policies provide a framework for accommodating these differences in traffic packets as they pass through a network, and can help you investigate issues such as packet loss, excessive bandwidth utilization, and so on.

MAPS can monitor FCIP QoS on three Brocade devices: the Brocade 7800, the Brocade 7840, and the Brocade FX8-24. On the Brocade 7800 and FX8-24, MAPS can monitor circuit level QoS, packet loss, and throughput, as well as circuit-level jitter and round-trip duration. On the Brocade 7840, in addition to the circuit-level monitoring available on the Brocade 7800 and Brocade FX8-24 MAPS can monitor tunnel-level QoS and throughput.

QoS monitoring using MAPS uses the predefined QoS priorities of High, Medium, Low, and F-class. You can configure the values used by these priorities at both the FCIP tunnel and circuit level. The attributes monitored by MAPS for QoS at circuit level are throughput and packet loss. Throughput is the percentage of QoS circuit or tunnel utilization in a configured time period (hour, minute, or day); packet loss is the percentage of the total number of packets that have had to be retransmitted. MAPS monitors the state changes and throughput of each individual tunnel using these QoS priorities. QoS monitoring is not High Availability (HA)-capable.

For tunnel-level monitoring, MAPS can monitor the predefined groups ALL_TUNNELS, ALL_TUNNEL_HIGH_QOS, ALL_TUNNEL_MED_QOS, ALL_TUNNEL_LOW_QOS, and ALL_TUNNEL_F_QOS. These groups correspond to the FCIP tunnels.

For circuit-level monitoring, MAPS can monitor the predefined groups ALL_CIRCUIT_HIGH_QOS, ALL_CIRCUIT_MED_QOS, ALL_CIRCUIT_LOW_QOS, ALL_CIRCUIT_F_QOS, as well as the ALL_CIRCUITS group for round-trip time (RTT) and connection variance (Jitter) in addition to the CIR_STATE, CIR_UTIL, and CIR_PKTLOSS parameters. Monitoring the RTT and Jitter values helps to alert you to possible network disruptions and congestion in the network.

The available statistics are broken out in the following table, and rules corresponding to these statistics are in the default policies.

TABLE 27 Brocade FCIP monitoring parameters and groups

Parameter	Groups where the parameter is used as a metric
State change (STATE_CHG)	**Brocade 7840 only:** ALL_TUNNELS , ALL_TUNNEL_HIGH_QOS, ALL_TUNNEL_MED_QOS, ALL_TUNNEL_LOW_QOS , ALL_TUNNEL_F_QOS
Percent utilization (UTIL)	ALL_CIRCUIT_HIGH_QOS, ALL_CIRCUIT_MED_QOS, ALL_CIRCUIT_LOW_QOS, ALL_CIRCUIT_F_QOS **Brocade 7840 only:** ALL_TUNNELS, ALL_TUNNEL_HIGH_QOS, ALL_TUNNEL_MED_QOS, ALL_TUNNEL_LOW_QOS, and ALL_TUNNEL_F_QOS

TABLE 27 Brocade FCIP monitoring parameters and groups (Continued)

Parameter	Groups where the parameter is used as a metric
Percentage of packets lost in transmission (PKTLOSS)	ALL_CIRCUIT_HIGH_QOS, ALL_CIRCUIT_MED_QOS, ALL_CIRCUIT_LOW_QOS, ALL_CIRCUIT_F_QOS
Round-trip time in milliseconds (RTT)	ALL_CIRCUITS
Variance in RTT in milliseconds (Jitter)	ALL_CIRCUITS
CIR_STATE, CIR_UTIL, and CIR_PKTLOSS	Refer to FCIP Health on page 38 for descriptions of these parameters.

The following table breaks the parameters out by supported platform.

TABLE 28 Brocade platform support for FCIP monitoring parameters

Platform	Supported parameters
All FCIP platforms: • Brocade 7800 • Brocade 7840 • Brocade FX8-24	• CIR_UTIL • CIR_STATE • CIR_PKTLOSS • RTT • JITTER • PKTLOSS (Circuit QoS) • UTIL (Circuit QoS)
Brocade 7840 only	All of the above parameters, plus the following: • PKTLOSS (Tunnel and Tunnel QoS) • UTIL (Tunnel and Tunnel QoS) • STATE_CHG (Tunnel)

Quality of Service monitoring example

The following MAPS rule states that when the packet loss percentage for ALL_CIRCUIT_HIGH_QOS group members becomes greater than or equal to 0.5 in a given minute, a RASLog entry will be posted.

```
switch:admin> mapsrule --create urule -group ALL_CIRCUIT_HIGH_QOS -monitor PKTLOSS -t min -op ge -
value .5 -action raslog
```

On triggering the rules, the corresponding RASLogs will appear under the summary section of the dashboard. In the following example, there is one RASLog, triggered by the rule “low_tunnel_mon”. This rule has the format `-group ALL_TUNNEL_LOW_QOS -monitor PKTLOSS -timebase HOUR -op ge -value 30 -action raslogs`. The “Conditions contributing to health” column headings have been edited so as to allow the example to display clearly.

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

Conditions contributing to health:
=================================
Category(RuleCnt)|RptCnt|Rule Name                            |Execution Time |Trigger Val(Units)|
---------------------------------------------------------------------------------------
FCIP Health (1)  |1     |defALL_CIRCUIT_HIGH_QOS_UTIL_75      |8/11/14 06:19:6|Circuit Qos 23/0  |
FCIP Health (1)  |1     |defALL_CIRCUIT_HIGH_QOS_PKTLOSS_PER_1|8/11/14 07:02:5|Circuit Qos 23/0  |
```

Updating monitoring policies for devices with 4 PSUs

If you have a chassis that supports more than 2 power supply units (PSUs), when you increase the number of PSUs and want to enable the Call Home feature to be activated for all the PSUs, you must change the active MAPS power supply “switchstatus” policy settings. This feature requires that you have a Brocade Direct Support maintenance contract.

NOTE
This procedure applies only to the following switches:

- Brocade DCX
- Brocade DCX 8510-8
- Brocade DCX 8510XT-8

To change the active MAPS power supply “switchstatus” policy settings, complete the following steps.

1. Display the MAPS policy rules using one of the following commands:
   - **mapspolicy --show -all**
   - **mapsrule --show fw_CHASSISBAD_PWRCrit_3**

The following examples show the results of each of these commands when the policy is set for two PSUs.

```
switch:admin> mapspolicy --show -all
fw_CHASSISBAD_PWRCrit_3 SW_CRITICAL
CHASSIS(BAD_PWR/none>=3)
.
.
.
Active Policy is 'fw_active_policy'.
```

```
switch:admin> mapsrule --show fw_CHASSISBAD_PWRCrit_3
Rule Data:
 ----------
RuleName: fw_CHASSISBAD_PWRCrit_3
Condition: CHASSIS(BAD_PWR/none>=3)
Actions: SW_CRITICAL
Associated Policies: fw_active_policy
```

2. Verify that the chassis has the condition BAD_PWR/none>=3 (**CHASSIS(BAD_PWR/none>=3)**).
3. Record the Active Policy name. This name is required to complete the following step.
4. Use the **mapsrule --config fw_CHASSISBAD_PWRCrit_3 -policy** *policy_name from Step 3* **-monitor BAD_PWR -group CHASSIS -timebase none -op ge -value 1 -action SW_CRITICAL** command to change the CHASSIS(BAD_PWR/none>=3) setting to CHASSIS(BAD_PWR/none>=1) . When you change a policy, you must enter all the values for the policy, even if you are changing only one value.
   In this example, the **-policy** name is **fw_active_policy** that 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
   ```

   Now if you have a Brocade Direct Support maintenance contract, all four PSUs can use the Call Home event notification capability to automatically send an e-mail or dial into a support center to report system problems.

MAPS Threshold Values

Viewing monitoring thresholds

You can use the CLI to view the thresholds for a policy, or for a group within a policy.

To view monitoring thresholds, complete the following steps.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter **mapspolicy --show** *policy_name*. To see only the thresholds for a specific group in a policy, use **--show** *policy_name* | **grep** *group_name*.

   The following example shows all the thresholds for the ALL_D_PORTS group in the policy named "dflt_conservative_policy".

```
switch:admin> mapspolicy --show dflt_conservative_policy | grep ALL_D_PORTS
defALL_D_PORTSCRC_3            RASLOG,SNMP,EMAIL   ALL_D_PORTS(CRC/MIN>3)
defALL_D_PORTSPE_3             RASLOG,SNMP,EMAIL   ALL_D_PORTS(PE/MIN>3)
defALL_D_PORTSITW_3            RASLOG,SNMP,EMAIL   ALL_D_PORTS(ITW/MIN>3)
defALL_D_PORTSLF_3             RASLOG,SNMP,EMAIL   ALL_D_PORTS(LF/MIN>3)
defALL_D_PORTSLOSS_SYNC_3      RASLOG,SNMP,EMAIL   ALL_D_PORTS(LOSS_SYNC/MIN>3)
defALL_D_PORTSCRC_H90          RASLOG,SNMP,EMAIL   ALL_D_PORTS(CRC/HOUR>90)
defALL_D_PORTSPE_H90           RASLOG,SNMP,EMAIL   ALL_D_PORTS(PE/HOUR>90)
defALL_D_PORTSITW_H90          RASLOG,SNMP,EMAIL   ALL_D_PORTS(ITW/HOUR>90)
defALL_D_PORTSLF_H90           RASLOG,SNMP,EMAIL   ALL_D_PORTS(LF/HOUR>90)
defALL_D_PORTSLOSS_SYNC_H90    RASLOG,SNMP,EMAIL   ALL_D_PORTS(LOSS_SYNC/HOUR>90)
defALL_D_PORTSCRC_D1500        RASLOG,SNMP,EMAIL   ALL_D_PORTS(CRC/DAY>1500)
defALL_D_PORTSPE_D1500         RASLOG,SNMP,EMAIL   ALL_D_PORTS(PE/DAY>1500)
defALL_D_PORTSITW_D1500        RASLOG,SNMP,EMAIL   ALL_D_PORTS(ITW/DAY>1500)
defALL_D_PORTSLF_D1500         RASLOG,SNMP,EMAIL   ALL_D_PORTS(LF/DAY>1500)
defALL_D_PORTSLOSS_SYNC_D1500  RASLOG,SNMP,EMAIL   ALL_D_PORTS(LOSS_SYNC/DAY>1500)
```

   The first column is the name of the statistic being monitored. The second is the actions for that statistic that will be triggered if the threshold is passed. The third column lists the group being monitored, followed by the metric, followed by the threshold. This means that "defALL_D_PORTSCRC_3 RASLOG,SNMP,EMAIL ALL_D_PORTS(CRC/MIN>3)" :

   - Is named "defALL_D_PORTSCRC_3"
   - Has the actions RASLog, SNMP, and e-mail

- Applies to all D_Ports
- Measures CRC errors per minute. The threshold to trigger the listed actions is “more than three errors in a minute”.

Back-end port monitoring thresholds

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

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

TABLE 29 Back-end port monitoring default thresholds and actions

Errors	Thresholds	Actions
CRC	• 10 per 5 minutes • 100 per day	RASLOG, SNMP, EMAIL, FMS
Link Reset	• 10 per 5 minutes • 100 per day	RASLOG, SNMP, EMAIL, FMS
ITW	• 10 per 5 minutes • 100 per day	RASLOG, SNMP, EMAIL, FMS
BAD_OS	• 10 per 5 minutes • 100 per day	RASLOG, SNMP, EMAIL, FMS
Frame too long	• 10 per 5 minutes • 100 per day	RASLOG, SNMP, EMAIL, FMS
Frame truncated	• 10 per 5 minutes • 100 per day	RASLOG, SNMP, EMAIL, FMS

Fabric state change monitoring thresholds

All the fabric state change monitors support the Minute, Hour, and Day timebases. They do not support the “None” timebase.

The following table lists the default monitoring thresholds for fabric state change criteria used by MAPS. All thresholds are measured per minute and actions are triggered when the reported value is greater than the threshold value.

TABLE 30 Default fabric state change monitoring thresholds and actions

	Fabric state change monitoring threshold values by policy			
Monitoring statistic	Aggressive	Moderate	Conservative	Actions
Domain ID change	1	1	1	RASLOG, SNMP, EMAIL

TABLE 30 Default fabric state change monitoring thresholds and actions (Continued)

	Fabric state change monitoring threshold values by policy			
Monitoring statistic	**Aggressive**	**Moderate**	**Conservative**	**Actions**
Fabric logins	4	6	8	RASLOG, SNMP, EMAIL
Fabric reconfigurations	1	2	4	RASLOG, SNMP, EMAIL
E_Ports down	1	2	4	RASLOG, SNMP, EMAIL
Segmentation changes	1	2	4	RASLOG, SNMP, EMAIL
Zone changes	2	5	10	RASLOG, SNMP, EMAIL
L2 Device Count	60	75	90	RASLOG, SNMP, EMAIL
LSAN Device Count	60	75	90	RASLOG, SNMP, EMAIL
Zone Configuration size	70	80	90	RASLOG, SNMP, EMAIL
FCR Count	12	12	12	RASLOG, SNMP, EMAIL

FCIP monitoring thresholds

These FCIP monitors support Minute, Hour, Day, and Week timebases: Tunnel state change, Tunnel throughput, Tunnel QoS throughput, Tunnel QoS Packet loss, FCIP Circuit State Changes, FCIP Circuit Utilization, FCIP Packet loss, FCIP Circuit Round Trip Time, and FCIP connection variance.

The following tables list the default monitoring thresholds for Fiber Channel over IP (FCIP) criteria used by MAPS. All actions are triggered when the reported value is greater than the threshold value.

TABLE 31 Default FCIP monitoring thresholds and actions

		FCIP monitoring threshold values by policy			
Monitoring statistic	**Units**	**Aggressive**	**Moderate**	**Conservative**	**Actions**
Circuit state change (CIR_STATE)	Changes per minute	0	3	5	RASLOG, SNMP, EMAIL, FENCE
Circuit utilization percentage (CIR_UTIL)	Percentage per hour	50	75	90	RASLOG, SNMP, EMAIL
Circuit packet loss percentage (CIR_PKTLOSS)	Percentage per minute	0.01	0.05	0.1	RASLOG, SNMP, EMAIL
Circuit round-trip times (RTT)	Total delay in milliseconds	250	250	250	RASLOG, SNMP, EMAIL

TABLE 31 Default FCIP monitoring thresholds and actions (Continued)

		FCIP monitoring threshold values by policy			
Monitoring statistic	**Units**	**Aggressive**	**Moderate**	**Conservative**	**Actions**
Circuit jitter (JITTER)	Percentage of delay, calculated from the difference of two successive minutes. The total delay in milliseconds is averaged per minute for each minute. Values less than 5 ms in the converted percentage are ignored.	5	15	20	RASLOG, SNMP, EMAIL
Tunnel (STATE_CHG)	Changes per minute	0	1	3	RASLOG, SNMP, EMAIL
Tunnel or circuit QoS (UTIL)	Percentage per hour	60	75	90	RASLOG, SNMP, EMAIL
QoS Packet loss percentage (PKTLOSS)	Percentage per minute	0.01	0.05	0.1	RASLOG, SNMP, EMAIL

FRU state monitoring thresholds

For all FRU monitoring statistics, the default MAPS thresholds are part of blade and WWN rules for Brocade DCX and Brocade DCX+ systems. All threshold conditions are absolute, and actions are triggered when the statistic value either does or does not match the value (depending on how the rule is written). FRU monitoring statistics do not use any timebases.

TABLE 32 FRU monitoring statistics, states, and actions

Monitored Statistic	Supported States	Actions
Power Supply (PS_STATE)	ON, OUT, FAULTY	RASLOG, SNMP, EMAIL
Fan (FAN_STATE)	ON, OUT, FAULTY	RASLOG, SNMP, EMAIL
Slot (BLADE_STATE)	ON, OFF, OUT, FAULTY	RASLOG, SNMP, EMAIL
SFP (SFP_STATE)	IN, OUT, FAULTY	RASLOG, SNMP, EMAIL
WWN (WWN)	ON, OUT, FAULTY	RASLOG, SNMP, EMAIL

Port Health monitoring thresholds

All Port Health monitoring thresholds used by MAPS are triggered when they exceed the listed value. For thresholds that have both an upper value and a lower value, the threshold is triggered when it exceeds the upper value or drops below the lower value. All thresholds other than RXP, TXP, and Utilization percentage are measured per minute. The RXP, TXP, and Utilization percentage thresholds are measured per hour.

The following Port Health monitors support Minute, Hour, and Day timebases: CRC Errors, Invalid Transmit Words, Loss of sync, Link Failure, Loss of Signal, Protocol Errors, Link Reset, C3 Time outs, and State change. The SFP Current, SFP Receive Power, SFP Transmit Power, SFP Voltage, SFP Temperature, and SFP Power On Hours monitors support only the “None” timebase.

D_Port default Port Health monitoring thresholds

The following tables list the default D_Port Port Health monitoring threshold values and actions, broken out by policy.

TABLE 33 Aggressive policy default D_Port Port Health monitoring threshold values and actions

Monitoring statistic	Unit	Threshold	Actions
CRC Errors (defALL_D_PORTSCRC_1)	Min	1	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_1)	Min	1	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_1)	Min	1	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_1)	Min	1	EMAIL, SNMP, RASLOG
CRC Errors (defALL_D_PORTSCRC_H30)	Hour	30	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_H30)	Hour	30	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_H30)	Hour	30	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_H30)	Hour	30	EMAIL, SNMP, RASLOG
CRC Errors (defALL_D_PORTSCRC_D500)	Day	500	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_D500)	Day	500	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_D500)	Day	500	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_D500)	Day	500	EMAIL, SNMP, RASLOG

TABLE 34 Moderate policy default D_Port Port Health monitoring threshold values and actions

Monitoring statistic	Unit	Threshold	Actions
CRC Errors (defALL_D_PORTSCRC_2)	Min	2	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_2)	Min	2	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_2)	Min	2	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_2)	Min	2	EMAIL, SNMP, RASLOG
CRC Errors (defALL_D_PORTSCRC_H60)	Hour	60	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_H60)	Hour	60	EMAIL, SNMP, RASLOG

TABLE 34 Moderate policy default D_Port Port Health monitoring threshold values and actions (Continued)

Monitoring statistic	Unit	Threshold	Actions
Link Failure (defALL_D_PORTSLF_H60)	Hour	60	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_H60)	Hour	60	EMAIL, SNMP, RASLOG
CRC Errors (defALL_D_PORTSCRC_D1000)	Day	1000	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_D1000)	Day	1000	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_D1000)	Day	1000	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_D1000)	Day	1000	EMAIL, SNMP, RASLOG

TABLE 35 Conservative policy default D_Port Port Health monitoring threshold values and actions

Monitoring statistic	Unit	Threshold	Actions
CRC Errors (defALL_D_PORTSCRC_3)	Min	3	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_3)	Min	3	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_3)	Min	3	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_3)	Min	3	EMAIL, SNMP, RASLOG
CRC Errors (defALL_D_PORTSCRC_H90)	Hour	90	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_H90)	Hour	90	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_H90)	Hour	90	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_H90)	Hour	90	EMAIL, SNMP, RASLOG
CRC Errors (defALL_D_PORTSCRC_D1500)	Day	1500	EMAIL, SNMP, RASLOG
Invalid Transmit Words (defALL_D_PORTSITW_D1500)	Day	1500	EMAIL, SNMP, RASLOG
Link Failure (defALL_D_PORTSLF_D1500)	Day	1500	EMAIL, SNMP, RASLOG
Sync Loss (defALL_D_PORTSLOSS_SYNC_D1500)	Day	1500	EMAIL, SNMP, RASLOG

E_Port default Port Health monitoring thresholds

The following table lists the default E_Port Port Health monitoring threshold values and actions.

TABLE 36 Default E_Port Port Health monitoring threshold values and actions

	E_Port monitoring threshold values by policy			
Monitoring statistic	**Aggressive**	**Moderate**	**Conservative**	**Actions**
C3 Time out (C3TX_TO)	5	10	20	EMAIL, SNMP, RASLOG
CRC Errors (CRC)	Low: 0 High: 2	Low: 10 High: 20	Low: 21 High: 40	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Invalid Transmit Words (ITW)	Low: 15 High: 20	Low: 21 High: 40	Low: 41 High: 80	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Link Reset (LR)	Low: 2 High: 4	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
State Change (STATE_CHG)	Low: 2 High: 4	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Loss of signal (LOSS_SIGNAL)	0	3	5	EMAIL, SNMP, RASLOG
Link Failure (LF)	0	3	5	EMAIL, SNMP, RASLOG
Sync Loss (LOSS_SYNC)	0	3	5	EMAIL, SNMP, RASLOG

F_Port default Port Health monitoring thresholds

The following table lists the default Host F_Port Port Health monitoring threshold values and actions.

TABLE 37 Default Host F_Port Port Health monitoring threshold values and actions

Monitoring statistic	**Host F_Port monitoring threshold values by policy**			**Actions**
	Aggressive	**Moderate**	**Conservative**	
C3 Time out (C3TX_TO)	Low: 2 High: 4	Low: 3 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLOG, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS
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
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
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

TABLE 37 Default Host F_Port Port Health monitoring threshold values and actions (Continued)

Monitoring statistic	Host F_Port monitoring threshold values by policy			Actions
	Aggressive	Moderate	Conservative	
State Change (STATE_CHG)	Low: 2 High: 4	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Loss of signal (LOSS_SIGNAL)	0	3	5	EMAIL, SNMP, RASLOG
Link Failure (LF)	0	3	5	EMAIL, SNMP, RASLOG
Sync Loss (LOSS_SYNC)	0	3	5	EMAIL, SNMP, RASLOG

The following table lists the default Target F_Port Port Health monitoring threshold values and actions.

TABLE 38 Default Target F_Port Port Health monitoring threshold values and actions

	Target F_Port monitoring threshold values by policy			
Monitoring statistic	Aggressive	Moderate	Conservative	Actions
C3 Time out (C3TX_TO)	Low: 0 High: 2	Low: 3 High: 5	Low: 6 High: 10	Low: EMAIL, SNMP, RASLOG, FMS High: EMAIL, SNMP, FENCE, DECOM, FMS
CRC Errors (CRC)	Low: 0 High: 2	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Invalid Transmit Words (ITW)	Low: 5 High: 10	Low: 11 High: 20	Low: 21 High: 40	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Link Reset (LR)	Low: 0 High: 2	Low: 3 High: 5	Low: 6 High: 10	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
State Change (STATE_CHG)	Low: 0 High: 2	Low: 3 High: 7	Low: 8 High: 15	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Loss of signal (LOSS_SIGNAL)	0	3	5	EMAIL, SNMP, RASLOG
Link Failure (LF)	0	3	5	EMAIL, SNMP, RASLOG
Sync Loss (LOSS_SYNC)	0	3	5	EMAIL, SNMP, RASLOG

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

Non-F_Port default Port Health monitoring thresholds

The following table lists the default non-F_Port Port Health monitoring threshold values and actions.

TABLE 39 Default non-F_Port Port Health monitoring threshold values and actions

	Non-F_Port monitoring threshold values by policy			
Monitoring statistic	**Aggressive**	**Moderate**	**Conservative**	**Actions**
C3 Time out (C3TX_TO)	N/A	N/A	N/A	N/A
CRC Errors (CRC)	Low: 0 High: 2	Low: 10 High: 20	Low: 21 High: 40	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Invalid Transmit Words (ITW)	Low: 15 High: 20	Low: 21 High: 40	Low: 41 High: 80	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Link Reset (LR)	Low: 2 High: 4	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
State Change (STATE_CHG)	Low: 2 High: 4	Low: 5 High: 10	Low: 11 High: 20	Low: EMAIL, SNMP, RASLOG High: EMAIL, SNMP, FENCE, DECOM
Loss of signal (LOSS_SIGNAL)	0	3	5	EMAIL, SNMP, RASLOG
Link Failure (LF)	0	3	5	EMAIL, SNMP, RASLOG
Sync Loss (LOSS_SYNC)	0	3	5	EMAIL, SNMP, RASLOG

Resource monitoring thresholds

The only timebase the Resource monitors support is “None”.

The following table lists the default monitoring threshold values and associated actions for the switch resource criteria monitored by MAPS. All thresholds are measured per minute and are triggered when they are greater than the shown value.

TABLE 40 Default resource monitoring thresholds and actions

	Resource monitoring threshold values by policy			
Monitoring statistic	**Aggressive**	**Moderate**	**Conservative**	**Actions**
Flash memory percentage used (FLASH_USAGE)	90	90	90	RASLOG, SNMP, EMAIL
CPU percentage used (CPU)	80	80	80	RASLOG, SNMP, EMAIL

TABLE 40 Default resource monitoring thresholds and actions (Continued)

	Resource monitoring threshold values by policy			
Monitoring statistic	Aggressive	Moderate	Conservative	Actions
Memory percentage used (MEMORY_USAGE)	75	75	75	RASLOG, SNMP, EMAIL
Ethernet management port state (ETH_MGMT_PORT_STATE)	Up/Down	Up/Down	Up/Down	RASLOG, SNMP, EMAIL
Temperature Sensor (TEMP)	OUT_OF_RANGE	OUT_OF_RANGE	OUT_OF_RANGE	RASLOG, SNMP, EMAIL

Security monitoring thresholds

All the Security Health monitors support the Minute, Hour, and Day timebases. They do not support the “None” timebase.

The following table lists the default monitoring thresholds for security criteria used by MAPS. Unless noted otherwise, all thresholds are measured per minute and actions are triggered when the reported value is greater than the threshold value.

TABLE 41 Default security monitoring thresholds and actions

	Security monitoring threshold values by policy			
Monitoring statistic	Aggressive	Moderate	Conservative	Actions
DCC violations	0	2	4	RASLOG, SNMP, EMAIL
HTTP violation	0	2	4	RASLOG, SNMP, EMAIL
Illegal command	0	2	4	RASLOG, SNMP, EMAIL
Incompatible security DB	0	2	4	RASLOG, SNMP, EMAIL
Login violations	0	2	4	RASLOG, SNMP, EMAIL
Invalid certifications	0	2	4	RASLOG, SNMP, EMAIL
No-FCS	0	2	4	RASLOG, SNMP, EMAIL
SCC violations	0	2	4	RASLOG, SNMP, EMAIL
SLAP failures	0	2	4	RASLOG, SNMP, EMAIL
Telnet violations	0	2	4	RASLOG, SNMP, EMAIL
TS out of sync	1 per hour 2 per day	2 per hour 4 per day	4 per hour 10 per day	RASLOG, SNMP, EMAIL

SFP monitoring thresholds

SFP monitoring statistics do not use any timebases.

All SFP monitoring thresholds used by MAPS are triggered when the reported value exceeds the threshold value. For thresholds with both an upper value and a lower value, actions are triggered when the reported value exceeds the upper threshold value or drops below the lower threshold value.

10 Gbps and 16 Gbps SFP monitoring threshold defaults

The following table lists the default thresholds for 10 Gbps and 16 Gbps SFPs.

TABLE 42 Default SFP monitoring thresholds and actions for 10 Gbps and 16 Gbps SFPs

	10 Gbps and 16 Gbps SFP monitoring thresholds for all policies				
Monitoring statistic	**10 Gbps SWL**	**10 Gbps LWL**	**16 Gbps SWL**	**16 Gbps LWL**	**Actions**
Current (CURRENT) (mA)	10	95	12	70	SFP_MARGINAL, RASLOG, SNMP, EMAIL
Receive Power (RXP) (µW)	1999	2230	1259	1995	SFP_MARGINAL, RASLOG, SNMP, EMAIL
Transmit Power (TXP) (µW)	1999	2230	1259	1995	SFP_MARGINAL, RASLOG, SNMP, EMAIL
Voltage (VOLTAGE) (mV)	3000 to 3600	2970 to 3600	3000 to 3600	3000 to 3600	SFP_MARGINAL, RASLOG, SNMP, EMAIL
Temperature (TEMP) (°C)	-5 to 90	-5 to 90	-5 to 85	-5 to 90	SFP_MARGINAL, RASLOG, SNMP, EMAIL

Quad SFPs and all other SFP monitoring threshold defaults

The following table lists the default threshold and actions for Quad SFPs and all other SFPs that are neither 10 Gbps nor 16 Gbps.

TABLE 43 Default SFP monitoring thresholds and actions for QSFPs and all other SFPs

	QSFP and other SFP monitoring thresholds for all policies		
Monitoring statistic	**All QSFPs**	**All Other SFPs**	**Actions**
Current (CURRENT) (mA)	10	50	RASLOG, SNMP, EMAIL
Receive Power (RXP) (µW)	2180	5000	RASLOG, SNMP, EMAIL
Transmit Power (TXP) (µW)	-	5000	RASLOG, SNMP, EMAIL
Voltage (VOLTAGE) (mV)	2940 to 3600	2960 to 3630	RASLOG, SNMP, EMAIL
Temperature (TEMP) (°C)	-5 to 85	-13 to 85	RASLOG, SNMP, EMAIL

Fabric Performance Impact thresholds

The following FPI monitors support the Minute, Hour, and Day timebases: Receive Bandwidth usage percentage, Transmit Bandwidth usage percentage, and Trunk Utilization percentage. The Fabric Performance Impact monitor (DEV_LATENCY_IMPACT) supports only the “None” timebase.

The following table lists the default latency threshold values for FPI monitoring. They are binary, in that the threshold value is either present or it is not.

TABLE 44 Default Fabric Performance Impact latency monitoring threshold values and actions

Monitoring statistic	Value (Y/N) for all policies	Actions
Fabric Performance Impact (DEV_LATENCY_IMPACT)	IO_FRAME_LOSS IO_PERF_IMPACT	RASLOG, SNMP, EMAIL, SDDQ, TOGGLE

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 45 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**
Receive Bandwidth usage percentage (RX)	60	75	90	EMAIL, SNMP, RASLOG
Transmit Bandwidth usage percentage (TX)	60	75	90	EMAIL, SNMP, RASLOG
Trunk Utilization percentage (UTIL)	60	75	90	EMAIL, SNMP, RASLOG

Switch status policy monitoring thresholds

The only timebase the Switch status monitors support is the “None” timebase. The following tables list the default switch status policy monitoring thresholds used by MAPS. All threshold conditions are absolute and actions are triggered when the reported value is greater than or equal to the threshold value. For thresholds with both an upper value and a lower value, an action is triggered when the reported value exceeds the upper threshold value or drops below the lower threshold value. The following tables list the default Switch Status monitoring threshold values and actions, broken out by policy.

The following table lists the default Switch Status monitoring thresholds and actions for the default aggressive policy.

TABLE 46 Aggressive policy default Switch Status monitoring thresholds and actions

Monitoring statistic	Switch status threshold values (Marginal/Critical)	Actions
Absent or faulty power supply (BAD_PWR)	DCX, DCX+: -/3 DCX-4S, DCX-4S+: -/1 All other platforms: 1/2	SW_CRITICAL, SNMP, EMAIL
Temperature sensors outside range (BAD_TEMP)	1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
Absent or faulty fans (BAD_FAN)	1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
Flash usage (FLASH_USAGE)	90	RASLOG, SNMP, EMAIL
Percentage of marginal ports (MARG_PORTS)	-/5	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Percentage of error ports (ERR_PORTS)	-/5	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Percentage of faulty ports (FAULTY_PORTS)	-/5	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Faulty blades (FAULTY_BLADE)	DCX, DCX+: 1/- DCX-4S/4S+: 1/-	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: No Action
Faulty WWN (WWN_DOWN)	DCX, DCX+: -/1 DCX-4S/4S+: -/1	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Core blade monitoring (DOWN_CORE)	DCX, DCX+: 1/2 DCX-4S, DCX-4S+: 1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
HA Sync (HA_SYNC)	DCX, DCX+, DCX-4S, DCX-4S+: sync=0	SW_MARGINAL, SNMP, EMAIL

The following table lists the default Switch Status monitoring thresholds and actions for the default moderate policy.

TABLE 47 Moderate policy default Switch Status monitoring thresholds and actions

Monitoring statistic	Switch status threshold values (Marginal/Critical)	Actions
Absent or faulty power supply (BAD_PWR)	DCX, DCX+: -/3 DCX-4S, DCX-4S+: -/1 All other platforms: 1/2	SW_CRITICAL, SNMP, EMAIL

TABLE 47 Moderate policy default Switch Status monitoring thresholds and actions (Continued)

Monitoring statistic	Switch status threshold values (Marginal/Critical)	Actions
Temperature sensors outside range (BAD_TEMP)	1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
Absent or faulty fans (BAD_FAN)	1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
Flash usage (FLASH_USAGE)	90	RASLOG, SNMP, EMAIL
Percentage of marginal ports (MARG_PORTS)	6/10	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
Percentage of error ports (ERR_PORTS)	6/10	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
Percentage of faulty ports (FAULTY_PORTS)	6/10	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
Faulty blades (FAULTY_BLADE)	DCX, DCX+: 1/- DCX-4S, DCX-4S+: 1/-	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: No Action
Faulty WWN (WWN_DOWN)	DCX, DCX+: -/1 DCX-4S, DCX-4S+: -/1	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Core blade monitoring (DOWN_CORE)	DCX, DCX+: 1/2 DCX-4S, DCX-4S+: 1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
HA Sync (HA_SYNC)	DCX, DCX+, DCX-4S, DCX-4S+: sync=0	SW_MARGINAL, SNMP, EMAIL

The following table lists the default Switch Status monitoring thresholds and actions for the default conservative policy.

TABLE 48 Conservative policy default Switch Status monitoring thresholds and actions

Monitoring statistic	Switch status threshold values (Marginal/Critical)	Actions
Absent or faulty power supply (BAD_PWR)	DCX, DCX+: -/3 DCX-4S, DCX-4S+: -/1 All Other Platforms: 1/2	SW_CRITICAL, SNMP, EMAIL
Temperature sensors outside range (BAD_TEMP)	1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL

TABLE 48 Conservative policy default Switch Status monitoring thresholds and actions (Continued)

Monitoring statistic	Switch status threshold values (Marginal/Critical)	Actions
Absent or faulty fans (BAD_FAN)	1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
Flash usage (FLASH_USAGE)	90	RASLOG, SNMP, EMAIL
Percentage of marginal ports (MARG_PORTS)	11/25	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Percentage of error ports (ERR_PORTS)	11/25	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Percentage of faulty ports (FAULTY_PORTS)	11/25	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Faulty blades (FAULTY_BLADE)	DCX, DCX+: 1/- DCX-4S, DCX-4S+: 1/-	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: No Action
Faulty WWN (WWN_DOWN)	DCX, DCX+: -/1 DCX-4S, DCX-4S+: -/1	Marginal: No Action Critical: SW_CRITICAL, SNMP, EMAIL
Core blade monitoring (DOWN_CORE)	DCX, DCX+: 1/2 DCX-4S, DCX-4S+: 1/2	Marginal: SW_MARGINAL, SNMP, EMAIL Critical: SW_CRITICAL, SNMP, EMAIL
HA Sync (HA_SYNC)	DCX, DCX+, DCX-4S, DCX-4S+: sync=0	SW_MARGINAL, SNMP, EMAIL

Traffic Performance thresholds

The following Traffic Performance monitors support the Minute, Hour, and Day timebases: Receive throughput, Transmit Frame Count, Receive Frame Count, Transmit throughput, IO Read Command Count, IO Write Command Count, IO Read Data, IO Write Data. The “Throughput Degradation” monitor supports Minute, Hour, and Day timebases as well as the “None” timebase. The “Time tx cred zero” monitor does not support any timebase.

The following table lists the default monitoring thresholds for traffic performance used by MAPS. All thresholds are measured per minute and are triggered when they are greater than the shown value. There are no default monitoring thresholds for any other Traffic Performance monitors.

TABLE 49 Default Traffic Performance monitoring thresholds and actions

Monitoring statistic	Traffic Performance monitoring threshold values by policy			Actions
	Aggressive	**Moderate**	**Conservative**	
Receive Bandwidth usage percentage (RX)	60	75	90	RASLOG, SNMP, EMAIL
Transmit Bandwidth usage percentage (TX)	60	75	90	RASLOG, SNMP, EMAIL
Trunk Utilization percentage (UTIL)	60	75	90	RASLOG, SNMP, EMAIL