Brocade® Fabric OS® MAPS User Guide, 9.2.x

User Guide
April 28, 2023

Table of Contents

Introduction

This section provides you with an overview of Monitoring and Alerting Policy Suite (MAPS), and information about its structural elements, set up, and use of its basic and advanced features.

About This Document

This document provides information on creating custom MAPS elements to address the strict monitoring requirements of a high-performance fabric. Although many different software and hardware configurations are tested and supported by Broadcom Inc. for Fabric OS™ v9.2.x, documenting all possible configurations and scenarios is beyond the scope of this document.

Supported Hardware and Software

The following hardware platforms are supported by Brocade Fabric OS 9.2.x.

Brocade Gen 7 (64G) Fixed-Port Switches

- Brocade G720 Switch
- Brocade G730 Switch
- Brocade 7850 Extension Switch

Brocade Gen 7 (64G) Directors

- Brocade X7-4 Director
- Brocade X7-8 Director

Brocade Gen 6 (32G) Fixed-Port Switches

- Brocade G610 Switch
- Brocade G620 Switch
- Brocade G630 Switch
- Brocade 7810 Extension Switch
- Brocade G648 Blade Server SAN I/O Module
- Brocade MXG610 Blade Server SAN I/O Module

Brocade Gen 6 (32G) Directors

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

Contacting Technical Support for Your Brocade® Product

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

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

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

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

Document Feedback

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

MAPS Overview

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

MAPS helps the SAN administrators perform the following:

- Proactively monitor the health and performance of the SAN infrastructure to ensure the application uptime and availability by leveraging predefined and custom policies and rules.
- Automate the policy-based real-time monitoring and alert thresholds on a port or group basis.
- Configure each MAPS-enabled switch to constantly monitor itself for potential faults and automatically alert you to problems before they become failures.
- Configure an entire fabric at one time using common rules and policies or custom policies for specific ports or switch elements with simplified fabric-wide threshold configuration, monitoring, and alerting.
- Simplify the configuration of a fabric with a common policy and remediate the conditions that require the attention.
- Automatically fence, toggle, or quarantine ports without a manual intervention.
- Monitor fabric-wide events, ports, FRUs, environmental parameters, security, traffic flows, and performance impacts.

MAPS is a licensed feature and requires a Fabric Vision® license. This guide assumes that the administrator has activated (enabled) the Fabric Vision license in the Fabric OS software to use the full set of MAPS options. For more information about licensing, refer to the *Brocade Fabric OS Software Licensing Guide*.

MAPS is shipped with predefined rules, groups, and policies. However, you can create custom rules, groups, and policies. MAPS enables the cloning of predefined rules, groups, and policies to facilitate monitoring. Predefined rules and groups are system-specific, and each system has its own rules and groups. All MAPS configuration is persistent, is retained across reboots or HA failovers, and can be uploaded and downloaded.

The following diagram describes the basic elements of MAPS and how they are structurally connected.

Figure 1: MAPS Elements

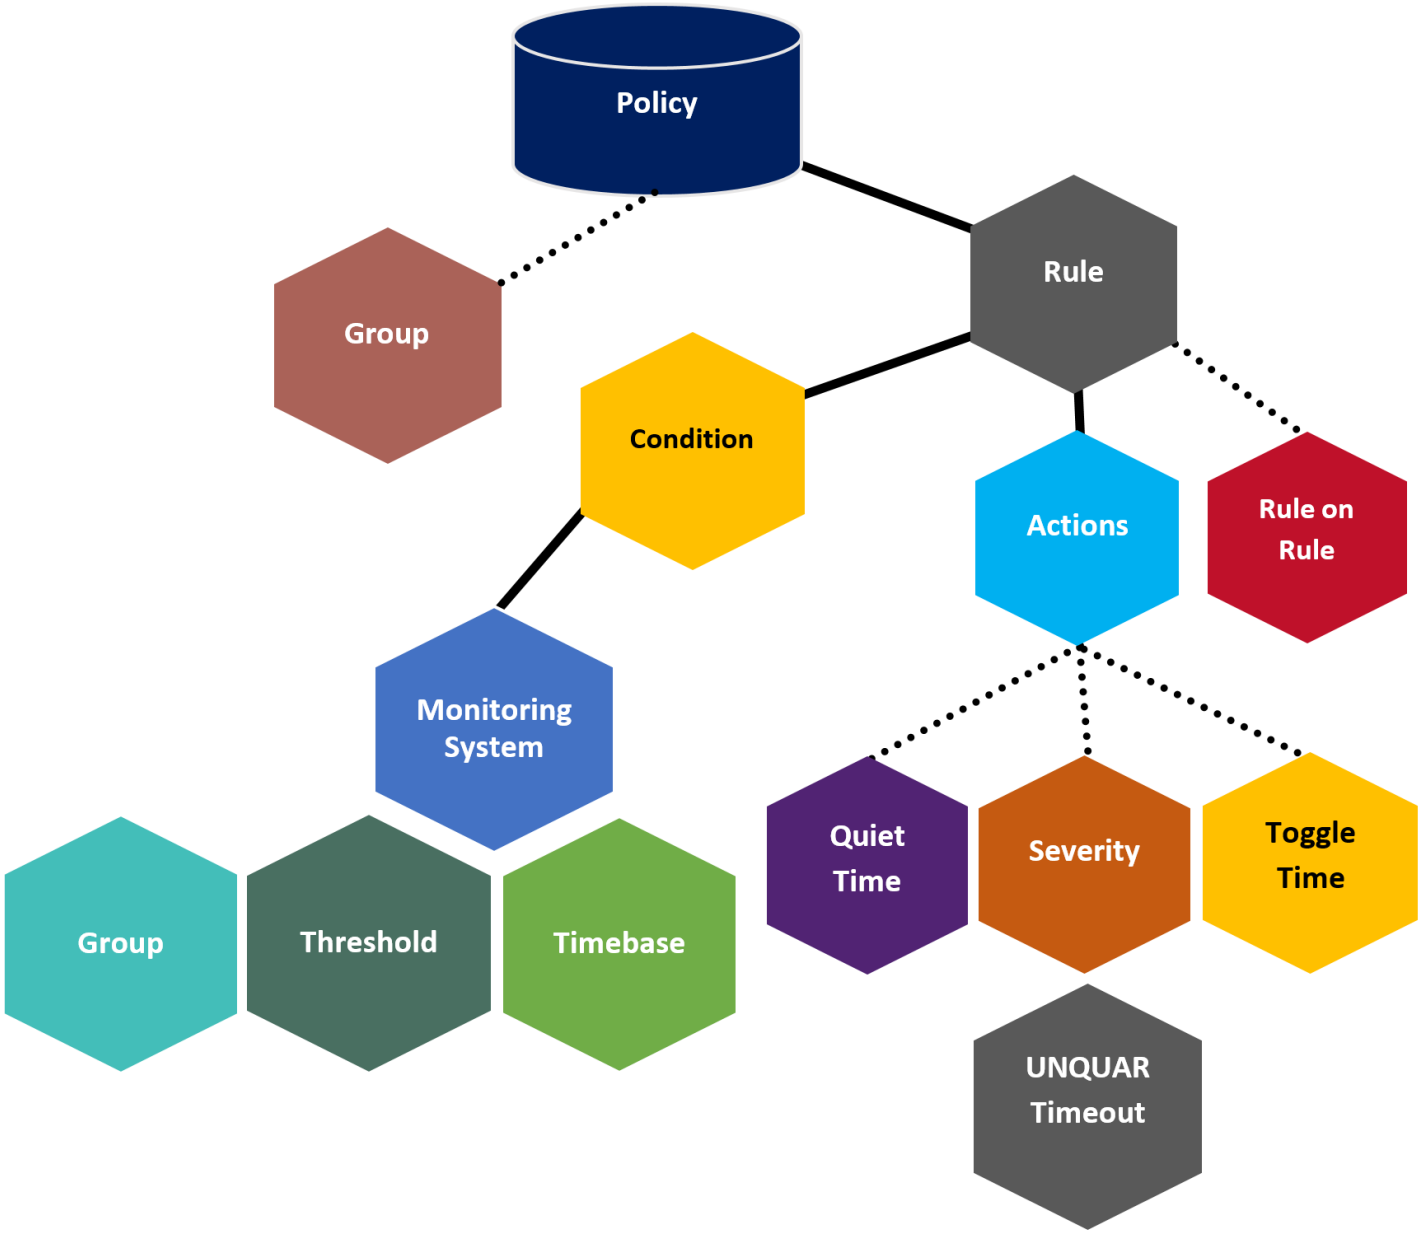

The dotted lines in the diagram represent the logical association between the MAPS elements. The black lines represent the physical association between MAPS elements; configuring these elements is mandatory to perform the monitoring. For more information about MAPS basic elements, see MAPS Basic Elements.

MAPS Initial Setup and Operation

You can easily configure MAPS to monitor a switch using one of the default policies.

Perform the following steps to configure MAPS to monitor the switch:

1. Connect to the switch and log on using an account with admin permissions.
2. Enter the `mapspolicy --enable` command followed by the name of the policy that you want to enable. For example, enter the `mapspolicy --enable dflt_conservative_policy` command to enable the conservative policy. You can enable the following default policies:
   - **dflt_conservative_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** – Contains rules with threshold values between aggressive and conservative policies.
   - **dflt_aggressive_policy** – Contains rules with stringent thresholds.
   - **dflt_base_policy** – Contains rules based on the features that can be monitored without a Fabric Vision license.
   - **dflt_always_active_policy** – Contains all default system rules to manage critical system resources. This system policy is always active and cannot be customized or deactivated.
3. Enter the `mapspolicy --show -summary` command to confirm that the specified policy is active.

   The following example shows the default policies available in MAPS; the active policy is dflt_aggressive_policy.

```
switch:admin> mapspolicy --show -summary
        Policy Name                    Number of Rules
-------------------------------------------------------------
dflt_aggressive_policy           :               398
dflt_moderate_policy             :               402
dflt_conservative_policy         :               402
dflt_base_policy                 :                44
dflt_always_active_policy        :                 6

Active Policy is 'dflt_aggressive_policy'. The policy 'dflt_always_active_policy' always monitors the
 system.
```

   Now, you have the basic policy set up and are ready to start monitoring your system using MAPS.
4. Enter the `mapsconfig --actions raslog` command to set the RASLog alert. Configuring the basic alert notification is important to get alerts from MAPS after enabling the basic policy.

   The following example shows the RASLog message when the system threshold is exceeded.

```
2023/01/23-16:57:00, [MAPS-1003], 569, WARNING, sw0, Temperature Sensor 2,
 Condition=ALL_TS(TEMP==OUT_OF_RANGE), Current Value:[TEMP, OUT_OF_RANGE],
 RuleName=defALL_TSTEMP_OUT_OF_RANGE, Dashboard Category=Switch Resource.
```

   The following example shows the dashboard output when the system threshold is exceeded.

```
switch:admin> mapsdb --show

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

DB start time:                  Mon Jan 23 16:54:43 2023
Active policy:                  dflt_aggressive_policy
Configured Notifications:       RASLOG,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL
```

```
Fenced Ports :                  None
Decommissioned Ports :          None
Fenced circuits :               N/A
Quarantined Ports :             None
Top Zoned PIDs <pid(it-flows)>:

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

Current Switch Policy Status: HEALTHY


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

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


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

Category(Violation Count)|RepeatCount|Rule Name      |Execution Time   |Object     |Triggered
 Value(Units)|
-------------------------------------------------------------------------------------------------------------
Switch Resource(1)       |1          |defALL_TSTEMP_ |01/23/23 16:57:00|Temperature|OUT_OF_RANGE
 |             |           |   OUT_OF_RANGE|                |   Sensor 2|
```

MAPS Basic Elements

You must understand the basic elements of MAPS to effectively set up MAPS operations before you start monitoring with MAPS.

MAPS contains structural elements that help you monitor a switch or fabric effectively.

Policy

A MAPS policy is a set of rules. A switch can have multiple policies, but you can activate or enable only one policy at a time. Once the policy is active, all the rules in the active policy take effect to monitor the switch. One policy must always be active on the switch along with dflt_always_active_policy.

The `mapspolicy --show -summary` command lists out the policies available in MAPS. The following example displays all the default policies available in MAPS for your use. You can define your custom policies to monitor your switch or chassis when the basic setup is running.

```
switch:admin> mapspolicy --show -summary
        Policy Name                    Number of Rules
-----------------------------------------------------------
dflt_aggressive_policy           :              398
dflt_moderate_policy             :              402
dflt_conservative_policy         :              402
dflt_base_policy                 :               44
dflt_always_active_policy        :                6

Active Policy is 'dflt_aggressive_policy'. The policy 'dflt_always_active_policy' always monitors the system.
```

You can have an active policy with no rules, but you must have an active policy. An active policy cannot be disabled. You can only change the active policy by enabling a different policy. Do not use policies without any rules.

You can use the `mapspolicy --show -all` command to view the complete list of MAPS policies with the associated rules.

To view the specific policy, enter the `mapspolicy --show` command followed by the policy name. The following example shows the `dflt_aggressive_policy`.

```
switch:admin> mapspolicy --show dflt_aggressive_policy
Policy Name: dflt_aggressive_policy

Rule Name                              |Condition                                    |Actions                    |
---------------------------------------------------------------------------------------------------------------
defALL_100M_16GSWL_QSFPCURRENT_1       |ALL_100M_16GSWL_QSFP(CURRENT/NONE<=1)        |SFP_MARGINAL,RASLOG,SNMP,EM|
                                       |                                             |AIL qt=1 hour              |
defALL_100M_16GSWL_QSFPCURRENT_10      |ALL_100M_16GSWL_QSFP(CURRENT/NONE>=10)       |SFP_MARGINAL,RASLOG,SNMP,EM|
                                       |                                             |AIL qt=1 hour              |
defALL_100M_16GSWL_QSFPRXP_2187        |ALL_100M_16GSWL_QSFP(RXP/NONE>=2187)         |SFP_MARGINAL,RASLOG,SNMP,EM|
                                       |                                             |AIL qt=1 hour              |
defALL_100M_16GSWL_QSFPRXP_44          |ALL_100M_16GSWL_QSFP(RXP/NONE<=44)           |SFP_MARGINAL,RASLOG,SNMP,EM|
                                       |                                             |AIL qt=1 hour              |

                          -------------------Output truncated---------------------

Active Policy is 'dflt_aggressive_policy'. The policy 'dflt_always_active_policy' always monitors the system.
```

```
Unmonitored Rules are prefixed with "*"
System rules are prefixed with "+"
```

MAPS helps you to filter the information based on the groups, monitoring systems, thresholds, and so on.

Select from the following default policies:

- **dflt_conservative_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** – Contains rules with thresholds values between aggressive and conservative policies.
- **dflt_aggressive_policy** – Contains rules with strict thresholds.
- **dflt_base_policy** – Contains rules that based on the features which can be monitored without a license.
- **dflt_always_active_policy** – Contains all default system rules to manage critical system resources. This system policy is always active and cannot be customized or deactivated.

For more information about maximum of Default, Custom, and RoR rules that are allowed per policy, see MAPS Scale Numbers.

Group

A MAPS group is a collection of similar objects that helps administrators to monitor as a single entity. You can create a group of objects and then use that group in rules, thus simplifying rule configuration and management. For example, you can create a group of ports that are connected to Windows servers, and then create specific rules for monitoring this group. Objects can be a member of multiple groups. Multiple rules can use the same group. A group is also referred to as a *logical group*.

MAPS allows you to view the information for all logical groups collectively or for a single specific group. To view a summary of all logical groups on a switch, enter the `logicalgroup --show` command. This command returns the group name and the information on whether the group is predefined. The output presents a table with columns that list characteristics for each group:

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

The following example shows the output of the `logicalgroup --show` command.

```
switch:admin> logicalgroup --show

--------------------------------------------------------------------------------
Group Name             |Predefined   |Type     |Member Count     |Members
--------------------------------------------------------------------------------
ALL_100M_16GSWL_QSFP   |Yes          |Sfp      |4                |52-55
ALL_32GSWL_QSFP        |Yes          |Sfp      |0                |
ALL_TARGET_PORTS       |Yes          |Port     |4                |10-11,26-27
ALL_ASICS              |Yes          |Asic     |0                |
ALL_16GLWL_SFP         |Yes          |Sfp      |0                |
ALL_HOST_PORTS         |Yes          |Port     |7                |16-18,21-22,38-39
ALL_SFP                |Yes          |Sfp      |27               |0,2-3,7,9-12,16-23
NON_E_F_PORTS          |Yes          |Port     |43               |1,4-6,8,13-15,28-31
ALL_10GSWL_SFP         |Yes          |Sfp      |0                |
SWITCH                 |Yes          |         |1                |0
------------------------Output truncated-----------------------------------------
ALL_QSFP               |Yes          |Sfp      |0                |
```

Rule

A MAPS rule associates a condition with actions that are triggered when the specified condition is evaluated to be true for a specified object.

Each rule specifies the following items:

- A group of objects to be evaluated.
- The condition being monitored. Each rule specifies a single condition that includes a timebase and a threshold.
- The actions to take when the condition or threshold value is attained.

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

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

```
switch:admin> mapsrule --show -all
Rule Name                        |Condition                        |Actions                        |
---------------------------------------------------------------------------------------------------
defALL_10GLWL_SFPCURRENT_8       |ALL_10GLWL_SFP(CURRENT/NONE<=8)  |SFP_MARGINAL,RASLOG,SNMP,EMAIL|
defALL_10GLWL_SFPCURRENT_95      |ALL_10GLWL_SFP(CURRENT/NONE>=95) |SFP_MARGINAL,RASLOG,SNMP,EMAIL|
defALL_10GLWL_SFPRXP_14          |ALL_10GLWL_SFP(RXP/NONE<=14)     |SFP_MARGINAL,RASLOG,SNMP,EMAIL|
defALL_10GLWL_SFPRXP_2230        |ALL_10GLWL_SFP(RXP/NONE>=2230)   |SFP_MARGINAL,RASLOG,SNMP,EMAIL|
                                 ----------Output truncated----------
```

Condition

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

For example, you can create a rule that monitors loss-of-signal errors on all host ports and that triggers actions when the counter is greater than 14. When the counter reaches 15, the rule triggers the actions. The following example shows a rule with the parameters that define a condition such as a monitoring system, threshold, and timebase are highlighted:

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

Monitoring System

A monitoring system is a value (measure or statistic) that is monitored by MAPS. MAPS monitors the changes in the monitoring system. These monitoring systems have configurable threshold values that MAPS can monitor and send alerts using RASLog notifications, SNMP messages, Email, and so on.

A monitoring system is part of a MAPS condition that also includes a timebase and a threshold. If the condition is evaluated as true, the actions that are specified in the rule are triggered.

The monitoring systems present in the BE Port Health category, such as CRC, LR, FRM_LONG, FRM_TRUNC, and ITW, help you monitor the errors on back-end switch ports. As an example, the CRC monitoring system lets you monitor the CRC errors for back-end ports.

The following example shows the CRC errors that are monitored by the CRC monitoring system for back-end ports.

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

             ------------------------Output truncated-----------------------------

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

```
Category                |Today                    |Last 7 days              |
-------------------------------------------------------------------------------
BE Port Health          |Out of operating range   |In operating range       |



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

Category(Violation Count)|RepeatCount|Rule Name               |Execution Time   |Object       |Triggered
 Value(Units)|
-------------------------------------------------------------------------------------------------------------------------------------------
BE Port Health(1)        |1          |defALL_BE_PORTSCRC_5M_10 |09/30/23 02:36:02|BE Port 11/88|26 CRC
         |



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

Stats(Units)        Current     09/29/23    --/--/--    --/--/--    --/--/--    --/--/--    --/--/--
-------------------------------------------------------------------------------------------------------
CRC(CRCs)           -           -           -           -           -           -           -
ITW(ITWs)           -           -           -           -           -           -           -


5 History Data for Backend ports:
=================================
Stats(Units)        Current     09/29/23    --/--/--    --/--/--    --/--/--    --/--/--    --/--/--
-------------------------------------------------------------------------------------------------------

CRC(CRCs)    11/88(26)          -           -           -           -           -           -
ITW(ITWs)           -           -           -           -           -           -           -

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

Timebase

The timebase (timebase value) specifies the time interval between two samples to be compared. For example, if you specify the timebase as minute in a rule that monitors CRC, the number of CRC errors that are observed in a minute time window is used to compare against the configured threshold. Similarly, if the timebase is specified as hour, the number of CRC errors in an hour is used to compare against the threshold. You can set the timebase for the following durations.

Table 1: Timebase Durations

Timebase	Description	How Often Samples Are Compared
10SEC	Samples that are used for comparison are 10 seconds apart.	Once every 10-seconds
minute	Samples that are used for comparison are one minute apart.	Once every minute

Timebase	Description	How Often Samples Are Compared
hour	Samples that are used for comparison are one hour apart.	Once every hour
day	Samples that are used for comparison are one day apart.	Once a day
week	Samples that are used for comparison are one week apart. (This timebase is valid only for rule-on-rule (RoR) rules except for the ASC_UPLOAD_FAILURE rule).	Once a week
IO	Samples that are used for monitoring flows per I/O in the hardware. All MAPS software violations are monitored at the port level and notified at the 10-seconds timebase. The MAPS software violations provide only the aggregated data and do not notify the type of the issue. The hardware violations are notified in real time with the information on the number of violated I/Os and the violated threshold at the monitored flow level. When the learning is enabled in the hardware on IT/ITL/VITL flows, real-time violations are aggregated in the parent flows, port, and device entity hierarchy.	Once in 10 seconds or None
none	A comparison is made between the latest value and the configured threshold value.	N/A

The following example shows using a timebase in a rule that monitors the number of times signal loss was observed on host ports in a day.

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

Threshold

Thresholds are the configured values for when potential problems might occur and alerts are generated.

While creating a rule using `mapsrule --create` , use the `-value` option to specify a threshold value that triggers specified actions when exceeded.

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

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

Each monitoring system supports only one type of threshold. You can configure any of the following different types of threshold values:

- integer
- float
- unsigned int
- strings

Operator

Logical operators allow you to create a condition with a monitoring system and threshold value. You can use the following operators to define the condition. These operators determine the value at which potential problems might occur.

- >
- >=
- =
- <=
- <

For example, the following rule shows the operator value as 14.

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

Actions

MAPS provides various actions that are triggered when a rule violation occurs. Some of the actions are simple notifications such as RASLog and simple network management protocol (SNMP) traps, and others are complex actions such as Port Toggle or Port Fencing.

For example, if you define a rule with an SNMP action and a violation of that rule occurs, the switch sends that SNMP trap to all SNMP trap receivers configured on the switch.

Table 2: MAPS Actions

Action	Description
DECOM	The port decommissioning (DECOM) provides an automated mechanism to remove an E_Port or F_Port from use. The DECOM action automatically takes ports offline when configured thresholds in a given rule are exceeded. Port decommissioning takes a port offline without a loss of traffic. This action is disabled by default. Port decommissioning and port fencing can be configured only for the port health monitoring systems for which decommissioning is supported. Port decommissioning cannot be configured by itself in a MAPS rule or action. It requires port fencing to be enabled in the same rule.
EMAIL	An Email alert sends information about a switch event to one or more specified Email addresses. The Email alert specifies the threshold and describes the event similar to an error message. For example, abc@12.com is a valid Email address; abc@12 is not.
FENCE	The port fencing action takes the ports offline when the user-defined thresholds are exceeded. This action immediately takes ports offline, which might cause a loss of traffic. If fencing is part of the overall switch configuration, it occurs anytime the port fails; whereas if it is part of a rule, the port is fenced if triggered. Multiple rules can have port fencing as an action; it occurs if any of the rules are triggered.
FMS	A FICON Management Server (FMS) notification is sent to the FICON management service as an action from MAPS. The FICON management service uses MAPS events to create a health summary report. Rules with a FICON notification action are part of all four default policies.

Action	Description
FPIN	Fabric Performance Impact Notification (FPIN) is a default MAPS action that enables FPI event reporting to members of the fabric that are participating in fabric notifications. The MAPS FPIN action triggers fabric notifications under the following performance impact conditions in a fabric: • **Congestion** – Credit-stall and oversubscription conditions that are detected on an F_Port. • **Link Integrity** – Signal integrity degradation issues such as CRC and ITW that are detected on an F_Port. The following MAPS monitoring systems support FPIN: • CRC • ITW For more information about fabric notifications, refer to the "Fabric Notifications" chapter in the *Brocade Fabric OS Administration Guide*. The FPIN action is a default MAPS action with the Fabric OS v9.1.x and later versions. The MAPS FPIN action is supported on Access Gateway (AG) switches.
RASLog	The RASLog action adds an entry to the RASLog event log for an individual switch. The RASLog stores infrequent event information and errors but does not actively send alerts. RASLog acts as a centralized logging mechanism for Fabric OS and covers error detection, reporting, and presenting useful data for administrators.
RE_BALANCE	The RE_BALANCE action brings the port group state back to the balanced state from the im_balance state. Administrators can configure an automatic re_balance action or can perform a manual re_balance action from im_balance state using the `devicelogin` command.
SDDQ	The Slow-Drain Device Quarantine (SDDQ) action automatically moves the traffic that is destined to the F_Port connected to a slow-draining device to a low-priority virtual channel. This action does not disable the port; instead, it reduces the effect of its latency on other flows in the fabric. SDDQ actions can be configured to only monitor rules, for example, the DEV_LATENCY_IMPACT state of IO_PERF_IMPACT and IO_FRAME_LOSS. • The SDDQ and Toggle actions are mutually exclusive. • Apply the SDDQ action universally (when activated) on all ports in the fabric.
SFP_MARGINAL	The SFP_MARGINAL action places the SFP into a marginal operating state. When a threshold is triggered, the SFP transceiver status changes to a marginal status. Port SFP transceiver measures include Current, Receive Power, Transmit Power, Voltage, Temperature, and Power On Hours. The SFP_MARGINAL action is enabled by default. SFP Marginal is available for SFP transceiver measures in the Port Health category.
SNMP	In environments in which you have a high number of messages coming from various switches, you might want to receive them in a single location and view them using a graphical user interface (GUI). In this scenario, SNMP notifications can be the most efficient notification method. You can avoid logging in to each switch individually, as you would have to do for error log notifications. SNMP performs an operation that is called a *trap*, which notifies a management station using SNMP when specific events occur on a switch. An SNMP trap delivered to the SNMP management station contains the following information: • Name of the element with the counter that registered an event • Class, 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 You must configure the software to receive trap information from the network device. You must also configure the SNMP trap receiver on the switch using the `snmpConfig` command to send the trap to the management station.

Action	Description
SW_CRITICAL	The SW_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 it affects what is displayed in the dashboard. This action is valid only in the context of rules that are related to Switch Status Policy.
SW_MARGINAL	The SW_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 it affects what is displayed in the dashboard. This action is valid only in the context of rules that are related to Switch Status Policy.
TOGGLE	The port toggling (TOGGLE) action temporarily disables a port and then re-enables it. Thus, it allows the port to reset and helps to restore the expected fabric performance. It attempts to recover the congestion conditions that are caused by slow-draining devices using the FPI (DEV_LATENCY_IMPACT) measure.
UNINSTALL_VTAP	The UNINSTALL_VTAP action uninstalls the vTAP feature if the mirrored frame count exceeds 250K IOPS and encryption is enabled on the 16G-capable ASIC. If encryption is not enabled on the ASIC, vTAP is not uninstalled.
UNQUAR	The UNQUAR (unquarantined) action is used to unquarantine the ports automatically under the following conditions: • When a port is moved to the quarantined state, user intervention is required to move the port back to the normal state. • When the Slow-Drain behavior is infrequent. For more information about quarantining, see Slow-Drain Device Quarantining.

The following example enables the RASLog notification by activating the RASLOG action using the `mapsconfig --actions <RASLOG>` command.

```
switch:admin> mapsconfig --actions RASLOG

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

After enabling the RASLog action, you can see RASLog alerts similar to the following example.

```
switch:admin> errshow
[MAPS-1003],1173, SLOT 1 | FID 1 | PORT 7/36, WARNING,X6-4_140_061, slot7 port36, F-Port 7/36,
 Condition=ALL_F_PORTS(DEV_LOGIN_DIST==IMBALANCED), Current Value:[DEV_LOGIN_DIST,IMBALANCED],
 RuleName=defALL_F_PORTSDEV_LOGIN_DIST_IMBALANCED, Dashboard Category=Fabric Performance Impact.

[MAPS-1003],1174, SLOT 1 | FID 1 | PORT 7/36, WARNING, X6-4_140_061, slot7 port36, F-Port
 7/36, Condition=ALL_F_PORTS(DEV_LOGIN_DIST==BALANCED), Current Value:[DEV_LOGIN_DIST,
 BALANCED],RuleName=defALL_F_PORTSDEV_LOGIN_DIST_BALANCED, Dashboard Category=Fabric Performance Impact.
```

You can use the following command to enable the port toggle action.

```
switch:admin> mapsconfig -actions TOGGLE
```

Quiet Time

The Quiet Time feature allows you to configure a time interval to suppress recurring alerts after the first alert is sent. Any successive alerts are disabled for the configured Quiet Time period. There are two ways you can configure the Quiet Time duration:

- **Rule Quiet Time** – Configured in a MAPS rule and applicable to the rule.
- **Global Quiet Time** – Configured at the global level and applicable to the current enabled policy.

Rule Quiet Time

Quiet Time is an optional rule parameter for the `mapsrule` command. Defining the `-qt <value>` parameter in a rule blocks MAPS from sending another alert based on the same rule for the specified number of seconds after it sends the initial alert.

The following actions support Quiet Time duration in a rule:

- RASLog
- E-Mail
- SNMP
- FPIN
- FMS

The Quiet Time feature suppresses an action for the specified duration when a rule is violated. MAPS supports different units for the Quiet Time duration: second, minute, hour, and day. The default unit is seconds.

The following example shows a rule that includes a 20-hour Quiet Time period.

```
mapsrule --create toggle_crc_rule -group ALL_PORTS -monitor CRC -timebase hour -op ge -value 0 -action
 raslog,email -qt 20 -unit hour
```

This rule generates one action every 20 hours, and then after 20 hours, it sends a summary with an alert of the MAPS-1005 RASLog message. You can modify a rule using the `mapsrule --config` command. See the following example.

```
mapsrule --config toggle_crc_rule -qt 27 -unit hour
```

The maximum Quiet Time is 30 days, regardless of the unit used to configure the Quiet Time.

Global Quiet Time

Global Quiet Time allows you to configure the Quiet Time duration at the global level and then apply it to the enabled MAPS policy. This process helps to avoid configuring the Quiet Time duration for multiple rules in a policy. However, if a default rule that is configured with a Quiet Time duration is part of an enabled policy, then the following results occur:

- Global Quiet Time overrides the Quiet Time duration present in the default rule.
- Global Quiet Time maintains the Quiet Time duration for the custom rule.

The Global Quiet Time configuration is valid only for logical switches, and this configuration is specific to each logical switch.

A Global Quiet Time of 24 hours is supported only on Brocade G730 switches running Fabric OS v9.2.x. In other platforms, the Global Quiet Time is not configured by default.

The following CLI output shows the default Global Quiet Time of 24 hours.

```
switch:admin> mapsconfig --show
Configured Notifications:       RASLOG,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL
Mail Recipient:                 ****@broadcom.com
Mail From Address:              vesp@bsn.in
Raslog Mode:                    Custom
Decom Action Config:            Impair (No Disable)
Global Quiet Time:              24 Hours
Notification Behavior:          Default
Paused members :
================
PORT :
```

```
CIRCUIT :
SFP :
SWITCH :
```

The unit for the default Global Quiet Time is hour. The minimum Quiet Time duration is 1 hour and the maximum duration is 30 days.

The following actions support Global Quiet Time:

- RASLog
- E-Mail
- SNMP
- FPIN
- FMS

Use the `mapsconfig --qt -value <value> [-unit hour|day]` command to configure the Global Quiet Time for a specific duration. The following command configures the Quiet Time duration for 5 hours.

```
switch:admin> mapsconfig --qt -value 5 -unit hour
WARNING: Quiet Time feature will be deprecated in future version.
```

The example above, a warning message indicates that the Quiet Time feature will be deprecated in a future release. To suppress recurring and multiple MAPS notifications, you can activate the adaptive notification feature. For more information, see Adaptive Notifications.

In the following example, the 5-hour Quiet Time is configured.

```
switch:admin> mapsconfig --show
Configured Notifications:       SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,DECOM,FPIN

   ------------------Output truncated-------------------------------
Global Quiet Time:              5 Hours
```

The following example show the RASLog alert, which indicates that the Global Quiet Time duration is configured for 5 hours.

```
2023/02/22-11:48:08, [MAPS-1209], 297, SLOT 2 | FID 128, INFO, Gen7-8, Global Quiet Time enabled with 5 Hours
```

Adaptive Notifications

In pre-Fabric OS v9.2.x release versions, the Quiet Time feature is used to configure a time interval for suppressing the following MAPS notifications after sending the first alert:

- RASLog
- SNMP
- Email

When a MAPS rule is violated and the first alert is sent, all subsequent alerts are disabled for the configured Quiet Time. Multiple notifications for multiple rule violations might occur before the issue is resolved. These notifications are suppressed only for the configured Quiet Time.

The Fabric OS v9.2.x release introduces the **Adaptive Notification** feature. It effectively suppresses recurrent and multiple MAPS notifications in a sequence of Quiet Time intervals following the initial notification. With the adaptive notification feature enabled, the sequence of Quiet Time intervals is determined automatically based on changing conditions until the issue is resolved. The MAPS violations are summarized at the end of the Quiet Time interval and thereafter the subsequent violations result in a new alert with the next Quiet Time interval. By modulating the number of notifications, it enhances the impact of the notified information.

The adaptive notification automatically uses the timebase that is configured in the rule and MAPS object as the base to determine the following sequential Quiet Time intervals:

- 1 hour
- 6 hours
- 1 day
- 7 days
- 30 days

The advantages of adaptive notifications include the following benefits:

- Applying sequential Quiet Time intervals automatically in response to changing conditions
- Suppressing recurring and continuous MAPS notifications
- Avoiding configurations at the rule level
- Avoiding a static Quiet Time configuration that is not progressive

Fabric OS v9.2.x disables both the Global Quiet Time and Rule Quiet Time features when the adaptive notification feature is enabled. The adaptive notification feature is not enabled by default. Once the adaptive notification is enabled, the Quiet Time configurations are no longer considered.

The configuration is chassis-wide, and any changes made at any logical switch will affect the behavior of the entire system. To enable the MAPS notification mode to either the default Quiet Time configuration or adaptive notifications, use the following command:

```
mapsconfig --notification {default|adaptive}
```

The following CLI output shows that the adaptive notification is enabled:

```
switch:admin> mapsconfig --notification adaptive

WARNING: Adaptive MAPS Notification feature is enabled. Quiet Time configurations will be ignored.
```

The following RASLog message shows that the MAPS notification mode is set to adaptive notification:

```
2023/01/29-14:56:43 (GMT), [MAPS-1220], 71, FID 103, INFO, Vesper120_103, MAPS Notification mode is set to
 Adaptive.
```

To switch to the pre-Fabric OS v9.2.x Quiet Time configuration with continuous notifications, use the following command:

```
switch:admin> mapsconfig --notification default

WARNING: Quiet Time feature will be deprecated in future version.
```

The following RASLog message shows that the MAPS notification mode is set to the default Quiet Time configuration with continuous notifications:

```
2023/01/29-14:57:11 (GMT), [MAPS-1220], 72, FID 103, INFO, Vesper120_103, MAPS Notification mode is set to
 Default.
```

The following CLI output shows an example of the status of the MAPS notification behavior:

```
switch:admin> mapsconfig --show
Configured Notifications:       RASLOG,SNMP,FENCE,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,DECOM,SDDQ
Mail Recipient:                 joe@company.com
Mail From Address:              vesp@bsn.in
Raslog Mode:                    Custom
Decom Action Config:            Impair (No Disable)
Global Quiet Time:              24 Hours (Not Effective)
Notification Behavior:          Adaptive
Paused members :
================
```

```
PORT :
CIRCUIT :
SFP :
SWITCH :
```

The following example shows the RASLog message that was generated when CRC frames were detected. When the Adaptive Notification feature is enabled, a real-time alert is triggered on port 12 within the first Quiet Time interval of 1 hour. Therefore, notifications will not be triggered for 1 hour following the initial alert.

```
2023/01/09-08:36:30 (GMT), [MAPS-1003], 34110, FID 29 | PORT 0/12, WARNING, REMOTE_FABRIC_TYR, port12, F-
Port 12, Condition=ALL_HOST_PORTS(CRC/min>0), Current Value:[CRC, 10 CRCs], RuleName=defALL_HOST_PORTSCRC_0,
 Dashboard Category=Port Health, Quiet Time=1 hour.
```

The subsequent notifications are suppressed until the first Quiet Time interval, which is 1 hour, if the CRC issue cannot be resolved by either replacing the cable, transceiver, or by checking the switch port. The Quiet Time summary notification is triggered if the issue is not resolved within 1 hour. As shown in the following example, a second alert with the next 6-hour Quiet Time interval is generated when the issue is not resolved.

```
2023/01/09-10:06:30 (GMT), [MAPS-1003], 34119, FID 29 | PORT 0/12, WARNING, REMOTE_FABRIC_TYR, port12, F-Port
 12, Condition=ALL_HOST_PORTS(CRC/min>0), Current Value:[CRC, 100 CRCs], RuleName=defALL_HOST_PORTSCRC_0,
 Dashboard Category=Port Health,Quiet Time=6 hour.
```

Upon failure of the CRC issue to be resolved within the 6-hour Quiet Time interval, the real-time notification is triggered with the subsequent 1-day Quiet Time interval. The following example shows the Quiet Time summary with the 1-day Quiet Time interval:

```
2023/01/09-17:06:36 (GMT), [MAPS-1003], 34126, FID 29 | PORT 0/12, WARNING, REMOTE_FABRIC_TYR, port12, F-Port
 12, Condition=ALL_HOST_PORTS(CRC/min>0), Current Value:[CRC, 350 CRCs], RuleName=defALL_HOST_PORTSCRC_0,
 Dashboard Category=Port Health, Quiet Time=1 day.
```

The following RASLog examples show the real-time notifications with subsequent Quiet Time intervals, such as 7 days and 30 days, if the CRC issue is not resolved:

```
2023/01/11-03:48:54 (GMT), [MAPS-1003], 34144, FID 29 | PORT 0/12, WARNING, REMOTE_FABRIC_TYR, port12, F-Port
 12, Condition=ALL_HOST_PORTS(CRC/min>0), Current Value:[CRC, 200 CRCs], RuleName=defALL_HOST_PORTSCRC_0,
 Dashboard Category=Port Health,Quiet Time=7 day.

2023/01/18-04:32:57 (GMT), [MAPS-1003], 34144, FID 29 | PORT 0/12, WARNING, REMOTE_FABRIC_TYR, port12, F-Port
 12, Condition=ALL_HOST_PORTS(CRC/min>0), Current Value:[CRC, 400 CRCs], RuleName=defALL_HOST_PORTSCRC_0,
 Dashboard Category=Port Health,Quiet Time=30 day.
```

After the last Quiet Time interval has expired, and the CRC issue has not been resolved, the adaptive notification is set to remain in the last Quiet Time interval, which is 30 days, until the issue has been resolved.

The Adaptive Quiet Time restarts if there have been no violations for 24 hours following the last violation.

Toggle Time

Toggle time is the duration for which the port is temporarily disabled and then re-enabled. MAPS supports toggle time, which allows Fabric OS software to recover a congested port by a device. While there are many reasons why the device could be congested, 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 it to reset and recover from the issue. To enable recovering ports using port toggling, MAPS assumes a redundant path to the device. It does not check to see if there is one, nor can it check to see if traffic to or from the device is switched over to a redundant path. MAPS also assumes that while the port is being toggled, the operational state of the port is not changed by any other mechanism, such as an administrator disabling or moving the port or a port fencing operation.

The following example defines a rule that toggles a port offline for 180 seconds (3 minutes) when the number of CRC errors on the port in a minute is greater than 0:

```
switch:admin> mapsrule --create toggle_rule -group ALL_F_PORTS -monitor DEV_LATENCY_IMPACT -timebase none -op
 eq -value IO_PERF_IMPACT -action TOGGLE -tt 180
```

Category

A dashboard monitoring category is the grouping of a similar switch or system elements that can be monitored, for example, Security Violations.

The following monitoring categories are available in MAPS.

Table 3: Monitoring Categories

Monitoring Category	Description
Port Health	Allows you to monitor port statistics and take action based on the configured thresholds and actions.
BE Port Health	Allows you to monitor the health of the back-end switch ports for the following: • CRC • Link Reset • Error rates • Invalid Transmission Words (ITW) • BAD_OS • Frame Length (either too long or truncated)
Extension Gigabit Ethernet Ports	Allows you to monitor the statistics for GE ports.
FRU Health	Allows you to monitor FRU events such as INSERT and ON.
Security Violations	Allows you to monitors various security violations such as logon violations.
Fabric State Changes	Allows you to monitor the potential fabric or switch-related issues such as zone changes, fabric segmentation, E_Port down, fabric reconfiguration, domain ID changes, and fabric logons.
Switch Resource	Allows you to monitor the temperature, flash usage, memory usage, and CPU usage of your system.
Extension Health	Allows you to define rules for Extension health, including circuit state changes, circuit state utilization, and packet loss.
Fabric Performance Impact (FPI)	Allows you to monitor the congestion-related issues on all physical E_Ports and F_Ports. Specifically, FPI monitors abnormal device latency and oversubscription issues.
IO Health	Allows you to monitor the health of SCSI/NVME flow exceptions and their associated responses.
IO Latency	Allows you to monitor an I/O operation to complete between an initiator and a target over a transmission medium.
Switch Status Policy	Allows you to monitor the health of the switch by defining the number of types of errors that transition the overall switch state into a state that is not healthy.

The following examples display the available monitoring categories.

```
switch:admin> mapsdb --show

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

DB start time:                  Thu Sep 30 22:29:43 2021
Active policy:                  dflt_aggressive_policy
Configured Notifications:       RASLOG,SNMP,EMAIL,FENCE,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,DECOM,SDDQ,FPIN
Fenced Ports :                  None
```

```
Decommissioned Ports :          None
Fenced circuits :               None
Quarantined Ports :             None
Top Zoned PIDs <pid(it-flows)>: 0x71a080(4) 0x7103c1(1) 0x7100c0(1) 0x7101c0(1) 0x7102c0(1)

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

Current Switch Policy Status: MARGINAL
Contributing Factors:
---------------------
*EXPIRED_CERTS (MARGINAL).


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

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


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

Category(Violation Count)|RepeatCount|Rule Name                 |Execution Time   |Object          |
Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------------------------------
Security Violations(1)   |1          |defCHASSISCERTS_EXPIRED   |09/30/21 22:30:03|Chassis         |1 certs
              |
```

Dashboard

The MAPS dashboard provides a summary view of the switch health status that allows you to determine whether everything is working according to policy or whether you must investigate further. The dashboard provides an integrated dashboard view that helps administrators easily identify trends and quickly isolate issues occurring on a switch or fabric.

The dashboard view contains the following reports:

- **Dashboard Information** – Displays the following basic dashboard data:

  - the time the dashboard was started
  - the name of the active policy
  - any fenced, decommissioned, or quarantined ports
  - the list of FCIP circuits that are fenced
  - top PIDs
- **Switch Health Report** – 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.
- **Summary Report** – Contains two subsections:
  a. **Category Report** – Collects and summarizes the various switch statistics that are monitored by MAPS into multiple categories, displays the current status of each category since midnight and the status of each category for the past seven days.
  b. **Rules Affecting Health** – Displays if a rule violation has caused a change in the status of a category and rule-related information.

The following example displays the MAPS dashboard view.

```
switch:admin> mapsdb --show

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

DB start time:                  Mon Jan 22 02:18:11 2021
Active policy:                  dflt_base_policy
Configured Notifications:       SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,FPIN

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

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

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

3.2 Rules Affecting Health:
===========================
Category(Violation Count) |RepeatCount|Rule Name                  |Execution Time   |Object           |
Triggered Value(Units)|
```

```
Fabric Performance Impact |1          |defALL_PORTS_IO_LATENCY_CLE|03/17/21 12:01:01|E-Port 8/35      |
IO_LATENCY_CLEAR       |

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

Unquarantine Timeout

The unquarantine (UNQUAR) action is used to unquarantine the ports automatically. The unquarantine timeout is the value at which MAPS performs an automatic unquarantine if the UNQUAR action is specified in the FPI rule and the UNQUAR action is globally enabled in `mapsconfig` , when the SDDQ state is clear.

The following example displays the RASLog messages for the devices that are in quarantined and unquarantined states.

```
[NS-1026], 372, FID 128, INFO, G620, Local domain 1, port index 40: A zoned device 0x012800 has been
 quarantined.
[NS-1024], 373, FID 128, INFO, G620, Flow from device 0x012200 to quarantined device 0x012800 of Local domain
 1 has been successfully quarantined.
[NS-1024], 374, FID 128, INFO, G620, Flow from device 0x012300 to quarantined device 0x012800 of Local domain
 1 has been successfully quarantined.
[MAPS-1022], 375, FID 128 | PORT 0/40, WARNING, G620, Port 40 (Port index 40) has been marked as Slow Drain
 Device.
[NS-1027], 377, FID 128, INFO, G620, Local domain 1, port index 40: A zoned quarantined device 0x012800 has
 been unquarantined.
[NS-1025], 378, FID 128, INFO, G620, Flow from device 0x012200 to quarantined device 0x012800 of Local domain
 1 has been successfully unquarantined.
[NS-1025], 379, FID 128, INFO, G620, Flow from device 0x012300 to quarantined device 0x012800 of Local domain
 1 has been successfully unquarantined.
[MAPS-1025], 380, FID 128 | PORT 0/40, WARNING, G620, Port 40 (Port index 40) removed from the Slow Drain
 Device Quarantine Group.
[AG-1086], 513, FID 128 | PORT 0/8, INFO, Switch, Slow drain device 0x011205 connected to F-port (8) have
 been quarantined.
[AG-1087], 528, FID 128 | PORT 0/8, INFO, Switch, Slow drain device 0x011205 connected to F-port (8) have
 been unquarantined.
```

Severity

You can configure the severity of the actions in a rule. By configuring the severity of actions on rules, you can ensure that certain actions correspond to how you filter alerts according to your policies.

The following severity levels are available in MAPS:

- critical
- error
- warning
- info
- default

The following example creates a rule with a critical severity level.

```
switch:admin> mapsrule --create crc_sev_rule -monitor CRC -group ALL_PORTS
                        -timebase min -op g -value 10 -action raslog,email
                        -severity critical
```

Table 4: Default Alert Severity for State-Based and Non-State-Based Monitoring Systems

Condition Description	RASLog Message Category	Rule Example
C3TXO and IO Frameloss rules with “>” or “=” threshold	Critical	defALL_E_PORTSC3TXTO_10 defALL_PORTS_IO_FRAME_LOSS_UNQUAR
FRU state monitoring with the Faulty state	Error	defALL_PORTSSFP_STATE_FAULTY defALL_PORTSSFP_STATE_IN
Util/Tx/Rx rules with the “>” threshold	Info	defALL_E_PORTSRX_75 defALL_E_PORTSTX_75 defALL_E_PORTSUTIL_75
Remaining all other default rules with the “>”, “>=” threshold	Warning	defALL_E_PORTSCRC_21

Rule on Rule

The rules created to monitor the performance of other rules are called rule-on-rule (RoR) rules. A RoR rule allows a rule to monitor the frequency of another rule being triggered. The frequency allows a threshold to be exceeded a certain number of times and not generate alerts or actions. However, if the threshold is exceeded several times, alerts can be generated and actions can be triggered.

The RoR rules allow you to take different actions based on the frequency of a base rule execution. For example, you can specify the RoR rule to ignore the first CPU spike and then take a RASLog action. For example, `chassis(CPU/none>=75) -action none` is the base rule. For this base rule, you can specify a RoR rule `chassis(cpu_rule/hour>=5) -action raslog` to alert you if the base rule is executed five times in an hour and not to alert if the base rule is triggered only once in an hour.

The following example shows that the CPUpolicy75 policy is active with the CPU base rule and RoR rule:

```
sw0:admin> mapspolicy --show CPUpolicy75
Policy Name: CPUpolicy75

Rule Name              |Condition                          |Actions                  |
-----------------------------------------------------------------------------------------
memory_rule            |chassis(MEMORY_USAGE/none>=60)     |raslog                   |
memory_ror             |chassis(memory_rule/hour>=5)       |raslog                   |
cpu_rule               |chassis(CPU/none>=75)              |raslog                   |
cpu_ror                |chassis(cpu_rule/hour>=5)          |raslog                   |

Active Policy is 'CPUpolicy75'.
Unmonitored Rules are prefixed with "*"
System rules are prefixed with "+"
```

The following output shows an example of the RASLog message that is generated when the base rule is violated:

```
2021/06/10-03:04:00, [MAPS-1003], 2107, FID 128, WARNING, sw0, Chassis, Condition=CHASSIS(CPU>=75), Current
 Value:[CPU, 81.50 %], RuleName=cpu_rule, Dashboard Category=Switch Resource.
```

The following output shows an example of the RASLog message that is generated when the RoR rule is violated:

```
2021/06/10-02:22:00, [MAPS-1003], 2093, FID 128, WARNING, sw0, Chassis, Condition=RULE(cpu_rule/hour>=5),
 Current Value:[cpu_rule, 5], RuleName=cpu_ror, Dashboard Category=Switch Resource.
```

The following example shows the dashboard both of output when the base and RoR rules are violated:

```
3.2 Rules Affecting Health:
```

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

Category(Violation Count)|RepeatCount|Rule Name     |Execution Time   |Object    |Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------
Switch Resource(28)      |2          |cpu_rule      |06/10/20 03:10:00|Chassis   |76.00 %               |
                         |           |              |                 |Chassis   |76.40 %               |
                         |2          |memory_rule   |06/10/20 03:10:00|Chassis   |60.17 %               |
                         |           |              |                 |Chassis   |61.04 %               |
                         |1          |cpu_ror       |06/10/20 02:22:00|Chassis   |5 Trigger cnt.        |
                         |1          |memory_ror    |06/10/20 02:22:00|Chassis   |5 Trigger cnt.        |
                         |10         |cpu_rule      |06/10/20 02:58:00|Chassis   |76.00 %               |
                         |           |              |                 |Chassis   |77.20 %               |
                         |           |              |                 |Chassis   |78.10 %               |
                         |           |              |                 |Chassis   |79.60 %               |
                         |           |              |                 |Chassis   |81.50 %               |
```

If the default policy does not meet your requirements, you can clone an existing default policy and add your custom rule. When you create a custom RoR rule, you need to account for the constraints associated with each rule parameter. To obtain detailed information, see Monitoring Systems Support Matrix.

FPI Profile

The Fabric Performance Impact (FPI) Profile is the MAPS feature that is used to manage and customize the FPI congestion thresholds. If the predefined thresholds are suitable for your monitoring requirements, you can continue to use them. In addition to the default FPI profile, MAPS allows you to create two more custom profiles with the transmit queue latency and credit latency values for IO_PERF_IMPACT and IO_FRAME_LOSS to suit your custom requirements.

The following example shows the default FPI profile that displays the static threshold configurations.

```
switch:admin> fpiprofile --show
-------------------------------------------------------------------------------------------------
FPI Profile         |Profile Type |IO_PERF_IMPACT      |IO_FRAME_LOSS       |OVERSUBSCRIBED(*)     |
-------------------------------------------------------------------------------------------------
dflt_fpi_profile    |F & E Ports  |TXQL          = 10ms|TXQL        = 80ms  |TX            = 70%   |
                    |             |CRED-ZERO-1s  = 70% |                    |                      |
                    |             |CRED-ZERO-5s  = 50% |                    |                      |
                    |             |CRED-ZERO-10s = 30% |                    |                      |
-------------------------------------------------------------------------------------------------
```

MAPS Categories

This section contains information on the MAPS monitoring categories.

MAPS provides the following categories for monitoring:

- Port Health
- BE Port Health
- Extension GE Port Health
- FRU Health
- Security Violations
- Fabric State Changes
- Switch Resource
- Extension Health
- Fabric Performance Impact
- IO Latency
- IO Health
- Switch Status Policy
- Uncategorized Monitoring Systems

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

Port Health

The Port Health category monitors port statistics, SFP diagnostic parameters, and physical fiber channel (FC) ports. This monitoring category helps to take action based on the configured thresholds and actions. You can configure thresholds by the port type, and you can apply the configuration to all ports of the specified type.

Table 5: Port Health Category Parameters

Monitoring System	Description
C3TXTO	Describes the number of Class 3 discard frames due to timeouts. C3TXTO is monitored for resident N_Ports using the NON_E_F_PORTS group.
ROUTING_ERR	Describes the monitoring system that is used for timely identification of frame drops due to the following factors: • Frames that are not routable • Frames with unreachable destinations The bad routes are identified only when the fabric is stable, and the errors during a fabric reconfiguration are ignored.
CRC	Describes the number of times an invalid cyclic redundancy check (CRC) error occurs on a port, or a frame that computes 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.
DEV_NPIV_LOGINS	Describes the number of NPIV logons to the device.

Monitoring System	Description
ITW	Describes the number of times an invalid transmission word (ITW) error occurs on a port. A word that is not transmitted successfully results in encoding errors. Invalid word messages usually indicate a hardware problem. The ITW counter includes any physical coding sub-layer (PCS) violations. ITW violations can occur due to an encoding in or an encoding out violation, a PCS violation, or all of these violation types. Encoding violations occur only at slow (8G or lower) speeds, and PCS violations occur only at high (10G or higher) speeds.
LF	Describes the number of times a link failure (LF) occurs on offline ports, or sends or receives the Not Operational Primitive Sequence (NOPS). Both physical and hardware problems can cause link failures. Link failures also frequently occur due to a loss of synchronization or a loss of signal.
LOSS_SIGNAL	Describes the number of times that a signal loss occurs in offline ports. Signal loss indicates that no data is moving through the port. A loss of signal usually indicates a hardware problem.
LOSS_SYNC	Describes the number of times a synchronization error occurs on the port. Two devices failed to communicate at the same speed. A link failure always accompanies synchronization errors. Loss of synchronization errors frequently occurs due to a faulty SFP transceiver or cable.
LR	Describes the ports where the number of link resets (LRs) exceeds the specified threshold value.
PE	Describes the number of times a protocol error (PE) occurs on a port. Occasionally, PEs occur due to software glitches. Persistent errors generally occur due to hardware problems.
CURRENT	Describes the amperage that is supplied to the SFP transceiver in milliamperes (mA). Current area events indicate hardware failures.
STATE_CHG	Describes the state of the port has changed due to any of the following reasons: • The port has gone offline. • The port has come online. • The port is faulty.
PWR_HRS	Describes the number of hours that an SFP transceiver has been powered up.
RXP	Describes the power of the incoming laser (receive power) in microwatts (µW). This value 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 deteriorated.
SFP_TEMP	Describes the temperature of the SFP transceiver in degrees Celsius. High temperatures indicate that the SFP transceiver could be in danger of damage.
TXP	Describes the power of the outgoing laser (transmit power) in microwatts (µW). This value is used to determine if the SFP transceiver is in good working condition. If the counter often exceeds the threshold, the SFP transceiver is deteriorated.
VOLTAGE	Describes the voltage that is supplied to the SFP transceiver in millivolts (mV). If this value exceeds the threshold, the SFP transceiver is deteriorating.
ENCR_BLK	Describes the encryption blocks.
ENCR_DISC	Describes the encryption discard.
ENCR_SHRT_FRM	Describes the encryption short frames.

MAPS provides default rules for monitoring the port health statistics. For more information about these rules, see Default Monitoring Rules.

If the default policy does not meet your requirements, you can clone an existing default policy and add your custom rule. When you create a custom rule, you must account for the constraints that are associated with each parameter of that rule. To obtain detailed information, see Monitoring Systems Support Matrix.

Rule Data

The following example shows the default rules for monitoring port health. These rules trigger the configured actions when the number of CRC errors reaches the threshold.

```
Rule Data:
----------
RuleName: defALL_HOST_PORTSCRC_21
Condition: ALL_HOST_PORTS(CRC/MIN>21)
Actions: RASLOG,SNMP,EMAIL,FPIN
Associated Policies: dflt_conservative_policy
```

RASLog Example

The following example shows the RASLog message that is generated when the number of CRC errors reaches the threshold.

```
2023/01/11-09:34:37 (GMT), [MAPS-1003], 432, SLOT 3 | FID 1 | PORT 3/33, WARNING, swd77, slot3 port33, F-Port
 3/33, Condition=ALL_HOST_PORTS(CRC/min>21), Current Value:[CRC, 22 CRCs], RuleName=defALL_HOST_PORTSCRC_21,
 Dashboard Category=Port Health, Quiet Time=None.
```

MAPS Dashboard Example

The following example shows the MAPS dashboard of the rules that affect the Port Health category status.

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

Category(Violation Count)|RepeatCount|Rule Name                  |Execution Time   |Object           |
Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------------------------------------------
Port Health(1)           |1          |defALL_HOST_PORTSCRC_21    |01/11/23 09:34:37|F-Port 3/33      |22 CRCs
               |
```

BE Port Health

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

The following table lists the monitored parameters in this category.

Table 6: BE Port Health Category Parameters

Monitoring System	Description
BAD_OS	Describes the number of invalid ordered sets (OS) outside the frame.
CRC	Describes the number of CRC errors.
FRM_LONG	Describes the number of detected frames that are longer than expected (greater than 2148 bytes).

Monitoring System	Description
FRM_TRUNC	Describes the number of detected frames that are truncated and shorter than expected (less than 36 bytes).
ITW	Describes the number of times an invalid transmission word (ITW) error occurs on a port. This count indicates that a word did not transmit successfully, resulting in encoding errors. Invalid word messages usually indicate a hardware problem.
LR	Describes the number of times a LR error occurs on a back-end port.

MAPS provides default rules to monitor the back-end port statistics. For more information about these rules, see Default Monitoring Rules.

If the default policy does not meet your requirements, you can clone an existing default policy and create your custom rule. When you create a custom rule, you need to account for the constraints that are associated with each rule parameter. To obtain detailed information, see Monitoring Systems Support Matrix.

Rule Data

The following example shows one of the default rules for monitoring the back-end port health. This rule triggers the configured actions when the OS count reaches more than 10 in 5 minutes.

```
Rule Data:
----------
RuleName: defALL_BE_PORTSBAD_OS_5M_10
Condition: ALL_BE_PORTS(BAD_OS/5MIN>10)
Actions: RASLOG,SNMP,EMAIL,FMS
Associated Policies: dflt_aggressive_policy, dflt_moderate_policy, dflt_conservative_policy
```

RASLog Example

The following example shows the RASLog message that is generated when the OS count reaches the threshold.

```
2023/01/11-09:46:01 (GMT), [MAPS-1003], 433, SLOT 2 | FID 1, WARNING, swd77, BE Port 3/97,
 Condition=ALL_BE_PORTS(BAD_OS/5MIN>10), Current Value:[BAD_OS, 12 Errors, (Conn. BE port: 6/64) ],
 RuleName=defALL_BE_PORTSBAD_OS_5M_10, Dashboard Category=BE Port Health, Quiet Time=None.
```

MAPS Dashboard Example

The following example shows the MAPS dashboard output that displays the rules affecting the health of the back-end port.

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

Category(Violation Count)|RepeatCount|Rule Name                  |Execution Time   |Object          |
Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------------------------------
Port Health(1)           |1          |defALL_HOST_PORTSCRC_21    |01/11/23 09:34:37|F-Port 3/33     |22 CRCs
              |
BE Port Health(1)        |1          |defALL_BE_PORTSBAD_OS_5M_10|01/11/23 09:46:01|BE Port 3/97    |12
 Errors             |
```

Extension GE Port Health

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

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

Table 7: Extension GE Port Health Category Parameters

Monitoring System	Description
GE_CRC	Describes the number of times an invalid cyclic redundancy check (CRC) error occurs on a GE port, or a frame computes 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.
GE_LOS_OF_SIG	Describes the number of times that signal loss occurs in offline GE ports. Signal loss indicates that no data is moving through the port. A loss of signal usually means a hardware problem.

MAPS provides default rules to monitor for the Extension GE Port Health statistics. For more information about these rules, see Default Monitoring Rules.

If the default policy does not meet your requirements, you can clone an existing default policy and add your custom rule. When you create a custom rule, you need to account for the constraints that are associated with each rule parameter. To obtain detailed information, see Monitoring Systems Support Matrix.

Rule Data

The following example shows one of the default rules for monitoring the Extension GE port health. This rule triggers the configured actions when the number of GE_CRC errors exceeds the threshold.

```
Rule Data:
----------
RuleName: defALL_EXT_GE_PORTSCRC_1
Condition: ALL_EXT_GE_PORTS(GE_CRC/MIN>1)
Actions: RASLOG,SNMP,EMAIL
Associated Policies: dflt_moderate_policy, dflt_conservative_policy
```

RASLog Example

The following example shows the RASLog message that is generated when the GE_CRC error occurs.

```
2023/01/11-10:09:03 (GMT), [MAPS-2062], 105, FID 128 | PORT 0/GE6, WARNING, AWING_157-49-10, GE Port ge6,
 Condition=ALL_EXT_GE_PORTS(GE_CRC/min>1), Current Value:[GE_CRC, 50 CRCs], RuleName=defALL_EXT_GE_PORTSCRC_1,
 Dashboard Category=Extension GE Port Health, Quiet Time=None.
```

MAPS Dashboard Example

The following example shows the MAPS dashboard output that displays the Extension GE Port Health category status.

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

Category                 |Today                      |Last 7 days               |
-------------------------------------------------------------------------------
Port Health              |No Errors                  |No Errors                 |
Extension GE Port Health |Out of operating range     |No Errors                 |
```

```
Fru Health               |In operating range        |In operating range        |
Security Violations      |In operating range        |No Errors                 |
Fabric State Changes     |In operating range        |No Errors                 |
Switch Resource          |In operating range        |In operating range        |
Extension Health         |In operating range        |No Errors                 |
Fabric Performance Impact|In operating range        |In operating range        |


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

Category(Violation Count)|RepeatCount|Rule Name                  |Execution Time   |Object           |
Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------------------------------------------
Extension GE Port Health(|1          |defALL_EXT_GE_PORTSCRC_1   |01/11/23 10:09:03|GE Port ge6      |50 CRCs
              |
1)                       |           |                           |                 |                 |
              |
```

FRU Health

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

The following table lists the monitored parameters in this category.

Table 8: FRU Health Category Parameters

Monitoring System	Description
PS_STATE	Describes changes in the state of a power supply (PS).
FAN_STATE	Describes changes in the state of a fan.
BLADE_STATE	Describes changes in the state of a blade.
SFP_STATE	Describes changes in the state of the SFP transceiver (for both FC or FCoE).
WWN	Describes changes in the state of a WWN card.

MAPS provides default rules to monitor the FRU Health statistics. For more information about these rules, see Default Monitoring Rules.

If the default policy does not meet your requirements, you can clone an existing default policy and add your custom rule. When you create a custom rule, you need to account for the constraints that are associated with each rule parameter. To obtain detailed information, see Monitoring Systems Support Matrix.

Rule Data

The following example shows one of the default rules to monitor FRU health. This rule triggers the configured actions when the fan state attains the FAULTY state.

```
RuleName: defALL_FANFAN_STATE_FAULTY
Condition: ALL_FAN(FAN_STATE/NONE==FAULTY)
Actions: RASLOG,SNMP,EMAIL
Associated Policies: dflt_aggressive_policy, dflt_moderate_policy, dflt_conservative_policy, dflt_base_policy
```

RASLog Example

The following example shows the RASLog message that is generated when the fan state is changed to FAULTY.

```
2023/02/13-21:01:28:957314, [MAPS-1002], 7717/2842, FID 128, ERROR, sw0, Fan 2,
 Condition=ALL_FAN(FAN_STATE==FAULTY), Current Value:[FAN_STATE, FAULTY], RuleName=defALL_FANFAN_STATE_FAULTY,
 Dashboard Category=Fru Health.
```

MAPS Dashboard Example

The following example shows the MAPS dashboard output that displays the FRU Health category status.

```
3.1 Summary Report:
===================
Category                 |Today                      |Last 7 days               |
-------------------------------------------------------------------------------
Fru Health               |Out of operating range     |In operating range        |

3.2 Rules Affecting Health:
===========================
Category(Violation Count)|RepeatCount|Rule Name           |Execution Time   |Object |Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------------
Fru Health(1)            |1          |defALL_FANFAN_STATE_|02/13/23 21:01:28|Fan 2  |FAULTY                |
                                      FAULTY
```

Security Violations

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

The following table lists the monitored parameters in this category.

Table 9: Security Violations Category Parameters

Monitoring System	Description
DAYS_TO_EXPIRE	Describes the number of days before a certificate expires. This parameter notifies the user when the threshold is reached. For more information about security certificate monitoring, refer to the section about configuring protocols in the *Brocade Fabric OS Administration Guide*.
EXPIRED_CERTS	Describes the number of expired certifications detected.
SEC_AUTH_FAIL	Describes the number of authentication failures that occur when a port comes up after a reboot.
SEC_CMD	Describes the number of times that commands that are only permitted on the primary Fibre Channel Switch (FCS) were executed (illegal commands).
SEC_DCC	Describes the number of times an unauthorized device attempted to log in to a secure fabric (DCC violations).
SEC_HTTP	Describes the number of HTTP policy violations that occurred on a switch. HTTP policy violations occur when the HTTP port (80) is blocked by the `ipfilter` feature, and the user tries to establish a connection.
SEC_LV	Describes the number of logon violations that occurred when the switch detected a logon failure. A logon failure occurs due to a wrong user ID or a wrong password.
SEC_SCC	Describes the number of SCC violations that occurred when an unauthorized switch tried to join a secure fabric. The WWN of the unauthorized switch appears in the ERRLOG.

Monitoring System	Description
SEC_TELNET	Describes the number of telnet policy violations that occurred on a switch. Telnet policy violations occur when the telnet port (23) is blocked through the `ipfilter` feature, and the user tries to establish a connection.
SEC_TS	Describes the number of time server (TS) violations that occurred when an out-of-synchronization error was detected.

The statistics from the following monitoring systems are not supported on AG mode:

- SEC_SCC
- SEC_DCC
- SEC_TS
- SEC_AUTH_FAIL
- SEC_CMD

MAPS provides default rules to monitor the security violation statistics. For more information about these rules, see Default Monitoring Rules.

If the default policy does not meet your requirements, you can clone an existing default policy and you can add your custom rule. When you create a custom rule, you must account for the constraints that are associated with each rule parameter. To obtain detailed information, see Monitoring Systems Support Matrix.

Rule Data

The following example shows one of the default rules to monitor the security violations on the switch. This rule triggers the configured actions when the security violations occur within 4 minutes.

```
RuleName: defSWITCHSEC_TELNET_4
Condition: SWITCH(SEC_TELNET/MIN>4)
Actions: RASLOG,SNMP,EMAIL
Associated Policies: dflt_conservative_policy
```

RASLog Example

The following example shows the RASLog message that is generated when security violations occur.

```
2023/01/11-22:28:24:615652, [MAPS-1003], 10/10, SLOT 2 | FID 31, WARNING, sw0, Switch,
 Condition=SWITCH(SEC_TELNET/min>4), Current Value:[SEC_TELNET, 6 Violations], RuleName=defSWITCHSEC_TELNET_4,
 Dashboard Category=Security Violations.
```

MAPS Dashboard Example

The following example shows the MAPS dashboard output that displays the rules that affect the health of the security violation.

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

Category(Violation Count)|RepeatCount|Rule Name              |Execution Time   |Object  |Triggered
 Value(Units)|
-------------------------------------------------------------------------------------------------------------------
Security Violations(8)   |4          |defSWITCHSEC_TELNET_4  |01/11/23 22:29:24|Switch  |5 Violations
  |
                         |           |                       |                 |Switch  |6 Violations
  |
```

```
                          |4          |defSWITCHSEC_LV_4      |01/11/23 22:29:24|Switch  |5 Violations
   |
                          |           |                       |                 |Switch  |6 Violations
   |
                          |1          |defSWITCHSEC_HTTP_4    |01/11/23 22:49:48|Switch  |6 Violations
   |
```

Fabric State Changes

The Fabric State Changes category contains areas of potential fabric-related or switch-related problems such as zone changes, fabric segmentation, E_Port down, fabric reconfiguration, domain ID (DID) changes, and fabric logons.

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

Table 10: Fabric State Changes Category Parameters

Monitoring System	Description
BB_FCR_CNT	Describes the number (count) of Fibre Channel routers (FCR) configured in a fabric.
DID_CHG	Describes the number of forced domain ID changes. DID_CHG occurs when there is a conflict of domain IDs in a single fabric.
EPORT_DOWN	Describes the number of times that an E_Port or VE_Port goes down. E_Ports and VE_Ports go down when you remove a cable or an SFP transceiver (where there are SFP transceiver failures or transient errors). E_Ports that are members of the trunk increment the EPORT_DOWN count only when the last E_Port of the trunk goes down.
FAB_CFG	Describes the number of fabric reconfigurations. Fabric reconfiguration occurs during the following events: • 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
FAB_SEG	Describes the cumulative number of segmentation changes. Segmentation changes occur if any 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
FLOGI	Describes the initialization of ports and devices with the fabric (fabric logons).
L2_DEVCNT_PER	Describes the percentage of imported devices in a Fibre Channel fabric relative to the total number of devices supported in the fabric, whether active or not (Layer 2 device count percentage). The switches in a pure Layer 2 fabric do not participate in the metaSAN.
LSAN_DEVCNT_PER	Describes the percentage of active devices in a Fibre Channel router-enabled backbone fabric relative to the maximum number of devices permitted in the metaSAN (LSAN device count percentage). This percentage includes devices imported from any attached edge fabrics.
ZONE_CFGSZ_PER	Describes the used zone configuration size relative to the maximum zone configuration size in the fabric and per chassis.
ZONE_CHG	Describes the number of zone changes. As zoning is a security provision, frequent zone changes can indicate a security breach or weakness. Therefore, zone change messages occur whenever there is a change in zone configurations.

MAPS provides default rules to monitor for the Fabric state changes. For more information about these rules, see Default Monitoring Rules.

If the default policy does not meet your requirements, you can clone an existing default policy and add your custom rule. When you create a custom rule, you need to account for the constraints that are associated with each rule parameter. To obtain detailed information, see Monitoring Systems Support Matrix.

Rule Data

The following example shows one of the default rules to monitor the fabric state changes. This rule triggered the configured actions when the DID changes occurred in a minute.

```
RuleName: defSWITCHDID_CHG_1
Condition: SWITCH(DID_CHG/MIN>1)
Actions: RASLOG,SNMP,EMAIL
Associated Policies: dflt_aggressive_policy, dflt_moderate_policy, dflt_conservative_policy
```

RASLog Example

The following example shows the RASLog messages that are generated when the DID changes occur.

```
2023/01/11-18:54:21:431121, [CONF-1042], 41/14, FID 60, INFO, switch_60, Fabric Configuration Parameter Domain
 changed to 22.

2023/01/11-18:54:40:434370, [FABR-1030], 42/15, FID 60, INFO, switch_60, fabric: Domain 2 (was 22).

2023/01/11-18:54:42:656193, [MAPS-1003], 43/16, FID 60, WARNING, switch_60, Switch, Condition=SWITCH(DID_CHG/
min>1), Current Value:[DID_CHG, 2 Changes], RuleName=defSWITCHDID_CHG_1, Dashboard Category=Fabric State
 Changes.
```

MAPS Dashboard Example

The following example shows the MAPS dashboard output that displays the rules that affect the fabric state changes.

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

Category(Violation Count)|RepeatCount|Rule Name           |Execution Time   |Object   |Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------------
Fabric State Changes(1)  |1          |defSWITCHDID_CHG_1  |01/11/23 18:54:42|Switch   |2 Changes             |

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

Switch Resource

Switch Resource monitoring enables you to monitor the temperature, flash usage, memory usage, and CPU usage of the system.

The following table lists the monitored parameters in this category.

Table 11: Switch Resource Category Parameters

Monitoring System	Description
CPU	Describes the percentage of CPU used. This percentage is calculated by comparing the percentage of CPU consumed with the configured threshold value.
ETH_MGMT_PORT_STATE	Describes the status of the Ethernet management port (Bond0).
FLASH_USAGE	Describes the percentage of storage space used. This percentage is calculated by comparing the percentage of flash space that is consumed with the configured high threshold value.
MEMORY_USAGE	Describes the available memory that monitors all the components of physical memory in comparison to legacy memory monitoring systems.
MEMORY_USAGE_STATE	Describe the state of the available memory that includes the states CRITICAL, WARNING, and IN_RANGE.
TEMP	Describes the ambient temperature inside the switch. The TEMP monitoring system monitors the temperature sensor state; the sensor state is in IN_RANGE and OUT_OF_RANGE.
VTAP_IOPS	Describes the VTAP I/Os per second.
IT_RES_USAGE	Describes the usage percentage of maximum IT flows in the system.
ITL_RES_USAGE	Describes the usage percentage of maximum ITL flows in the system.

MAPS provides default rules to monitor for the switch resources statistics. For more information about these rules, see Default Monitoring Rules.

If the default policy does not meet your requirements, you can clone an existing default policy and add your custom rule. When you create a custom rule, you need to account for the constraints that are associated with each rule parameter. To obtain detailed information, see Monitoring Systems Support Matrix.

Rule Data

The following example shows the default rules that trigger the configured actions when the Ethernet management port state is changed to down.

```
RuleName: defCHASSISETH_MGMT_PORT_STATE_UP
Condition: CHASSIS(ETH_MGMT_PORT_STATE/NONE==UP)
Actions: RASLOG,SNMP,EMAIL
Associated Policies: dflt_aggressive_policy, dflt_moderate_policy, dflt_conservative_policy, dflt_base_policy
RuleName: defCHASSISETH_MGMT_PORT_STATE_DOWN
Condition: CHASSIS(ETH_MGMT_PORT_STATE/NONE==DOWN)
Actions: RASLOG,SNMP,EMAIL
Associated Policies: dflt_aggressive_policy, dflt_moderate_policy, dflt_conservative_policy, dflt_base_policy
```

RASLog Example

The following example shows the RASLog messages that are generated when the Ethernet management port state is changed to UP and DOWN.

```
2023/01/11-23:00:00:713812, [MAPS-1003], 39/38, SLOT 2 | FID 31, WARNING, sw0, Chassis,
 Condition=CHASSIS(ETH_MGMT_PORT_STATE==DOWN), Current Value:[ETH_MGMT_PORT_STATE, DOWN],
 RuleName=defCHASSISETH_MGMT_PORT_STATE_DOWN, Dashboard Category=Switch Resource.

2023/01/11-23:02:00:662660, [MAPS-1003], 40/39, SLOT 2 | FID 31, WARNING, sw0, Chassis,
 Condition=CHASSIS(ETH_MGMT_PORT_STATE==UP), Current Value:[ETH_MGMT_PORT_STATE, UP],
 RuleName=defCHASSISETH_MGMT_PORT_STATE_UP, Dashboard Category=Switch Resource.
```

MAPS Dashboard Example

The following example shows the MAPS dashboard output that displays the rules that affect the health of the Ethernet management port.

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

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


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

Category(Violation Count)|RepeatCount|Rule Name         |Execution Time   |Object   |Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------------
Switch Resource(2)       |1          |defCHASSISETH_MGMT|01/11/23 23:02:00|Chassis  |UP                    |
                         |           |_PORT_STATE_UP    |                 |         |                      |

                         |1          |defCHASSISETH_MGMT|01/11/23 23:00:00|Chassis  |DOWN                  |
                         |           |_PORT_STATE_DOWN  |                 |         |                      |
```

Extension Health

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

The following table lists the parameters in the Extension Health category that are monitored on all Brocade platforms.

Table 12: Extension Health Category Parameters

Monitoring System	Description
CIR_PKTLOSS	Describes the percentage of the total number of FCIP circuit packets that had to be retransmitted.
CIR_STATE	Describes the state of the FCIP circuit that has changed for one of the following reasons: • The circuit has gone offline. • The circuit has come online. • The circuit is faulty.
CIR_UTIL	Describes the percentage of FCIP circuit utilization in the configured period (this can be a minute, or an hour, or a day).
IP_UTIL	Describes the tunnel IP utilization percentage.

Monitoring System	Description
IP_JITTER	Describes the amount of jitter in an IP circuit that is calculated in percentage and only applies to the circuit group.
IP_RTT	Describes the IP circuit round-trip latency, which is an absolute value and only applies to the circuit group.
JITTER	Describes the amount of jitter in an FCIP circuit that is calculated in percentage and only applies to the circuit group.
PKTLOSS	Describes the percentage of the total number of packets that must be re-transmitted in each QoS level. This value applies to the FCIP QoS group only.
RTT	Describes the FCIP circuit round-trip latency, which is an absolute value and only applies to the circuit group.
IP_EXTN_FLOW	Describes the number of LAN IP Extension flows (TCP connections).
STATE_CHG	Describes the count of FCIP tunnel and FCIP QoS state changes that applies to the tunnel group only.
UTIL	Describes the percentage of FCIP utilization that applies to both the tunnel and the tunnel QoS groups.
IP_PKTLOSS	Describes the percentage of the total number of IP circuit packets that had to be retransmitted.

MAPS provides default rules to monitor for the Extension Health statistics. For more information about these rules, see Default Monitoring Rules.

If the default policy does not meet your requirements, you can clone an existing default policy and add your custom rule. When you create a custom rule, you need to account for the constraints that are associated with each rule parameter. To obtain detailed information, see Monitoring Systems Support Matrix.

Rule Data

The following example shows one of the default rules to monitor Extension Health. This rule triggers the configured actions when the number of CRC errors in a minute on the port is greater than 1.

```
RuleName: defALL_EXT_GE_PORTSCRC_1
Condition: ALL_EXT_GE_PORTS(GE_CRC/MIN>1)
Actions: RASLOG,SNMP,EMAIL
Associated Policies: dflt_moderate_policy, dflt_conservative_policy
```

RASLog Example

The following example shows the RASLog message that is generated when the number of CRC errors reaches the threshold.

```
2023/01/11-10:09:03 (GMT), [MAPS-2062], 105, FID 128 | PORT 0/GE6, WARNING, AWING_157-49-10, GE Port ge6,
Condition=ALL_EXT_GE_PORTS(GE_CRC/min>1), Current Value:[GE_CRC, 50 CRCs], RuleName=defALL_EXT_GE_PORTSCRC_1,
Dashboard Category=Extension GE Port Health, Quiet Time=None.
```

MAPS Dashboard Example

The following example shows the MAPS dashboard output that displays the rules that affect the Extension Health category status.

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

Category                |Today                    |Last 7 days              |
```

```
--------------------------------------------------------------------------------
Port Health              |No Errors                |No Errors                |
Extension GE Port Health |Out of operating range   |No Errors                |
Fru Health               |In operating range       |In operating range       |
Security Violations      |In operating range       |No Errors                |
Fabric State Changes     |In operating range       |No Errors                |
Switch Resource          |In operating range       |In operating range       |
Extension Health         |In operating range       |No Errors                |
Fabric Performance Impact|In operating range       |In operating range       |


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

Category(Violation Count)|RepeatCount|Rule Name                 |Execution Time   |Object           |
Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------------------------------------------
Extension GE Port Health(|1          |defALL_EXT_GE_PORTSCRC_1  |01/11/23 10:09:03|GE Port ge6      |50 CRCs
               |
1)                       |           |                          |                 |                 |
               |

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

Fabric Performance Impact

Fabric Performance Impact (FPI) monitors latency due to the congestion or oversubscription conditions, frame loss, and throughput on ports in a fabric using the MAPS policies. The FPI detects the performance impacts automatically and provides alerts based on the severity of the impact that is caused by credit-stalled devices and congested ports.

The following table lists the monitoring systems in this category:

Table 13: Fabric Performance Impact Category Parameters

Monitoring System	Description
BE_LATENCY_IMPACT	Describes the latency condition that occurred on back-end ports. A port or device in the IO_PERF_IMPACT state can negatively impact the overall network performance.
DEV_LATENCY_IMPACT	Describes the moderate latency condition that results in the IO_PERF_IMPACT state. A severe or frame loss condition due to C3TXTO, that is the count of discarded frames due to timeout of the transmission hold time, results in the IO_FRAME_LOSS state.
IT_FLOW	Describes the initiator to target flow ratio. If any device is zoned to 124 or more zones, then the device is not monitored.
RX	Describes the percentage of port bandwidth being used by received and incoming (RX) traffic. For example, if the port speed is 10G and the port receives 5G of data in one second, then the percentage of RX utilization is 50% (5G × 100/(10G × 1 second)). For a controller trunk port, this incoming traffic percentage indicates the RX percentage for the entire trunk.
TX	Describes the percentage of port bandwidth being used by transmitted, outgoing (TX) traffic. For example, if the port speed is 10G and the port sends 5G of data in one second, the percentage of TX utilization is 50% (5G × 100/(10G × 1 second)). For a controller trunk port, this outgoing traffic percentage indicates the TX percentage for the entire trunk.

Monitoring System	Description
DEV_LOGIN_DIST	Describes the port group state. This monitoring system has the following threshold values: • **BALANCED** – All ports in the group have close to an equal number of devices that are connected to them. • **IMBALANCED** – At least one port in the group does not have any connected devices or has more than one connected device but less than the connected devices to the rest of the ports. • **BALANCE_FAILED** – MAPS performed the re_balance action but failed to rebalance the group.
UTIL	Describes the percentage of the individual port (or trunk) bandwidth being used at the time of the most recent poll.
PORT_BANDWIDTH	Describes the port oversubscription states. In addition, this monitoring system has the following threshold values: • **OVERSUBSCRIBED** – This state shows that the port is in the oversubscription condition. • **OVERSUBSCRIPTION_CLEAR** – This state indicates whether the oversubscription state is cleared or not.

MAPS provides default rules for monitoring the FPI statistics. For more information about these rules, see Default Monitoring Rules.

If the default policy does not meet your requirements, you can clone an existing default policy and can add your custom rule. When you create a custom rule, you must account for the constraints that are associated with each rule parameter. To obtain detailed information, see Monitoring Systems Support Matrix.

Rule Data

The following example shows the default rules to monitor the FPI. This rule shows the frequency that a port is counted as oversubscribed.

```
RuleName: defALL_PORTS_OVERSUBSCRIBED
Condition: ALL_PORTS(PORT_BANDWIDTH/NONE==OVERSUBSCRIBED)
Actions: RASLOG,SNMP,EMAIL,FPIN qt=15 min
Associated Policies: dflt_aggressive_policy, dflt_moderate_policy, dflt_conservative_policy, dflt_base_policy
```

RASLog Example

The following example shows the RASLog messages that are generated when a port is counted as oversubscribed, and the oversubscription state is cleared.

```
2023/01/11-03:08:39 (MST), [MAPS-1003], 503, FID 1 | PORT 0/42, WARNING, TYR_1_159, port42, E-Port 42,
 Condition=ALL_PORTS(PORT_BANDWIDTH/NONE==OVERSUBSCRIBED), Current Value:[PORT_BANDWIDTH, OVERSUBSCRIBED,
 (TXQL=14 us, TX=70.2%) ], RuleName=defALL_PORTS_OVERSUBSCRIBED, Dashboard Category=Fabric Performance Impact,
 Quiet Time=15 min.

2023/01/11-03:20:59 (MST), [MAPS-1004], 504, FID 1 | PORT 0/42, INFO, TYR_1_159, port42, E-Port 42,
 Condition=ALL_PORTS(PORT_BANDWIDTH/NONE==OVERSUBSCRIPTION_CLEAR), Current Value:[PORT_BANDWIDTH,
 OVERSUBSCRIPTION_CLEAR], Rule defALL_PORTS_OVERSUBSCRIPTION_CLEAR, Dashboard Category=Fabric Performance
 Impact, Quiet Time=15 min.
```

MAPS Dashboard Example

The following example shows a dashboard output of the `PORT_BANDWIDTH==OVERSUBSCRIBED` state.

```
3.2 Rules Affecting Health:
```

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

Category(Violation Count)|RepeatCount|Rule Name                   |Execution Time    |Object    |Triggered
 Value(Units)|
-------------------------------------------------------------------------------------------------------------------------------
Fabric Performance Impact|10         |defALL_PORTS_OVERSUBSCRIPTI|01/11/23 03:19:59 |E-Port 42 |
OVERSUBSCRIPTION_CLEAR|
                   (102) |           |ON_CLEAR                   |                  |          |
         |
                         |           |                           |                  |E-Port 42 |
OVERSUBSCRIPTION_CLEAR|
                         |           |                           |                  |E-Port 42 |
OVERSUBSCRIPTION_CLEAR|
                         |           |                           |                  |E-Port 42 |
OVERSUBSCRIPTION_CLEAR|
                         |           |                           |                  |E-Port 42 |
OVERSUBSCRIPTION_CLEAR|
                         |10         |defALL_PORTS_OVERSUBSCRIBED|01/11/23 03:19:19 |E-Port 42 |OVERSUBSCRIBED
        |
                         |           |                           |                  |E-Port 42 |OVERSUBSCRIBED
        |
                         |           |                           |                  |E-Port 42 |OVERSUBSCRIBED
        |
                         |           |                           |                  |E-Port 42 |OVERSUBSCRIBED
        |
                         |           |                           |                  |E-Port 42 |OVERSUBSCRIBED
        |
```

The following example shows the congestion dashboard output of the `PORT_BANDWIDTH==OVERSUBSCRIPTION_CLEAR` state.

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

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

DB start time: Fri Mar 31 10:05:59 2023
Time Window: 10:06 - 10:05
Total Credit-Stalled ports:    1
Total Oversubscribed ports:    0

2 Credit-Stall State Frequency Table:
=====================================

Port        |Current Minute State |Frame Loss   |Perf Impact  |Medium   |Low      |Info     |
-------------------------------------------------------------------------------------------
F-Port 95   |Frame Loss           |2            |0            |0        |0        |0        |

3 Oversubscription State Frequency Table:
=========================================

Port        |Current Minute State   |Frequency|
-----------------------------------------------
```

IO Latency

The IO Latency category monitors metrics that are associated with the I/O operations between an initiator and a target.

The following monitoring systems monitor IO latency statistics for each flow.

Table 14: IO Latency Category Parameters – Flow

Monitoring System	Description
MAX_RD_PENDING_IO	Describes the maximum read pending I/O per ITL/ITN/VITL/VITN on target ports.
MAX_WR_PENDING_IO	Describes the maximum write pending I/O per ITL/ITN/VITL/VITN on target ports.
RD_1stDATA_TIME	Describes the total completion time for the read violations percentage.
RD_STATUS_TIME	Describes the total completion time for the read I/O request.
RD_PENDING_IO	Describes the number of pending read I/O requests.
WR_1stXFER_RDY	Describes the time that the target responds to the first data transfer violation percentage in the ready state.
WR_STATUS_TIME	Describes the total completion time for the write I/O request.
WR_PENDING_IO	Describes the number of pending write I/O requests.

MAPS provides default rules for monitoring the IO Latency statistics. For more information about these rules, see Default Monitoring Rules.

If the default policy does not meet your requirements, you can clone an existing default policy and can add your custom rule. When you create a custom rule, you must account for the constraints that are associated with each rule parameter. To obtain detailed information, see Monitoring Systems Support Matrix.

Rule Data

The following example shows the default rules to monitor the I/O latency. This rule triggers the configured actions when the first data transfer for a write I/O request reaches greater than 1750 in a 10-seconds polling period.

```
Rule Data:
----------
RuleName: defWR_1stXFER_RDY_17500
Condition: sys_flow_monitor_nvme(WR_1stXFER_RDY/10SEC>1750)
Actions: RASLOG,SNMP,EMAIL qt=10 min
Associated Policies: dflt_aggressive_policy
```

RASLog Example

The following example shows the RASLog message that is generated when the WR_1stXFER_RDY for a write I/O request reaches 1750 in a 10-seconds polling period.

```
2023/08/02-00:24:19 (MDT), [MAPS-2982], 125, FID 128 | PORT 1/41, WARNING, TYR2_1_41, Flow
 (SID=0x290600,DID=0x718d00,Host Port=6), Condition=sys_flow_monitor_nvme(WR_1stXFER_RDY/10SEC>1750), Current
 Value:[WR_1stXFER_RDY, 20822 Microseconds], RuleName=defWR_1stXFER_RDY_17500, Dashboard Category=IO Latency,
 Quiet Time=10 min.
```

MAPS Dashboard Example

The following example shows the MAPS dashboard output that displays the rules that affect the IO Latency category status.

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

Category(Violation|RepeatCount|Rule Name               |Execution Time   |Object           |Triggered Value
 |
            Count)                                                                                 (Units)
-------------------------------------------------------------------------------------------------------------
IO Latency(281)   |4          |defWR_1stXFER_RDY_17500 |08/02/23 00:24:59|Flow (SID=0x29060|18K Microseconds
 |
                  |           |                        |                 |0,DID=0x718d00,Ho|
 |
                  |           |                        |                 |st Port=6)       |
 |

                  |           |                        |                 |Flow (SID=0x29060|20.8K
 Microseconds|
                  |           |                        |                 |0,DID=0x718d00,Ho|
 |
                  |           |                        |                 |st Port=6)       |
 |
                  |39         |defWR_1stXFER_RDY_17500 |08/01/23 23:19:59|Flow (SID=0x29070|17.7K
 Microseconds|
                  |           |                        |                 |0,DID=0x718d00,Ho|
 |
                  |           |                        |                 |st Port=7)       |
 |
                  |           |                        |                 |Flow (SID=0x29070|18.3K
 Microseconds|
                  |           |                        |                 |0,DID=0x718d00,Ho|
 |
                  |           |                        |                 |st Port=7)       |
 |
3.2 Rules Affecting Health:
===========================

Category(Violation Count)|RepeatCount|Rule Name                  |Execution Time   |Object           |
Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------------
IO Latency(281)          |4          |defWR_1stXFER_RDY_17500    |08/02/23 00:24:59|Flow (SID=0x29060|18K
 Microseconds      |
                          |           |                           |                 |0,DID=0x718d00,Ho|
               |
                          |           |                           |                 |st Port=6)       |
               |

                          |           |                           |                 |Flow (SID=0x29060|20.8K
 Microseconds    |
                          |           |                           |                 |0,DID=0x718d00,Ho|
               |
                          |           |                           |                 |st Port=6)       |
               |
                          |39         |defWR_1stXFER_RDY_17500    |08/01/23 23:19:59|Flow (SID=0x29070|17.7K
 Microseconds    |
```

```
                          |          |                          |                |0,DID=0x718d00,Ho|
                |
                          |          |                          |                |st Port=7)       |
                |
                          |          |                          |                |Flow (SID=0x29070|18.3K
  Microseconds    |
                          |          |                          |                |0,DID=0x718d00,Ho|
                |
                          |          |                          |                |st Port=7)       |
                |
```

IO Health

The IO Health category monitors the health of the SCSI/NVME flow exceptions and their associated responses.

The following table lists the monitoring systems in the IO Health category.

Table 15: IO Health Category Parameters

Monitoring System	Description
IO_TIMEOUT	Describes the SCSI/NVMe Timeout metric that occurs when no associated first response, status, or abort frame is received for that I/O frame.
IO_ABORT	Describes the frames that are issued by servers to abort the specific exchange in a congested fabric. You can identify the issues on the hosts and targets by monitoring the number of abort sequence (ABTS) frames.
RESERVE	Describes the SCSI/NVME reserve frames to detect if there are too many reservations going through a port that might affect the performance.
IO_ERROR	Describes the status of the SCSI/NVMe errors.

MAPS provides default rules for monitoring the IO Health statistics. For more information about these rules, see Default Monitoring Rules.

If the default policy does not meet your requirements, you can clone an existing default policy and you can add your custom rule. When you create a custom rule, you must account for the constraints that are associated with each rule parameter. To obtain detailed information, see Monitoring Systems Support Matrix.

Rule Data

The following example shows one of the default rules to monitor the I/O errors. This rule triggers the configured actions when the number of I/O errors for an I/O exchange is greater than 100 in a 10-second polling period.

```
Rule Data:
----------
RuleName: def_IO_ERROR
Condition: sys_flow_monitor(IO_ERROR/10SEC>100)
Actions: RASLOG,SNMP,EMAIL qt=10 min
Associated Policies: dflt_aggressive_policy, dflt_moderate_policy, dflt_conservative_policy
```

RASLog Example

The following example shows the RASLog message that is generated when the I/O error count for a flow is greater than 100 in a 10-second polling period.

```
2023/01/11-06:00:34 (MDT), [MAPS-1003], 889, SLOT 2 | FID 1, WARNING, X78_1_113, Flow
 (SID=0x2def80,DID=0x718540,LUN=0,Target Port=9/5), Condition=sys_flow_monitor(IO_ERROR/10SEC>100), Current
 Value:[IO_ERROR, 116 IO-ERRORS], RuleName=def_IO_ERROR, Dashboard Category=IO Health, Quiet Time=10 min.

2023/01/11-06:00:34 (MDT), [MAPS-1003], 890, SLOT 2 | FID 1, WARNING, X78_1_113, Flow
 (SID=0x2def80,DID=0x718540,Target Port=9/5), Condition=sys_flow_monitor(IO_ERROR/10SEC>100), Current Value:
[IO_ERROR, 116 IO-ERRORS], RuleName=def_IO_ERROR, Dashboard Category=IO Health, Quiet Time=10 min.
```

MAPS Dashboard Example

The following example shows the MAPS dashboard output that displays the rules that affect the IO Health category status.

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

Category(Violation Count)|RepeatCount|Rule Name            |Execution Time     |Object           |Triggered
 Value(Units)   |
-------------------------------------------------------------------------------------------------------------------
IO Health(18)            |18         |def_IO_ERROR        |01/11/23 06:01:54  |Flow (SID=0x2def8|1289 IO-
ERRORS           |
                         |           |                    |                   |0,DID=0x718540,Ta|
              |
                         |           |                    |                   |rget Port=9/5    |
              |
                         |           |                    |                   |Flow (SID=0x2def8|1289 IO-
ERRORS           |
                         |           |                    |                   |0,DID=0x71854L   |
              |
                         |           |                    |                   |UN=0,Target Port=|
              |
                         |           |                    |                   |9/5)             |
              |
```

Switch Status Policy

The Switch Status Policy category lets you monitor the health of the switch by defining the number of error types that will transition the overall switch state into a not healthy state. For example, you can specify a switch status policy where a power supply failure puts the switch in a marginal state, and if there are two power supply failures the switch is put in a critical (down) state.

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

Table 16: Switch Status Policy Category Parameters

Monitoring System	Description
BAD_FAN	Describes the number of problematic fans including missing fans and faulty fans.
BAD_PWR	Describes the number of times a power (PWR) supply threshold detected (absent or faulty) and power supplies that were not in the correct slot for redundancy.
BAD_TEMP	Describes the number of times temperature thresholds detected faulty temperature sensors.
DOWN_CORE	Describes the number of faulty (down) core blades (applies to modular switches only).

Monitoring System	Description
ERR_PORTS	Describes the percentage of ports with errors, including the following port types: • Ports that are fenced and decommissioned by MAPS or SANnav Management Portal • Loopback ports • Segmented ports • Ports that are segmented due to security violations
EXPIRED_CERTS	Describes the count of expired certificates.
FAULTY_BLADE	Describes the number of faulty (down) blades (applies to modular switches only).
FAULTY_PORTS	Describes the percentage of hardware-related port faults.
FLASH_USAGE	Describes the percentage of flash usage by the system.
HA_SYNC	Indicates whether the system is in or out of sync.
MARG_PORTS	Describes the percentage of physical ports, E_Ports, and F_Ports (optical and copper) that exceed threshold settings. Whenever these thresholds are persistently high, the port does not have sufficient credits to operate.
SYSTEM_TEMP	Describes the temperature of the system.
MISSING_SFP	Describes the percentage of ports that are missing SFP media.
WWN_DOWN	Describes the number of faulty (down) WWN cards (applies to modular switches only).
FAN_AIRFLOW_MISMATCH	Describes the mismatched air flow of fans.
MARG_SFPS	Describes the percentage of SFPs that exceed threshold settings.

Marginal ports, faulty ports, error ports, and missing SFP transceivers are calculated as a percentage of the physical FC ports. This calculation excludes logical ports, FCoE_Ports, and VE_Ports.

MAPS provides default rules to monitor the FPI statistics. For more information about these rules, see Default Monitoring Rules.

If the default policy does not meet your requirements, you can clone an existing default policy and can add your custom rule. When you create a custom rule, you must account for the constraints that are associated with each rule parameter. To obtain detailed information, see Monitoring Systems Support Matrix.

Rule Data

The following example shows one of the default rules to monitor the switch status policy. This rule shows the rule that enables the switch to be in a marginal state due to the temperature, faulty blade, and faulty port.

```
RuleName: defCHASSISSYSTEM_TEMP_MARG
Condition: CHASSIS(SYSTEM_TEMP/NONE==MARG_OUT_OF_RANGE)
Actions: SNMP,EMAIL,SW_MARGINAL
Associated Policies: dflt_aggressive_policy, dflt_moderate_policy, dflt_conservative_policy, dflt_base_policy
```

RASLog Example

The following example shows the RASLog message that is generated when a faulty blade or port occurs on a switch.

```
2023/01/11-22:01:31:010511, [MAPS-1021], 20651/8321, SLOT 2 | FID 70, WARNING, switch_70,
 RuleName=defCHASSISSYSTEM_TEMP_MARG, Condition=CHASSIS(SYSTEM_TEMP==MARG_OUT_OF_RANGE), Obj:Chassis
 [ SYSTEM_TEMP,MARG_OUT_OF_RANGE] has contributed to switch status MARGINAL.
```

MAPS Dashboard Example

The following example shows the MAPS dashboard output that displays the rules that affect the Switch Status Policy category status.

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

Current Switch Policy Status: CRITICAL
Contributing Factors:
---------------------
*FAULTY_PORTS (CRITICAL).
*FAULTY_BLADE (MARGINAL).
*SYSTEM_TEMP (MARGINAL).
```

Uncategorized Monitoring Systems

The following table lists the monitoring systems that are not associated with any MAPS monitoring category.

Table 17: Uncategorized Monitoring Systems

Monitoring System	Description
ASC_UPLOAD_FAILURE	Describes the Active Support Connectivity (ASC) that monitors the data upload failure.
AVG_PENDING_IOS	Describes the average pending I/Os of the target port.
MAX_PENDING_IOS	Describes the maximum pending I/Os of the target port.
MAX_ROS	Describes the maximum read oversubscription.
AVG_ROS	Describes the average read oversubscription.
TRUFOS_CERT_INSTALLED	Indicates the warning message when an invalid TruFOS certificate is installed or not installed.
TRUFOS_CERT_DAYS_TO_EXPIRE	Indicates the warning message before the TruFOS certificate expiration. An expiration warning is sent out each week 60 days before the expiration date.
TRUFOS_CERT_EXPIRED	Indicates that the support license is expired.

MAPS provides default rules for monitoring the IO Latency statistics. For more information about these rules, see Default Monitoring Rules.

If the default policy does not meet your requirements, you can clone an existing default policy and add your custom rule. When you create a custom rule, you must account for the constraints that are associated with each rule parameter. To obtain detailed information, see Monitoring Systems Support Matrix.

Rule Data

The following example shows the default rules for monitoring the ASC_UPLOAD_FAILURE. This rule shows the frequency that a data upload failure exceeded the specified threshold value.

```
Rule Data:
----------
RuleName: defSWITCHASC_UPLOAD_FAILURE_1
Condition: CHASSIS(ASC_UPLOAD_FAILURE/WEEK>=1)
Actions: RASLOG,SNMP,EMAIL
Associated Policies: dflt_aggressive_policy
```

RASLog Example

The following example shows the RASLog message that is generated when a data upload failure occurs.

```
2023/01/11-00:00:00, [MAPS-1003], 10, FID 128, WARNING, Wedge_37, Chassis,
 Condition=CHASSIS(ASC_UPLOAD_FAILURE/week>=1), Current Value:[ASC_UPLOAD_FAILURE, 2],
 RuleName=defSWITCHASC_UPLOAD_FAILURE_1, Dashboard Category=N/A.
```

MAPS Dashboard Example

The Active Support Connectivity (ASC) monitoring system is not associated with any MAPS dashboard category.

MAPS Groups and Objects

This section describes the MAPS groups and objects (group types).

MAPS Groups Overview

A MAPS group is a collection of similar objects that you can monitor as a single entity using a common threshold. You can create a group of objects and then use that group in rules, thus simplifying rule configuration and management. For example, you can create a group of FC ports and then create specific rules for monitoring a group.

MAPS supports the following groups:

- **Predefined Groups** – Each monitoring category has predefined system groups. For example, ALL_FAN, ALL_PS, and ALL_PORTS. These system groups cannot be modified, except for port groups. To get the complete port group to modify, see MAPS Default Groups Support Matrix.
- **User-Defined (Custom) Groups** – MAPS allows you to create a user-defined group and then use the same group in rules to simplify the rule configuration and management. Some monitoring categories such as Port Health, FCIP Health, and Fabric Performance Impact (FPI) allow user-defined custom groups to be created. You can specify groups that are defined by the characteristics you select. For example, you might need groups of elements that are more suitable for your environment than the predefined groups. User-defined groups contain the following subgroups:
  - **Dynamic Groups** – You can create a group to monitor, where membership is determined by meeting a filter value. Dynamic groups are only used to group ports or devices. Dynamic group contains the following feature names that you can be used when you create a group:
    - portname
    - nodeWWN

    The following example creates the portname for the attached devices in the `portname_grp` group and the portnames start with slot7:

    ```
    switch:admin> logicalGroup --create portname_grp -type PORT -feature portname -pattern "slot7*"
    ```

    The following example creates the nodeWWN for the attached devices in the `nodewwn_grp` group, and the nodeWWNs start with 10:00:

    ```
    switch:admin> logicalGroup --create nodewwn_grp -type PORT -feature nodewwn -pattern "10:00*"
    ```
  - **Static Groups** – You can use a static definition to create a group that can be monitored. The membership in this group is explicit and only changes if you redefine the group.

For more information about the user-defined groups, see User-Defined Groups.

MAPS Objects (Group Types)

MAPS contains multiple objects within a category that is monitored by a rule. A set of similar objects that combine to form a single entity using a common threshold is referred to as a *group* or *logical group*. For example, you can combine all the small form-factor pluggable (SFP) transceivers for the ALL_SFP group. Then you create specific rules for monitoring this group which facilitates the effective monitoring of the SFP ports.

The MAPS object view describes the types of MAPS groups and elements that can be monitored using MAPS. For example, you can use the `switchshow` command to view the switch ports, connections, and status of all ports without MAPS. The following example shows the switch ports and connections using the `switchshow` command.

```
switch:admin> switchshow
switchName:     fid_128
switchType:     180.0
switchState:    Online
switchMode:     Native
```

```
switchRole:     Principal
switchDomain:   1
switchId:       fffc01
switchWwn:      10:00:00:04:71:dx:c0:cd
zoning:         OFF
switchBeacon:   OFF
FC Router:      OFF
HIF Mode:       OFF
Allow XISL Use: OFF
LS Attributes:  [FID: 128, Base Switch: No, Default Switch: Yes, Ficon Switch: No, Address Mode 0]
Index Slot Port Address Media  Speed        State    Proto
==========================================================
  96    4    0   016000   --    N32       No_Module   FC
  97    4    1   016100   --    N32       No_Module   FC
  98    4    2   016200   --    N32       No_Module   FC
  99    4    3   016300   --    N32       No_Module   FC

    <The output is truncated to minimize the content>

1805   -1  1805   ------   --    --       Offline    FCoE
1806   -1  1806   ------   --    --       Offline    FCoE
1807   -1  1807   ------   --    --       Offline    FCoE
```

The following sections provide more information about the types of MAPS objects (group types) that are available for monitoring.

FC Front-End Ports

The FC Front-End ports object type allows you to monitor the FC physical ports. You can use the following monitoring systems to monitor the FC Front-End ports:

- C3TXTO
- RX
- TX
- UTIL
- DEV_LATENCY_IMPACT
- PORT_BANDWIDTH
- ROUTING_ERR
- DEV_LOGIN_DIST
- MAX_RD_PENDING_IO
- MAX_WR_PENDING_IO
- CRC
- DEV_NPIV_LOGINS
- ITW
- LF
- LOSS_SIGNAL
- LOSS_SYNC
- LR
- PE
- PWR_HRS
- STATE_CHG
- WR_1stXFER_RDY_VIOL
- RD_1stDATA_TIME_VIOL
- WR_STATUS_TIME_VIOL
- RD_STATUS_TIME_VIOL

The following table lists the predefined groups for the FC Front-End ports.

Table 18: Predefined Groups for the FC Front-End Ports Object Type

Predefined Group Name	Description
ALL_PORTS	All ports in the logical switch.
ALL_D_PORTS	All D_Ports in the logical switch.
ALL_E_PORTS	All E_Ports and EX_Ports in the logical switch. This group also includes all the ports in E_Port, EX_Port trunks, and AE ports.
ALL_F_PORTS	All F_Ports in the logical switch. This group also includes all the ports in the F_Port trunks. FCLAG ports are monitored as a part of the ALL_F_PORTS and ALL_OTHER_F_PORTS groups.
ALL_OTHER_F_PORTS	All F_Ports in the logical switch that is not a Host or a Target port. FCLAG ports are monitored as a part of the ALL_F_PORTS and ALL_OTHER_F_PORTS groups.
NON_E_F_PORTS	All ports in the logical switch that are not E_Ports, N_Ports, or F_Ports.
ALL_QUARANTINED_PORTS	All ports in the logical switch that have been quarantined for the slow-drain performance.
SFP_STATE	Any port where the state of the SFP transceiver has changed (for both FC or FCoE).

The following example shows a predefined group with a specific list of group members monitored by MAPS.

```
switch:admin> logicalgroup --show NON_E_F_PORTS
-----------------------------------------------------------------------------------------------------
```

```
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
NON_E_F_PORTS                   |Yes        |Port             |353          |3/1-7,3/10-15,3/17-27,
3/34-37,3/40-47,4/2,4/5-42,4/44-46,6/0-22,6/24-47,7/8-55,7/60-63,8/0-7,8/12-31,8/36-55,8/60-63,10/0-47,
11/4-63,12/0-15
```

The groups listed in the previous table can be used to monitor errors or statistics for the FC Front-End ports. For more information, see Port Health.

FC Back-End Ports

The FC Back-End Ports object type allows you to monitor the back-end FC physical ports. You can use the following monitoring systems to monitor the ports:

- BAD_OS
- BE_LATENCY_IMPACT
- CRC
- FRM_LONG
- FRM_TRUNC
- ITW
- LR

The following table lists the predefined groups monitored by categorized by the FC Back-End Ports object type.

Table 19: Predefined Groups for the FC Back-End Ports Object Type

Predefined Group Name	Description
ALL_BE_PORTS	All back-end ports in the physical switch.

The following example displays the group-specific to FC Back-End ports with the complete group members monitored by MAPS.

```
switch:admin> logicalgroup --show ALL_BE_PORTS
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_BE_PORTS                    |Yes        |BE Port          |288          |4/32-34,4/36-37,4/39,
4/41-46,4/48-50,4/52-54,4/56-58,4/61-63,4/96-98,4/100-101,4/103,4/105-110,4/112-114,4/116-118,
4/120-122,4/125-127,5/0-47,5/64-111,6/0-47,6/64-111,8/17-22,8/24,8/26-27,8/29-34,8/36-37,8/39,
8/41-46,8/81-86,8/88,8/90-91,8/93-98,8/100-101,8/103,8/105-110
```

For more information, see BE Port Health.

Gigabit Ethernet Ports

The Gigabit Ethernet (GE) Ports object type monitors statistics for GE ports and takes action based on the configured thresholds and actions. The following monitoring systems can be used to monitor the GE ports:

- GE_CRC
- GE_LOS_OF_SIG

The following table lists the predefined groups monitored by categorized by the GE Ports object type.

Table 20: Predefined Groups for the GE Port Object Type

Predefined Group Name	Description
ALL_EXT_GE_PORTS	All GE ports in the chassis.

The following example displays the group-specific to GE ports with the complete list of ALL_EXT_GE_PORTS group members monitored by the MAPS.

```
switch:admin> logicalgroup --show ALL_EXT_GE_PORTS
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_EXT_GE_PORTS                |Yes        |GE Port          |8            |ge0-ge7
```

For more information, see Extension GE Port Health.

SFP Ports

The SFP Ports object type allows you to monitor SFP-related parameters. The following monitoring systems can be used to monitor SFPs:

- Voltage
- Current
- RXP
- TXP
- PWR_HRS
- SFP_Temp

The following table lists the predefined groups monitored by the SFP Ports object type:

Table 21: Predefined Groups for the SFP Ports Object Type

Predefined Group Name	Description
ALL_SFP	All SFP transceivers.
ALL_10GSWL_SFP	All 10G short wavelength (SWL) SFP transceivers on FC ports in the logical switch.
ALL_10GLWL_SFP	All 10G long wavelength (LWL) SFP transceivers on FC ports in the logical switch.
ALL_16GSWL_SFP	All 16G SWL SFP transceivers in the logical switch.
ALL_16GLWL_SFP	All 16G LWL SFP transceivers in the logical switch.
ALL_25Km_16GLWL_SFP	All 25-kilometer-capable 16G LWL SFP transceivers in the logical switch.
ALL_25Km_32GELWL_JBA_SFP	All 25-kilometer-capable 32G extended long wavelength (ELWL) JBA SFP transceivers in the logical switch.
ALL_40Km_16GELWL_SFP	All 40-kilometer-capable 16G ELWL SFP transceivers in the logical switch.
ALL_25Km_32GELWL_SFP	All 25-kilometer-capable 32G ELWL SFP transceivers in the logical switch.
ALL_25Km_64GELWL_SFP	All 25-kilometer-capable 64G ELWL SFP transceivers in the logical switch.
ALL_40Km_32GELWL_SFP	All 40-kilometer-capable 32G ELWL SFP transceivers in the logical switch.
ALL_32GSWL_SFP	All 32G SWL SFP transceivers in the logical switch.
ALL_32GLWL_SFP	All 32G LWL SFP transceivers in the logical switch.
ALL_32GLWL_JDB_SFP	All 32G LWL JDB SFP transceivers in the logical switch.

Predefined Group Name	Description
ALL_32GSWL_QSFP	All 4 X 32G SWL quad SFP (QSFP) transceivers in the logical switch.
ALL_64GLWL_SFP	All 64G LWL SFP transceivers in the logical switch.
ALL_64GSWL_SFP	All 64G SWL SFP transceivers in the logical switch.
ALL_64GSWL_SFP_DD	All 64G SWL SFP transceivers in the logical switch that help to create a default group for SFP-DD SFPs.
ALL_OTHER_SFP	All SFP transceivers that do not belong to any other groups except ALL_SFP.
ALL_2K_QSFP	All 2-kilometer-capable 16G QSFP transceivers used for inter-chassis link (ICL) connections in the logical switch.
ALL_80Km_8GELWL_SFP	All 80-kilometer-capable 8G ELWL SFP transceivers in the logical switch to monitor 8G DWDM Smart Optics SFPs.
ALL_100M_16GSWL_QSFP	All 4 X 32G SWL QSFP transceivers in the logical switch.
ALL_2Km_32GLWL_QSFP	All 2-kilometer-capable 4 X 32G LWL QSFP transceivers in the logical switch.
ALL_2Km_GEN7LWL_QSFP	All 2-kilometer-capable Gen 7 LWL QSFP transceivers in the logical switch.
ALL_GEN7SWL_QSFP	All Gen 7 SWL QSFP transceivers in the logical switch.
ALL_FCOE_40G_QSFP	All 40GE SWL QSFP transceivers in the logical switch.
ALL_FCOE_100G_SR4_QSFP	All 100GE SWL QSFP transceivers in the logical switch.

The following example displays the SFP ports with the complete list of ALL_SFP group members monitored by MAPS:

```
switch:admin> logicalgroup --show ALL_SFP
-------------------------------------------------------------------------------------------------------------
Group Name    |Predefined |Type    |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_SFP       |Yes        |Sfp     |15           |3/48,3/50,3/56-59,7/1,7/13,7/20-21,7/28,7/30,7/36,8/8,8/19
```

For more information, see Port Health.

Temperature Sensor

The Temperature Sensor object type allows you to monitor the temperature of the system. The TEMP monitoring system can be used to monitor the temperature.

The following table lists the predefined groups monitored by the Temperature Sensor object type.

Table 22: Predefined Groups for the Temperature Sensor Object Type

Predefined Group Name	Description
ALL_TS	All temperature sensors in the system.

The following example displays the predefined groups in the ALL_TS object type:

```
switch:admin> logicalgroup --show ALL_TS
-------------------------------------------------------------------------------------------------
Group Name                    |Predefined |Type              |Member Count |Members
-------------------------------------------------------------------------------------------------
ALL_TS                        |Yes        |Temperature sensor|9            |0-8
```

For more information, see Switch Status Policy.

ASIC

The application-specific integrated circuit (ASIC) object type allows you to monitor the ASICs of the switch. The VTAP_IOPS monitoring system can be used to monitor the ASIC platforms.

The following table lists the predefined groups monitored by the ASIC object type.

Table 23: Predefined Groups for the ASIC Object Type

Predefined Group Name	Description
ALL_ASICS	All ASICs exist on the system. The content of the group depends on the system and blades.

The following example displays the predefined group with the specific list of ALL_ASICS group members monitored by the MAPS.

```
switch:admin> logicalgroup --show ALL_ASICS
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_ASICS                       |Yes        |Asic             |11           |3/0-1,4/0,5/0-1,6/0-1,7/0-1
```

For more information, see Switch Resource.

Blade

The Blade object type allows you to monitor the state of blades. The following monitoring systems can be used to monitor the blades:

- BLADE_STATE
- DOWN_CORE
- HA_SYNC

DOWN_CORE and HA_SYNC are monitored using the chassis group. However, DOWN_CORE monitors the number of bad core blades, and HA_SYNC monitors the state of the CP blades.

The following table lists the predefined groups monitored by the Blade object type.

Table 24: Predefined Groups for the Blade Object Type

Predefined Group Name	Description
ALL_SW_BLADES	All port and application blades in the chassis.
ALL_CORE_BLADES	All core blades in the chassis.

The following example displays the predefined group with the specific list of ALL_SW_BLADES group members monitored by the MAPS.

```
switch:admin> logicalgroup --show ALL_SW_BLADES
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_SW_BLADES                   |Yes        |Blade            |4            |3,4,7,8
```

The following example displays the predefined group with the specific list of ALL_CORE_BLADES group members monitored by the MAPS.

```
switch:admin> logicalgroup --show ALL_CORE_BLADES
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_CORE_BLADES                 |Yes        |Blade            |2            |5,6
```

For more information, see Port Health.

Fan

The Fan object type allows you to monitor the state of fans. The following monitoring systems can be used to monitor the fans:

- FAN_STATE
- BAD_FAN
- FAN_AIRFLOW_MISMATCH

NOTE
BAD_FAN and FAN_AIRFLOW_MISMATCH can be monitored using the chassis group.

The following table lists the predefined groups monitored by the Fan object type.

Table 25: Predefined Groups for the Fan Object Type

Predefined Group Name	Object Type	Description
ALL_FAN	Fan	All fans in the chassis.

The following example displays the predefined group with the specific list of ALL_FAN group members monitored by the MAPS.

```
switch:admin> logicalgroup --show ALL_FAN
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_FAN                         |Yes        |Fan              |3            |1,2,3
```

For more information, see Switch Status Policy and FRU Health.

Flash

The Flash object type allows you to monitor the state of the flash memory cards, the percentage of storage space used, and the available memory. The following monitoring systems can be used to monitor the flash memory cards:

- FLASH_USAGE
- MEMORY_USAGE

The following table lists the predefined groups monitored by the Flash object type.

Table 26: Predefined Groups for the Flash Object Type

Predefined Group Name	Description
ALL_FLASH	The flash memory card in the system.

The following example displays the predefined group with the specific list of ALL_FLASH group members monitored by the MAPS.

```
switch:admin> logicalgroup --show ALL_FLASH
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_FLASH                       |Yes        |Flash            |1            |0
```

Power Supply

The Power Supply object type allows you to monitor the state of power supplies. You can use the PS_STATE monitoring system to monitor the power supplies.

The following table lists the predefined groups monitored by the power supply object type.

Table 27: Predefined Groups for the Power Supply Object Type

Predefined Group Name	Object Type	Description
ALL_PS	Power Supply	All power supplies in the system.

The following example shows a predefined group with the specific list of ALL_PS group members monitored by MAPS.

```
switch:admin> logicalgroup --show ALL_PS
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_PS                          |Yes        |Power Supply     |4            |1,2,3,4
```

Certificate

The Certificate object type allows you to monitor the state of imported certificates. You can use the following monitoring systems to monitor the certificates:

- DAYS_TO_EXPIRE
- EXPIRED_CERTS

MAPS monitors the following Certificate object types:

- HTTPS
- LDAP (TLS client)
- RADIUS
- ASC
- FCAP
- SYSLOG
- KAFKA
- EXTN SW
- EXTN CA

The following table lists the predefined groups that are monitored by the Certificate object type.

Table 28: Predefined Groups for the Certificate Object Type

Predefined Group Name	Description
ALL_CERTS	All imported certificates.

The following example shows a predefined group with a specific list of the ALL_CERTS group members that are monitored by MAPS.

```
switch:admin> logicalgroup --show ALL_CERTS

-----------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-----------------------------------------------------------------------------------------------------------
ALL_CERTS                       |Yes        |Certificate      |14           |HTTPS SW Certificate,HTTPS CA
Certificates, SYSLOG Client certificate,SYSLOG Client CA Certificate,LDAP Client certificate,LDAP Client CA
Certificates,RADIUS Client certificate,RADIUS Client CA Certificate,FCAP SW certificate,FCAP CA certificates,
KAFKA Server CA Certificate, EXTN SW Certificates,EXTN CA Certificate
switch:admin>  logicalgroup --show ALL_CERTS -details

GroupName          : ALL_CERTS
Predefined         : Yes
Type               : Certificate
MemberCount        : 14
Members            : HTTPS SW Certificate,HTTPS CA Certificates,SYSLOG Client certificate,SYSLOG Client CA
                     Certificate,LDAP Client certificate,LDAP Client CA Certificates,RADIUS Client certificate,
                     RADIUS Client CA Certificate,FCAP SW certificate,FCAP CA certificates,KAFKA Server CA
                     Certificate, EXTN SW Certificates, EXTN CA Certificate
Added Members      :
Deleted Members    :
Feature            :
Pattern            :
```

Fabric OS v9.2.x supports monitoring certificates which are multiple certificates of the same type.

For more information, see Security Violations.

Circuit

The Circuit object type allows you to monitor the circuits. The following monitoring systems can be used to monitor the circuits:

- CIR_PKTLOSS
- CIR_STATE
- CIR_UTIL
- IP_JITTER
- IP_UTIL
- IP_PKTLOSS
- JITTER
- RTT
- IP_RTT

The following table lists the predefined groups monitored by the Circuit object type.

Table 29: Predefined Groups for the Circuit Object Type

Predefined Group Name	Description
ALL_CIRCUITS	All the circuits present in the switch.

The following example displays the predefined group with the specific list of ALL_CIRCUITS group members monitored by the MAPS.

```
switch:admin> logicalgroup --show ALL_CIRCUITS
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_CIRCUITS                    |Yes        |Circuit          |4            |14/0-3
```

For more information, refer to Extension Health.

Circuit QoS

The Circuit QoS object type allows you to monitor each QoS per circuit that allows prioritizing traffic flows. You can use the following monitoring systems to monitor the Circuit QoS:

- Packet Loss
- UTIL
- STATE_CHG

The following table lists the predefined groups monitored by the Circuit QoS object type.

Table 30: Predefined Groups for the Circuit QoS Object Type

Predefined Group Name	Description
ALL_CIRCUIT_F_QOS	All control QoS members of a circuit.
ALL_CIRCUIT_HIGH_QOS	All QoS members of a circuit with high levels of traffic.
ALL_CIRCUIT_MED_QOS	All QoS members of a circuit with medium levels of traffic.
ALL_CIRCUIT_LOW_QOS	All QoS members of a circuit with low levels of traffic.
ALL_CIRCUIT_IP_HIGH_QOS	All IP QoS members of a circuit with high levels of traffic.
ALL_CIRCUIT_IP_MED_QOS	All IP QoS members of a circuit with medium levels of traffic.
ALL_CIRCUIT_IP_LOW_QOS	All IP QoS members of a circuit with low levels of traffic.

The following example shows a predefined group with a specific list of ALL_CIRCUIT_F_QOS group members monitored by MAPS.

```
switch:admin> logicalgroup --show ALL_CIRCUIT_F_QOS
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_CIRCUIT_F_QOS               |Yes        |Circuit Qos      |4            |14/0-3
```

The following example shows a predefined group with the specific list of ALL_CIRCUIT_HIGH_QOS group members monitored by MAPS.

```
switch:admin> logicalgroup --show ALL_CIRCUIT_HIGH_QOS
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_CIRCUIT_HIGH_QOS            |Yes        |Circuit Qos      |4            |14/0-3
```

The following example shows a predefined group with the specific list of ALL_CIRCUIT_MED_QOS group members monitored by MAPS.

```
switch:admin> logicalgroup --show ALL_CIRCUIT_MED_QOS
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_CIRCUIT_MED_QOS             |Yes        |Circuit Qos      |4            |14/0-3
```

The following example shows a predefined group with the specific list of ALL_CIRCUIT_LOW_QOS group members monitored by MAPS.

```
switch:admin> logicalgroup --show ALL_CIRCUIT_LOW_QOS
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_CIRCUIT_LOW_QOS             |Yes        |Circuit Qos      |4            |14/0-3
```

For more information, see Extension Health.

ETH Port

The ETH Port object type allows you to monitor the ETH ports. The SFP_STATE monitoring system can be used to monitor the ETH ports.

The following table lists the predefined groups monitored by the ETH Port object type.

Table 31: Predefined Groups for the ETH Port Object Type

Predefined Group Name	Description
ALL_ETH_PORTS	All VF_Ports and VN_Ports in the logical switch. This also includes all the ports in the VF_Port port-channels.

The following example shows a predefined group with the specific list of ALL_ETH_PORTS group members monitored by the MAPS.

```
switch:admin> logicalgroup --show ALL_ETH_PORTS
-------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------
ALL_ETH_PORTS                   |Yes        |Eth Port         |7            |11/36-39,11/60,11/62-63
```

For more information, see Port Health.

Flow

A flow is a set of Fibre Channel (FC) frames or packets that share similar traits, such as the following:

- An ingress port identifier
- An egress port identifier
- Any other data that can be used to differentiate one set of related frames or packets from another set

The following monitoring systems can be used for monitoring the flow:

- RD_STATUS_TIME
- WR_STATUS_TIME
- RD_1stDATA_TIME
- WR_1stXFER_RDY
- RD_PENDING_IO
- WR_PENDING_IO
- IO_TIMEOUT
- IO_ERROR

The following monitoring systems are used for monitoring the flow using the port object types:

- RD_STATUS_TIME_VIOL
- WR_STATUS_TIME_VIOL
- RD_1stDATA_TIME_VIOL
- WR_1stXFER_RDY_VIOL

The following monitoring systems are used for monitoring the ports and flows:

- MAX_RD_PENDING_IO
- MAX_WR_PENDING_IO

The following table lists the predefined groups that are monitored by the *Flow* object type.

Predefined Group Name	Description
sys_flow_monitor	The predefined learning flow to provide Fibre Channel and IO Insight metrics.
sys_flow_monitor_scsi	The automated predefined flow to monitor SCSI flows and metrics.
sys_flow_monitor_nvme	The automated predefined flow to monitor NVMe flows and metrics.

The following example shows the predefined group with the specific list of the FLOW group members that are monitored by the MAPS.

```
switch:admin> logicalgroup --show sys_flow_monitor

-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
sys_flow_monitor                |Yes        |Flow             |4            |Monitored Flow
```

The following example shows the predefined sys_flow_monitor_scsi group with the FLOW group members that are monitored by the MAPS.

```
switch:admin> logicalgroup --show sys_flow_monitor_scsi

-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
sys_flow_monitor_scsi           |Yes        |Flow             |3            |Monitored Flow
```

The following example shows the predefined sys_flow_monitor_nvme group with the FLOW group members that are monitored by the MAPS.

```
switch:admin> logicalgroup --show sys_flow_monitor_nvme

-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
sys_flow_monitor_nvme           |Yes        |Flow             |1            |Monitored Flow
```

Tunnel

The Tunnel object type allows you to monitor each tunnel in a fabric. The following monitoring systems can be used to monitor the tunnels:

- UTIL
- STATE_CHG
- IP_UTIL

The following table lists the predefined groups monitored by the Tunnel object type.

Table 32: Predefined Groups for the Tunnel Object Type

Predefined Group Name	Description
ALL_TUNNELS	All tunnels configured on the system.

The following example shows the predefined group with the specific list of ALL_TUNNELS group members monitored by the MAPS.

```
switch:admin> logicalgroup --show ALL_TUNNELS
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_TUNNELS                     |Yes        |Tunnel           |1            |14
```

For more information, see Extension Health.

Tunnel QoS

The Tunnel QoS object type allows you to monitor each QoS per tunnel. The following monitoring systems can be used to monitor the tunnel QoS:

- PKTLOSS
- UTIL

The following table lists the predefined groups monitored by the Tunnel QoS object type.

Table 33: Predefined Groups for the Tunnel QoS Object Type

Predefined Group Name	Description
ALL_TUNNEL_F_QOS	All control QoS members of a tunnel.
ALL_TUNNEL_LOW_QOS	All QoS members of a tunnel with low levels of traffic.
ALL_TUNNEL_MED_QOS	All QoS members of a tunnel with medium levels of traffic.

Predefined Group Name	Description
ALL_TUNNEL_HIGH_QOS	All QoS members of a tunnel with high levels of traffic.
ALL_TUNNEL_IP_LOW_QOS	All QoS members of a tunnel IP with low levels of traffic.
ALL_TUNNEL_IP_MED_QOS	All QoS members of a tunnel IP with medium levels of traffic.
ALL_TUNNEL_IP_HIGH_QOS	All QoS members of a tunnel IP with high levels of traffic.

The following example shows the predefined group with the specific list of ALL_TUNNEL_F_QOS group members monitored by the MAPS.

```
switch:admin> logicalgroup --show ALL_TUNNEL_F_QOS
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_TUNNEL_F_QOS                |Yes        |Tunnel  Qos      |1            |14
```

The following example shows the predefined group with the specific list of ALL_TUNNEL_LOW_QOS group members monitored by the MAPS.

```
switch:admin> logicalgroup --show ALL_TUNNEL_LOW_QOS
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_TUNNEL_LOW_QOS              |Yes        |Tunnel  Qos      |1            |14
```

The following example shows the predefined group with the specific list of ALL_TUNNEL_MED_QOS group members monitored by the MAPS.

```
switch:admin> logicalgroup --show ALL_TUNNEL_MED_QOS
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_TUNNEL_MED_QOS              |Yes        |Tunnel  Qos      |1            |14
```

The following example shows the predefined group with the specific list of ALL_TUNNEL_HIGH_QOS group members monitored by the MAPS.

```
switch:admin> logicalgroup --show ALL_TUNNEL_HIGH_QOS
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_TUNNEL_HIGH_QOS             |Yes        |Tunnel  Qos      |1            |14
```

The following example shows the predefined group with the specific list of ALL_TUNNEL_IP_LOW_QOS group members monitored by the MAPS.

```
switch:admin> logicalgroup --show ALL_TUNNEL_IP_LOW_QOS


-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type             |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_TUNNEL_IP_LOW_QOS           |Yes        |Tunnel  Qos      |1            |8
```

For more information, see Extension Health.

WWN

The WWN object type allows you to monitor the state of the World Wide Name (WWN) card. The WWN monitoring system can be used to monitor the WWN cards.

The following table lists the predefined groups monitored by the WWN object type.

Table 34: Predefined Groups for the WWN Object Type

Predefined Group Name	Description
ALL_WWN	All WWN cards in the chassis.

The following example shows the predefined group with the specific list of ALL_TUNNELS group members monitored by the MAPS.

```
switch:admin> logicalgroup --show ALL_WWN
-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type              |Member Count |Members
-------------------------------------------------------------------------------------------------------------
ALL_WWN                         |Yes        |WWN               |2            |1,2
```

For more information, see FRU Health and Switch Status Policy.

Switch

The Switch object group type allows you to monitor the switches. You can use the following monitoring systems to monitor the switches:

- BB_FCR_CNT
- DID_CHG
- EPORT_DOWN
- ERR_PORTS
- FAB_CFG
- FAB_SEG
- MISSING_SFP
- FAULTY_PORTS
- FLOGI
- L2_DEVCNT_PER
- LSAN_DEVCNT_PER
- MARG_PORTS
- MARG_SFPS
- SEC_AUTH_FAIL
- SEC_CMD
- SEC_DCC
- SEC_HTTP
- SEC_LV
- SEC_SCC
- SEC_TELNET
- SEC_TS
- ZONE_CFGSZ_PER
- ZONE_CHG

The following table lists the predefined group that is monitored by the switch object type.

Table 35: Predefined Groups for the Switch Object Type

Predefined Group Name	Description
SWITCH	Default group used for defining rules on parameters that are global for the whole switch level, for example, security violations or fabric health.

The following example shows a predefined group with the specific list of SWITCH group members monitored by MAPS.

```
switch:admin> logicalgroup --show SWITCH
-------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type              |Member Count |Members
-------------------------------------------------------------------------------------------------
SWITCH                          |Yes        |                  |1            |0
```

Chassis

The Chassis object group type allows you to monitor the chassis systems. You can use the following monitoring systems to monitor the chassis systems:

- ASC_UPLOAD_FAILURE
- BAD_TEMP
- FAN_AIR_FLOW_MISMATCH
- BAD_FAN
- BAD_PWR
- EXPIRED_CERTS
- DAYS_TO_EXPIRE
- CPU
- DOWN_CORE
- ETH_MGMT_PORT_STATE
- FAULTY_BLADE
- FLASH_USAGE
- HA_SYNC
- IT_RES_USAGE
- ITL_RES_USAGE
- MEMORY_USAGE
- MEMORY_USAGE_STATE
- SYSTEM_TEMP
- WWN_DOWN
- TRUFOS_CERT_DAYS_TO_EXPIRE
- TRUFOS_CERT_EXPIRED
- TRUFOS_CERT_INSTALLED
- ZONE_CFGSZ_PER

The following table lists the predefined group that is monitored by the CHASSIS object type.

Table 36: Predefined Groups for the Switch Object Type

Predefined Group Name	Description
CHASSIS	Default group used for defining rules on parameters that are global for the whole chassis, for example, CPU or flash.

The following example shows a predefined group with the specific list of CHASSIS group members monitored by MAPS.

```
switch:admin> logicalgroup --show CHASSIS

-------------------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type              |Member Count |Members
-------------------------------------------------------------------------------------------------------------
CHASSIS                         |Yes        |                  |1            |0
```

Setting Up MAPS Operations

This section describes the advanced MAPS operations using MAPS default rules and policies.

In the previous sections, you learned about the basics of MAPS and how to apply the predefined MAPS rules and policies to set up basic MAPS monitoring and alerting. To perform basic MAPS operations, you must enable MAPS with the predefined groups and enable the dflt_conservative policy with the RASLog, SNMP, and email actions. You might prefer custom policies if you have stricter monitoring requirements for a high-performance fabric.

If you have stringent monitoring requirements for a high-performance fabric, you might prefer custom policies to meet your requirements. Creating, modifying, and cloning the appropriate MAPS groups, conditions, and policies are critical to setting up the advanced MAPS operations. Use the topics in this section to configure the advanced MAPS operations using default rules and policies or customize policies for specific ports or switch elements.

You can create custom MAPS elements and perform MAPS operations in the following order:

1. Create a policy using the following methods:
   – Clone the existing policy to create a custom policy.
     For example, you can clone dflt_conservative_policy. For more information, see Custom MAPS Policies.
   – Create a policy.
     For more information about the CLI commands that are used to clone MAPS elements, refer to the *Brocade Fabric OS Command Reference Manual*.
2. Create a rule and add it to the custom policy you created in Step 1.
   For more information, see Custom MAPS Rules.
3. Create a custom group by modifying a predefined group.
   For more information, see Custom MAPS Groups.
4. Enable notifications is critical for getting alerts after you create custom MAPS elements. For example, you can enable the FENCE, RASLog, DECOM, and SNMP actions. These MAPS actions help you to control the MAPS behavior. For more information, see Enabling or Disabling Rule Actions at a Global Level and Switch Level Configuration.
5. Activate the newly created custom policy to start the MAPS monitoring.
   To activate the policy, you must enable the custom policy that you created using `mapspolicy --enable policy <custom_policy_name>` for the modified rule to take effect.
6. Create rules to monitor the performance of other rules, that is, rule-on-rule (RoR) rule monitoring. For more information, see Rule on Rule.
7. Pause monitoring during maintenance operations such as device or server upgrades and perform a temporary suspension. For more information, see Pausing and Restarting MAPS.

Custom MAPS Policies

A MAPS policy is a set of rules. Whenyou enable a policy, all rules in the policy are in effect. You can create the policy by cloning an existing default policy, for example, dflt_conservative_policy, as follows:

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

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 rules in the newly created my_mon_policy:

```
switch:admin>  mapspolicy --show my_mon_policy
defALL_HOST_PORTSC3TXTO_10      |ALL_HOST_PORTS(C3TXTO/MIN>10)      |FENCE,SNMP,EMAIL |
defALL_HOST_PORTSC3TXTO_3       |ALL_HOST_PORTS(C3TXTO/MIN>3)       |RASLOG,SNMP,EMAIL|
```

```
defALL_HOST_PORTSCRC_10        |ALL_HOST_PORTS(CRC/MIN>10)        |RASLOG,SNMP,EMAIL|
         --------------------Output truncated--------------------------
defALL_HOST_PORTSTX_75         |ALL_HOST_PORTS(TX/HOUR>75)        |RASLOG,SNMP,EMAIL|
defALL_HOST_PORTSUTIL_75       |ALL_HOST_PORTS(UTIL/HOUR>75)      |RASLOG,SNMP,EMAIL|
```

For more information about CLI commands that are used to clone MAPS elements, refer to the *Brocade Fabric OS Command Reference Manual*.

Custom MAPS Rules

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

The combination of actions, conditions, and groups allow you to create a rule for almost any scenario required by your environment. When your requirements are different from the thresholds specified in the default rules, you can create your own rule and add it to the policy. You can delete the default rules from the custom policy. For example, you can create a custom rule if the 75% threshold defined in the defALL_HOST_PORTSUTIL_75 rule does not meet your requirements.

The following example command creates a custom rule and adds it to your custom policy:

```
mapsrule --clone defALL_HOST_PORTSUTIL_75 -rule my_mon_policy my_mon_rule_4_util_90 -value 90
```

You can delete a rule that is not required. The following example command deletes the rule defALL_HOST_PORTSUTIL_75.

```
mapspolicy --delrule my_mon_policy -rule defALL_HOST_PORTSUTIL_75
```

You can manage newly created rules using the `mapsrule --config` command.

For more information about configuring and managing MAPS monitoring rules, refer to the man pages of the Brocade Fabric OS CLI and the *Brocade Fabric OS Command Reference Manual*.

You can create, modify, or clone a rule using the `mapsrule --create` , `mapsrule --config` , or `mapsrule --clone` commands.

Enabling or Disabling Rule Actions at a Global Level

Allowable actions on a switch can be specified globally, and these global actions supersede any actions specified in individual rules.

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

1. Enter the `mapsconfig --show` command to display the actions that are currently allowed on the switch.
2. Enter the `mapsconfig --actions` command and specify all of the actions that you want to allow on the switch, for example, the `mapsconfig --actions <action1>,<action2>,<action3>...` command (up to the complete set of actions).
3. If you are changing the list of active actions, specify all the actions to be active. For example, if you add RASLog notifications to a switch that already has email notifications enabled, you must specify both email and RASLog as actions in the `mapsconfig` command.

   To disable all actions, enter the `mapsconfig --actions none` command. The keyword `none` cannot be combined with any other action. The following example shows that RASLog, EMail, DECOM, and FENCE notifications are not active actions on the switch. It then shows them being added to the list of allowed actions.

   ```
   switch:admin> mapsconfig --show
   Configured Notifications:     SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,FPIN
   Mail Recipient:               Not Configured
   Mail From Address:            Not Configured
   ```

```
Raslog Mode:                  Default
Decom Action Config:          With Disable
Global Quiet Time:            5 Hours (Not Effective)
Notification Behavior:        Adaptive
Paused members :
================
PORT :
CIRCUIT :
SFP :
SWITCH :

switch:admin> mapsconfig --actions raslog,email,decom,fence
2021/04/16-14:02:04, [MAPS-1130], 856, SLOT 2 | FID 128, INFO, Gen7-8, Actions
 raslog,email,decom,fence configured.

switch:admin> mapsconfig --show
Configured Notifications:       RASLOG,EMAIL,FENCE,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,DECOM
Mail Recipient:                 Not Configured
Mail From Address:              Not Configured
Raslog Mode:                    Custom
Decom Action Config:            Impair (No Disable)
Global Quiet Time:              5 Hours
Notification Behavior:          Adaptive
Paused members :
================
PORT :
CIRCUIT :
SFP :
SWITCH :
```

For the list of MAPS actions, see MAPS Basic Elements.

Email Alerts

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

To configure the email recipients, use the `mapsConfig --emailcfg` command. Multiple destination email addresses are possible; they must be separated using a comma between each address, and each address must be a complete email address. For example, abc@company.com is a valid email address; abc@company is not.

To clear all configured email addresses, enter the `mapsconfig --emailcfg -address none` command. All configured email addresses are erased.

Email Delivery Monitoring

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

Only one email failure RASLog generated within an hour for every configured email address. The RASLog contains the email address of the failed destination. Another RASLog that is sent every hour, which indicates the total number of failures in one hour. Sending only one RASLog avoids generating too many RASLog failure messages if there is an email server problem or network issues.

An email can be sent from a system, but it might not be successfully delivered to the destination email server. In MAPS, emails are sent using the `sendmail` command in Linux. The initial steps to debug the problem are to check the following:

- Whether the network is running
- Whether the routing tables are configured correctly
- Whether DNS is configured correctly

Even when the network connectivity is working properly, the `sendmail` command can have the following failure conditions:

- There are no routes in the system to use to send an email. For example, there is a network error, and TCP/IP service is not available.
- A fatal configuration problem occurs while reading the configuration file. Failure of a delivery agent to function correctly can lead to this kind of failure.
- Problems can result from various operating system errors.
  - The `sendmail` process cannot be forked by the `cron` job script due to a memory problem, and therefore, the command cannot be successfully executed.
  - Other internal software errors.

Examples of RASLog Messages

The following is an example of the RASLog that is generated when the `sendmail` command fails to send an email from the switch.

```
2021/02/01-19:49:00, [MAPS-1206], 1468, SLOT 6 CHASSIS, INFO, dcx_178, A MAPS notification sent from the
 switch to abc@company.com could not be delivered to the mail server.
```

The following is an example of the RASLog that is generated every hour to indicate the number of failures seen in one hour. This RASLog is generated only if more than one failure for any failed address during the one hour.

```
2021/02/02-20:49:00, [MAPS-1206], 1468, SLOT 6 CHASSIS, INFO, dcx_178, There were 12 or more notifications
 that could not be delivered in the last one hour.
```

MAPS SNMP Traps

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

In environments with a high number of messages coming from various switches, you might want to receive them in a single location and view them using a graphical user interface (GUI). In this type of scenario, simple network management protocol (SNMP) notifications can be the most efficient notification method because you can avoid having to log on to each switch individually as you do for error log notifications.

To get the traps, you must configure the SNMP software to receive the trap information from the network device, and you must configure the IP address of the SNMP agent on the switch to send the trap to the management station. Configure SNMP traps receiver using the `snmpconfig` command. For additional information on configuring the SNMP agent using the `snmpconfig` command, refer to the *Brocade Fabric OS Command Reference Manual*.

Complete details on the MIBs and traps are beyond the scope of this document. For more information, refer to the *Brocade Fabric OS MIB Reference Manual*.

Port Fencing and Port Decommissioning

MAPS supports port fencing for E_Ports, F_Ports, and FCIP circuits. MAPS also supports port decommissioning for E_Ports and F_Ports. However, decommissioning for F_Ports can only be done with MAPS in conjunction with Brocade SANnav™ Management Portal. These actions automatically take ports offline when configured thresholds in a given rule are exceeded. Port fencing immediately takes ports offline, which might cause a loss of traffic. Port decommissioning takes a port offline without loss of traffic. Both actions are disabled by default. Port decommissioning and port fencing can only be configured for the port health monitoring systems for which decommissioning is supported.

The decommissioning action identifies the target port and communicates the intention to decommission the port to those systems. Each affected system can agree or disagree with the action, and these responses are automatically collected before a port is decommissioned. Port decommissioning cannot be configured by itself in a MAPS rule or action. It requires port fencing to be enabled in the same rule. If you attempt to create a MAPS rule or action that has port decommissioning without port fencing, the rule or action will be rejected. MAPS can be configured to have only port fencing enabled in a rule; if this is the case, the port is taken offline immediately.

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

Port Decommissioning for E_Ports and F_Ports

For E_Ports, if port decommissioning fails, MAPS fences the port. Switches themselves can decommission E_Ports through MAPS. For F_Ports, port decommissioning works if Brocade SANnav Management Portal is actively monitoring the switch. If Brocade SANnav Management Portal is not configured on a switch, MAPS fences the F_Port.

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

The `mapsconfig --decomcfg` command has two flavors: impair and withdisable. These flavors help you to enable or disable the DECOM action without enabling the FENCE action.

The following example shows the default action with the impair flavor enabled by default.

```
switch:admin> mapsconfig --show
Configured Notifications:       FENCE,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,DECOM
Mail Recipient:                 Not Configured
Mail From Address:              Not Configured
Raslog Mode:                    Default
Decom Action Config:            Impair (No Disable)
Global Quiet Time:              5 Hours
Notification Behavior:          Adaptive
Paused members :
================
PORT :
CIRCUIT :
SFP :
SWITCH :
```

On the Brocade Gen 7 platforms, the default DECOM flavor is impair. The following MAPS action combinations are supported with the impair flavor:

- DECOM
- FENCE
- DECOM, FENCE

The impair flavor is applicable only for E_Ports.

On the Brocade Gen 6 platforms, the default DECOM flavor is withdisable. The following MAPS action combinations are supported with the withdisable flavor:

- FENCE
- DECOM, FENCE

On both Brocade Gen 6 and Gen 7 platforms, if the withdisable flavor is configured, the following actions are supported:

- FENCE
- DECOM, FENCE

On both Gen 6 and Gen 7 platforms, if the impair flavor is configured, the following actions are supported:

- DECOM
- FENCE
- DECOM, FENCE

The impair flavor is applicable only for E_Ports.

Custom RASLog IDs

You can use the custom RASLog IDs (starting from MAPS-2004) so that the third-party applications do not need to parse the RASLog message content to determine the monitoring statistics such as CRC or ITW, and thus automate the response based on the RASLog IDs. Each type of RASLog message supports up to four IDs based on severity levels such as Critical, Error, Warning, and Info. Refer to the *Brocade Fabric OS Message Reference* for the complete list of new RASLog message IDs.

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

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

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

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

In Fabric OS v9.1.x, RASLog messages generated when a rule is triggered contain additional information with the Quiet Time value enforced for the rule. This information helps with the information on the subsequent trigger for the rule on the same object is suppressed and the duration. The following is an example of the RASLog message with the Quiet Time value.

```
2021/11/11-10:56:24 (IST), [MAPS-2006], 99, FID 120 | PORT 0/17, WARNING, switch_120, port17, F-Port 17,
 Condition=ALL_HOST_PORTS(CRC/min>0), Current Value:[CRC, 100 CRCs], RuleName=defALL_HOST_PORTSCRC_0,
 Dashboard Category=Port Health, Quiet Time=None.
```

To enable the custom RASLog IDs, run the `mapsconfig --raslogmode custom` command. To switch back to the legacy RASLog mode, run the `mapsconfig --raslogmode default` command.

For more information, see MAPS Support Matrix and Default Rules and Actions.

For more information about custom RASLog related CLI commands, refer to the *Brocade Fabric OS Message Reference Manual*.

Managing RoR and Custom RASLog

A rule that monitors the frequency of the execution of existing rules is called a rule-on-rule (RoR) rule. A rule-on-rule is defined to monitor a base MAPS rule. A rule-on-rule evaluates the number of times within a given timebase that the base rule is violated. Actions defined for a rule-on-rule are triggered when the frequency of a base rule violation exceeds the threshold. Rule-on-rule monitoring allows you to apply different actions based on the frequency at which base rules are violated.

A RoR rule specifies the following items:

- The base rule that the RoR rule is monitoring.
- A group of objects to be evaluated. The group name is optional. If you do not specify a group name, the group for the base rule is used.
- The condition being monitored. Each RoR rule specifies a single condition. A condition includes a timebase. See Monitoring Systems Support Matrix for additional information.
- The actions to take if the condition is evaluated to be true.

When creating RoR rules, use all existing MAPS features including:

- Rule management
- Policy management
- Dashboard
- Alerts

Creating RoR Rules

1. Enter the `mapsrule --createRoR <RoR_rule_name> [-group <group_name> -monitor <name_of_rule_to_monitor>` command followed by the RoR rule parameters and optionally the policy that you want to assign it to. RoR rule names are not case-sensitive; My_Rule and my_rule are considered to be the same.
   The following is an example that shows creating a MAPS rule named memory_rule.

   ```
   switch:admin> mapsRule --create memory_rule -monitor MEMORY_USAGE -group chassis -timebase none -op ge -
   value 60 -action raslog
   2023/03/19-01:52:35 (EST), [MAPS-1100], 10467, FID 128, INFO, G22, Rule memory_rule is created.
   ```

   The following is an example that shows creating a MAPS RoR rule named memory_ror.

   ```
   switch:admin> mapsRule --createRoR memory_ror -monitor memory_rule -group chassis -timebase hour -op ge -
   value 5 -action raslog
   2023/03/19-01:53:30 (EST), [MAPS-1100], 10468, FID 128, INFO, G22, Rule memory_ror is created.
   ```

2. Optional: Enter the `mapsrule --show <rule_name>` command to display the rule.
3. Optional: If you added the RoR rule to an active policy, you must re-enable the policy for the RoR rule to take effect by entering the `mapspolicy --enable policy <policy_name>` command. You can add a RoR rule to an existing policy only. If you specify a group (optional), the group must exist.

   Examples of Using a RoR Rule to Monitor CPU Performance

   RoR rules allow you to take different actions based on the frequency of a base rule execution. For example, you can specify that the RoR rule ignores the first CPU spike and takes a RASLog action later. For example, the `chassis(CPU/none>=75) -action none` command is the base rule. For this base rule, you can specify a RoR rule `chassis(cpu_rule/hour>=5) -action raslog` to alert you if the base rule is executed 5 times in an hour and not to alert if the base rule is triggered fewer times in an hour. The following example shows that the CPUpolicy75 policy is active with the CPU base rule and RoR rule:

   ```
   switch:admin> mapspolicy --show CPUpolicy75
   ```

```
Policy Name: CPUpolicy75

Rule Name              |Condition                              |Actions                    |
-------------------------------------------------------------------------------------------
memory_rule            |chassis(MEMORY_USAGE/none>=60)         |raslog                     |
memory_ror             |chassis(memory_rule/hour>=5)           |raslog                     |
cpu_rule               |chassis(CPU/none>=75)                  |raslog                     |
cpu_ror                |chassis(cpu_rule/hour>=5)              |raslog                     |

Active Policy is 'CPUpolicy75'.
Unmonitored Rules are prefixed with "*"
System rules are prefixed with "+"
```

The following is an example of the RASLog message that is generated when the base rule is violated:

```
2023/03/10-03:04:00, [MAPS-1003], 2107, FID 128, WARNING, sw0, Chassis, Condition=CHASSIS(CPU>=75),
 Current Value:[CPU, 81.50 %], RuleName=cpu_rule, Dashboard Category=Switch Resource.
```

The following is an example of the RASLog message generated when the RoR rule is violated:

```
2023/03/10-02:22:00, [MAPS-1003], 2093, FID 128, WARNING, sw0, Chassis, Condition=RULE(cpu_rule/
hour>=5), Current Value:[cpu_rule, 5], RuleName=cpu_ror, Dashboard Category=Switch Resource.
```

The following example shows the dashboard output when the base and RoR rules are violated:

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

Category(Violation  |RepeatCount|Rule Name     |Execution Time   |Object    |Triggered Value|
Count)                                                                      (Units)
-------------------------------------------------------------------------------------------
Switch Resource(28) |2          |cpu_rule      |03/10/23 03:10:00|Chassis   |76.00 %        |
                    |           |              |                 |Chassis   |76.40 %        |
                    |2          |memory_rule   |03/10/23 03:10:00|Chassis   |60.17 %        |
                    |           |              |                 |Chassis   |61.04 %        |
                    |1          |cpu_ror       |03/10/23 02:22:00|Chassis   |5 Trigger cnt. |
                    |1          |memory_ror    |03/10/23 02:22:00|Chassis   |5 Trigger cnt. |
                    |10         |cpu_rule      |03/10/23 02:58:00|Chassis   |76.00 %        |
                    |           |              |                 |Chassis   |77.20 %        |
                    |           |              |                 |Chassis   |78.10 %        |
                    |           |              |                 |Chassis   |79.60 %        |
                    |           |              |                 |Chassis   |81.50 %        |
```

If the default policy does not suit your requirements, you can clone an existing default policy and add your custom rule. When you create a custom RoR rule, you must account for the constraints that are associated with each rule parameter. For detailed information, see Monitoring Systems Support Matrix.

Restrictions That Affect the Modification of Base Rules

If you have created one or more RoR rules to monitor a base rule, the following restrictions apply when updating or configuring the base rule:

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

RoR Rules with Custom RASLogs

You can use the basic elements of MAPS that are described in MAPS Basic Elements while defining the RoR rules. The following important aspects of the MAPS basic elements that are relevant to RoR rules are covered:

- Custom RASLogs
- Timebase
- Thresholds
- Actions

Custom RASLogs on the RoR rules inherit from the base rule. For example, the CRC RASLog messages use the following IDs for the RoR rules and base rules:

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

For the following MAPS elements, the values are derived from the thresholds configured in the Monitoring Systems Support Matrix:

- Timebase
- Operator (values are always > or =)
- Threshold (unsigned integers 0 and above)
- Actions

Custom MAPS Groups

Custom MAPS groups allow you to specify groups that are defined by the characteristics that you select.

In the previous sections, you created a custom policy, created a rule, and added the rule to the custom policy. MAPS provides the default groups that can meet most of your requirements. However, you can create custom groups in some scenarios where your requirements are different from the default members. Use the following steps to create custom groups:

1. Enter the following command to customize a specific MAPS group.

   ```
   logicalgroup --addmem <group name> -members <member port numbers>
   ```

   For example, to customize the ALL_HOST_PORTS group to ports 10 through 11, enter the `logicalgroup --addmem ALL_HOST_PORTS -members 10-11` command.
2. Use the `logicalgroup --create` command to create a group. The following example command creates a port group.

   ```
   switch:admin> logicalGroup --create port_group -type PORT -members 10-11
   ```

   The following example shows the newly created port group.

   ```
   switch:admin> logicalgroup --show port_group

   ------------------------------------------------------------------------------------------------------
   Group Name                      |Predefined |Type              |Member Count |Members
   ------------------------------------------------------------------------------------------------------
   port_group                      |No         |Port              |6            |10-11
   ```

User-Defined Groups

You might need groups of elements that are more appropriate 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 might want to create a separate group of SFP transceivers for each vendor. In another scenario, some ports might 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 filter 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.

Be aware of the following rules when you work with user-defined groups, also known as *custom groups*:

- Dynamic groups are only used to group ports.
- The device node WWN information is fetched from the NS database; group membership is validated against this database.
- If you create a group with the feature specified as device node WWN on an Access Gateway device, the ports to which the devices are connected will be part of the group.
- On a switch that is connected to Access Gateway (AG), the ports that are connected to AG are not grouped in a user-defined group.
- A port or device can be a member of multiple groups.
- A maximum of 64 user-defined groups and imported flows that are 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 the same.

Ports in the ETH mode are not included as a part of MAPS dynamic user-defined groups and cannot be included as a part of the static user-defined groups. There are three types of ports in ETH mode: VN_Ports, VF_Ports, and port-channels. The ports in ETH mode are not supported by the `logicalgroup --addmember` , `logicalgroup --create` , and `logicalgroup --update` commands.

Creating a Static User-Defined Group

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

As an example of a static definition, you can define a group called MY_CRITICAL_PORTS and specify its members as 2/1-10, 2/15, and 3/1-20. In this case, the group has a fixed membership. To add or remove a member from the group, use the `logicalGroup` command, and specify whether you want to add or remove a member. You can add port, sfp, or circuits.

To create a static group that contains a specific set of ports, complete the following steps:

1. Enter the `logicalGroup --create <group_name> -type port -members <member_list>` command.

   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.

2. Optional: Enter the `logicalGroup --show <group_name> -details` command to view the group membership.

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

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

   For more information about the `logicalGroup` command, refer to the *Brocade Fabric OS Command Reference Manual*.

Modifying a Static User-Defined Group

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

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

1. Use the following commands to add or delete the specific ports from a group:

   - To explicitly add ports to the group, enter the `logicalGroup --addmember <group_name> -members <member_list>` command.
   - To explicitly remove ports from the group, enter the `logicalGroup --delmember <group_name> -members <member_list>` command.

   You need to specify every element in the command. When adding or removing ports, 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 (2/15-64,16/1-15).

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

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

2. Optional: Enter the `logicalGroup --show <group_name> -details` command to view the group membership.

   For more information about the `logicalGroup` command, refer to the *Brocade Fabric OS Command Reference Manual*.

Creating a Dynamic User-Defined Group

By using a dynamic definition, you can create a group that can be monitored where the membership is determined by meeting a filter value. When the value is matched, the port or device is added to the group, and it is included in

the monitoring. When the value is not matched, the port or device is removed from the group, and it is not included in monitoring data for that group.

Dynamic groups allow you to group ports based on a filter value for an attribute of the port. When the attribute value matches the specified filter, it is included in the group. When the value of the attribute for a port does not match the filter, the port is automatically removed from the group.

Dynamic groups can be created by specifying a pattern filter that is applied to the port name or WWN of the device node. The filter can be added as input as a regular expression. The following table shows the special characters that can be used in regular expressions.

Table 37: Group-Definition Operators

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

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

1. Enter the `logicalgroup --create <group_name> -type <type> -feature <feature_type> -pattern <pattern>` command.

   For the `<feature_type>` option, 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.

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

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

2. Enter the `logicalgroup --show <group_name> -details` command to view the group membership. This is optional.

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

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

For more information about the `logicalgroup` command, refer to the *Brocade Fabric OS Command Reference Manual*.

Modifying a Dynamic User-Defined Group

MAPS allows you to change the definition pattern used to specify a dynamic user-defined group after you created it. Specify the ports that need to be added or deleted from the dynamic group in any of the following scenarios:

- You want to remove ports that were automatically grouped by the pattern filter.
- You want to add ports that were not included by pattern filter.

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

1. Enter the `logicalGroup --update <group_name> -feature <feature_name> -pattern <pattern>` command.

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

   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*"
   ```

2. Use the following commands to add or delete specific ports from the group. You can also use this command to modify the group membership of predefined groups.

   - To explicitly add ports to the group, enter the `logicalGroup --addmember <group_name> -members <member_list>` command.
   - To explicitly remove ports from the group, enter the `logicalGroup --delmember <group_name> -members <member_list>` command.

   You must specify every element in the command. When adding or removing ports, specify either a single port (2/15), or specify multiple ports as either individual IDs separated by commas (2/15, 2/16, 3/15), or a range of ports with the IDs separated by hyphens (2/15-16, 3/15).

3. Optional: Enter the `logicalGroup --show <group_name> -details` command to view the group membership.

   For more information about the `logicalGroup` command, refer to the *Brocade Fabric OS Command Reference Manual*.

Restoring a Group to Its Default Membership

MAPS allows you to restore the membership of any modified MAPS group to its default. You can perform this procedure with predefined groups and dynamic user-defined groups. This command does not impact groups with a static definition.

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

1. Enter the `logicalgroup --restore <group_name>` command. This command restores the group membership to its default.
2. Optional: Enter the `logicalgroup --show <group_name> -details` command to view the group membership.

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

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

   GroupName          : ALL_F_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 ALL_F_PORTS

switch:admin> logicalgroup --show ALL_F_PORTS -detail

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

Switch-Level Configuration

The information in this section helps you control MAPS behavior.

You can configure the switch to perform the following advanced MAPS functions using the `mapsconfig` option:

- Defining the list of allowable actions that can be taken on the switch when a threshold is triggered. For more information, see Enabling or Disabling Actions at a Global Level.
- Configuring email address to which the alerts must be delivered. For more information, see Email Alerts.
- Sending a test email message to check the correct email server configuration. For more information, see Email Delivery Monitoring.
- Configuring the RASLog mode to custom or default. In custom mode, MAPS generates different RASLogs for different monitoring systems; whereas in the default mode MAPS generates generic RASLogs from MAPS-1001 through MAPS-1004 for all the monitoring systems. For more information, see Custom RASLog IDs.
- Configuring the result of the DECOM action in MAPS to either impair the link or decommission with withdisable. For more information, see Port Fencing and Port Decommissioning.
- Configuring FPIN actions whenever congestion-related FPI rules are triggered and triggering the Fabric Notifications periodically to participating devices until the congestion criteria is cleared by MAPS. For more information about FPIN action, see Actions.
- Configuring the Quiet Time, that is the time interval for which the alerts are suppressed. For more information, see Quiet Time.

In addition to the advanced MAPS functions above, you can also configure the following optional MAPS functions using the `mapsconfig` option:

- Pausing or restarting MAPS monitoring specific elements. For more information, see Pausing and Restarting MAPS.
- Deleting all user-defined MAPS configurations related to rules, groups, policies, and so on. For more information, refer to the "mapsconfig" section in the *Brocade Fabric OS Command Reference Manual*.

For detailed instructions about using the `mapsconfig` command, refer to the *Brocade Fabric OS Command Reference Manual*.

MAPS Dashboard

The MAPS dashboard provides a summary view of the switch health status along with details on out-of-range conditions.

The dashboard information allows you to determine whether the switch or fabric is working according to policy or whether you must investigate further. The dashboard displays the following information:

- Basic information about the dashboard
- A status report on the current health of the switch
- A summary of the categories of the switch and its status for the last seven days
- A list of the rules currently affecting the health of the switch

To view the dashboard summary, enter the `mapsdb –show` command. Enter the `mapsdb --show all` command to display the basic dashboard information. If you want to view only history, enter the `mapsdb --show history` command.

Dashboard Basic Information Section

The Dashboard Information section displays the following basic dashboard data:

- The time when the dashboard was started
- The name of the active policy
- Any fenced, decommissioned, or quarantined ports
- The list of FCIP circuits that are fenced
- The top zoned PIDs

The following example shows the basic dashboard information.

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

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

DB start time:                  Wed Aug 25 11:34:46 2021
Active policy:                  dflt_aggressive_policy
Configured Notifications:       RASLOG,SNMP,EMAIL,FENCE,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,DECOM,SDDQ,FPIN
Fenced Ports :                  None
Decommissioned Ports :          None
Fenced circuits :               None
Quarantined Ports :             None
Top Zoned PIDs <pid(it-flows)>: 0x71a080(4) 0x7103c0(1) 0x7100c0(1) 0x7101c0(1) 0x7102c0(1)
                    ----------------Output truncated---------------------
```

Notes on Dashboard Data

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

- The following dashboard state conditions can be displayed:

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

Switch Health Report Section

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

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

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

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


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

See Switch Status Policy for more details on switch policies.

Summary Report Section

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

MAPS monitors the following categories:

- Port Health
- BE Port Health
- Extension GE Port Health
- FRU Health
- Security Violations
- Fabric State Changes
- Switch Resource
- Extension Health
- Fabric Performance Impact
- IO Latency
- IO Health
- Switch Status Policy
- Uncategorized Monitoring Systems

The following output extract shows a sample Summary Report section.

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

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

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

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

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

- For each rule, the five most recent distinct rule violations were occurred.
- For each rule, the five most recent entities on which the rules were triggered.
- Although a rule might be triggered multiple times within a given hour, only the timestamp of the latest violation is stored.
- Each rule violation individually is reflected in the rule count for that category and the repeat count for that rule.

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

the rule's violation). Alternatively, if a rule was triggered 12 times since midnight, but each violation happened in a different hour, each violation is logged separately in the dashboard.

The following example shows a sample Rules Affecting Health section. The column headings in the example are edited to display clearly.

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

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

History Data Section

The History Data section provides information on how the switch behaves regardless of whether rules were triggered. It contains only port-related statistics and is the basic 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. For more information about historical data, see Viewing Historical Data. The following `mapsdb --show history` output extract shows a sample History Data section.

```
(output truncated)
```

```
4 History Data:
===============
Stats(Units)         Current   17/02/23  15/01/23    --/--/--    --/--/--   --/--/--   --/--/--
                     Port(val) Port(val) Port(val)
-------------------------------------------------------------------------------------------
CRC(CRCs)            1/13(20)   -          -          -          -          -          -
ITW(ITWs)            -          1/13(612)  -          -          -          -          -
LOSS_SYNC(SyncLoss)  -          -          -          -          -          -          -
LF                   -          -          -          -          -          -          -
LOSS_SIGNAL(LOS)     -          -          -          -          -          -          -
PE(Errors)           -          -          -          -          -          -          -
STATE_CHG            -          -          -          -          -          -          -
C3TXTO(Timeouts)     -          -          -          -          -          -          -
RX(%)                -          -          -          -          -          -          -
TX(%)                -          -          -          -          -          -          -
UTIL(%)              -          -          -          -          -          -          -
BN_SECS(Seconds)     -          -          -          -          -          -          -

5 History Data for Backend ports:
=================================
Stats(Units)         Current   17/02/23   15/01/23   --/--/--   --/--/--   --/--/--   --/--/--
                     Port(val) Port(val)  Port(val)
-------------------------------------------------------------------------------------------
CRC(CRCs)            6/8(50)     -          -          -          -          -          -
                    ----------------------Output truncated-----------------------
```

MAPS Dashboard Display Options

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

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

Viewing the MAPS Dashboard

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

Perform the following steps to view the status of the switch:

1. Connect to the switch and log on using an account with admin permissions.
2. Enter the `mapsdb --show` command 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 result of the `mapsdb --show all` command.

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

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

DB start time:                  Tue Jan 01 16:38:55 2023
Active policy:                  dflt_conservative_policy
Configured Notifications:       RASLOG,EMAIL,FENCE,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,DECOM
```

```
Fenced Ports :                  None
Decommissioned Ports :          None
Fenced circuits :               None
Quarantined Ports :             None
Top Zoned PIDs <pid(it-flows)>: 0xbb3740(1) 0xbb1000(1) 0xbb3741(1) 0xbb8f00(1)0xbb0000(1)

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

Current Switch Policy Status: HEALTHY


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

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

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

Category(Violation Count)|RepeatCount|Rule Name                 |Execution Time   |Object
  |Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------------------------------
Fabric Performance Impact|3          |defALL_E_PORTSUTIL_95     |01/01/23 14:06:31|E-Port 7/1
  |99.62 %               |
(392)             |      |           |                          |                 |
                         |           |                          |                 |E-Port 7/1
  |99.89 %               |
                         |           |                          |                 |E-Port 7/1
  |95.76 %               |
                         |7          |defALL_E_PORTSTX_95       |01/01/23 14:06:37|E-Port 7/1
  |100.00 %              |
                         |           |                          |                 |E-Port 7/1
  |100.00 %              |
                         |           |                          |                 |E-Port 7/1
  |100.00 %              |
                         |           |                          |                 |E-Port 7/1
  |100.00 %              |
                         |           |                          |                 |E-Port 7/1
  |100.00 %              |
```

```
                        |1          |defALL_E_PORTSTX_90        |01/01/23 13:56:01|E-Port 7/1
  |91.76 %              |

                                    --------------Output Truncated---------------


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

Stats(Units)        Current     01/02/23    01/01/23    --/--/--    --/--/--    --/--/--
 --/--/--
-----------------------------------------------------------------------------------------------------

CRC(CRCs)           -           -           -           -           -           -           -
ITW(ITWs)           -           -           -           -           -           -           -
LOSS_SYNC(SyncLoss) -           -           -           -           -           -           -
LF(LFs)             -           -           -           -           -           -           -
LOSS_SIGNAL(LOS)    3/29(148)   3/29(240)   3/29(56)    -           -           -           -
PE(Errors)          -           -           -           -           -           -           -
STATE_CHG           -           -           -           -           -           -           -
LR(LRs)             -           -           -           -           -           -           -
C3TXTO(Timeouts)    -           -           -           -           -           -           -
RX(%)               6/23(34.67) 3/33(34.29) 3/33(10.96) -           -           -           -
                    3/33(34.12) 6/23(33.71) 6/23(10.77) -           -           -           -
                    8/33(25.40) 8/33(23.40) 8/33(6.88)  -           -           -           -
                    7/3(19.61)  7/3(19.61)  7/3(6.29)    -           -           -           -
                    8/34(17.04) 7/0(16.92)  7/0(5.42)    -           -           -           -
TX(%)               6/23(41.57) 6/23(40.41) 6/23(12.90) -           -           -           -
                    3/33(21.62) 3/33(22.38) 3/33(6.85)  -           -           -           -
                    3/16(12.25) 3/16(11.73) 3/16(3.74)  -           -           -           -
                    7/1(11.86)  3/38(6.37)  7/2(2.38)    -           -           -           -
                    7/2(6.42)   3/0(5.77)   3/38(1.94)  -           -           -           -
UTIL(%)             6/23(38.11) 7/6(3924.71) 6/23(11.83) -          -           -           -
                    3/33(27.07) 8/9(3924.71) 3/33(8.90)  -          -           -           -
                    3/16(13.47) 7/5(3924.63) 3/16(4.12)  -          -           -           -
                    7/1(12.96)  6/23(37.06) 8/33(3.44)  -           -           -           -
                    8/33(12.70) 3/33(16.21) 7/3(3.14)   -           -           -           -
BN_SECS(Seconds)    -           -           -           -           -           -           -

5 History Data for Backend ports:
=================================
Stats(Units)        Current     01/02/23    01/01/23    --/--/--    --/--/--    --/--/--
 --/--/--
-----------------------------------------------------------------------------------------------------

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

6 History Data for Gige Ethernet ports:
```

```
=========================================
Stats(Units)        Current     01/02/23    01/01/23    --/--/--    --/--/--    --/--/--
 --/--/--
-------------------------------------------------------------------------------------------------

GE_CRC(CRCs)        -           -           -           -           -           -           -
GE_LOS_OF_SIG(LOS)  -           -           -           -           -           -           -
```

See MAPS Monitoring Categories for explanations of the categories listed in the dashboard output.

Viewing a Summary Switch Status Report

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

To view a summary switch status report, perform the following steps:

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

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

```
switch:admin> mapsdb --show

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

DB start time:                  Thu Mar  2 08:32:24 2023
Active policy:                  max_rules
Configured Notifications:       RASLOG,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL
Fenced Ports :                  None
Decommissioned Ports :          None
Fenced circuits :               N/A
Quarantined Ports :             None
Top Zoned PIDs <pid(it-flows)>:

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

Current Switch Policy Status: HEALTHY


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

Category                |Today                      |Last 7 days                |
-------------------------------------------------------------------------------
Port Health             |No Errors                  |No Errors                  |
BE Port Health          |No Errors                  |No Errors                  |
Fru Health              |In operating range         |In operating range         |
Security Violations     |In operating range         |No Errors                  |
Fabric State Changes    |In operating range         |No Errors                  |
```

```
Switch Resource          |In operating range        |In operating range        |
Fabric Performance Impact|In operating range        |In operating range        |
IO Health                |In operating range        |In operating range        |
IO Latency               |In operating range        |In operating range        |


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

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

Viewing Historical Data

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

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

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

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

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

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

1 History Data:
===============

Stats(Units)        Current     --/--/--    --/--/--    --/--/--    --/--/--    --/--/--
 --/--/--
-----------------------------------------------------------------------------------------------

CRC(CRCs)           -           -           -           -           -           -           -
ITW(ITWs)           -           -           -           -           -           -           -
LOSS_SYNC(SyncLoss) -           -           -           -           -           -           -
LF(LFs)             -           -           -           -           -           -           -
LOSS_SIGNAL(LOS)    -           -           -           -           -           -           -
PE(Errors)          -           -           -           -           -           -           -
STATE_CHG           -           -           -           -           -           -           -
LR(LRs)             -           -           -           -           -           -           -
C3TXTO(Timeouts)    -           -           -           -           -           -           -
RX(%)               -           -           -           -           -           -           -
TX(%)               -           -           -           -           -           -           -
UTIL(%)             -           -           -           -           -           -           -
BN_SECS(Seconds)    -           -           -           -           -           -           -
```

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 ordinary historical data.

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

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

The following example displays historical port performance data for February 20, 2023, on a chassis-based platform. Because the health status of the current switch policy is CRITICAL, the sections Contributing Factors and Rules Affecting Health are displayed. If the current switch policy status was HEALTHY, neither of these sections are displayed. The column headings in the example are edited to display clearly.

```
switch:admin> mapsdb --show details -day 2/20/2023

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

DB start time:                  Thu Feb 20 15:28:01 2023
Active policy:                  dflt_aggressive_policy
------------------Output truncated--------------------------

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

Current Switch Policy Status: MARGINAL
Contributing Factors:
---------------------
*SYSTEM_TEMP (MARGINAL).


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

Category                 |Today                     |2/20/2023                |
-------------------------------------------------------------------------------
Port Health              |No Errors                 |In operating range       |
BE Port Health           |No Errors                 |No Errors                |
Extension GE Port Health |No Errors                 |No Errors                |
Fru Health               |In operating range        |In operating range       |
Security Violations      |No Errors                 |No Errors                |
Fabric State Changes     |No Errors                 |In operating range       |
Switch Resource          |In operating range        |In operating range       |
Extension Health         |In operating range        |In operating range       |
Fabric Performance Impact|In operating range        |Out of operating range   |
IO Latency               |In operating range        |In operating range       |
IO SCSI Health           |In operating range        |In operating range       |
```

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

Category(Violation Count)|RepeatCount|Rule Name                  |Execution Time   |Object           |
Triggered Value(Units)      |
-------------------------------------------------------------------------------------------------------------------------------
Fabric Performance Impact|56         |defALL_PORTS_IO_LATENCY_CLE|02/20/23 23:59:02|E-Port 8/32      |IO_LATE
(315)                    |           |AR                         |                 |
NCY_CLEAR                |

                        ----------------------Output truncated--------------------------------

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

Stats(Units)        Current     02/20/23    --/--/--    --/--/--    --/--/--    --/--/--    --/--/--
-------------------------------------------------------------------------------------------------

CRC(CRCs)           -           -           -           -           -           -           -
ITW(ITWs)           -           -           -           -           -           -           -
LOSS_SYNC(SyncLoss) -           -           -           -           -           -           -
LF(LFs)             -           3/32(2)     -           -           -           -           -
                    -           3/33(1)     -           -           -           -           -

       --------------------------Output truncated-------------------------------------

5 History Data for Backend ports:
=================================
Stats(Units)        Current     02/20/21    --/--/--    --/--/--    --/--/--    --/--/--    --/--/--
-------------------------------------------------------------------------------------------------

CRC(CRCs)           -           -           -           -           -           -           -
ITW(ITWs)           -           -           -           -           -           -           -
LR(LRs)             -           -           -           -           -           -           -

6 History Data for Gige Ethernet ports:
=======================================
Stats(Units)        Current     02/20/23    --/--/--    --/--/--    --/--/--    --/--/--    --/--/--
-------------------------------------------------------------------------------------------------

GE_CRC(CRCs)        -           -           -           -           -           -           -
GE_LOS_OF_SIG(LOS)  -           -           -           -           -           -           -
```

The following example displays historical port performance data for four hours on a chassis-based platform. The History Data section is truncated to display the output correctly.

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

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

DB start time:                  Thu Feb 20 15:28:01 2023
Active policy:                  dflt_base_policy
```

```
Configured Notifications:       SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,DECOM,FPIN
Fenced Ports :                  None
Decommissioned Ports :          None
Fenced circuits :               None
Quarantined Ports :             None
Top Zoned PIDs <pid(it-flows)>: 0xbb3740(1) 0xbb3741(1) 0xbb1000(1) 0xbb6100(1) 0xbb8f00(1)

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

Current Switch Policy Status: MARGINAL
Contributing Factors:
---------------------
*SYSTEM_TEMP (MARGINAL).


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

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


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

Category(Violation Count)|RepeatCount|Rule Name                  |Execution Time   |Object       |Triggered
 Value(Units)          |
-------------------------------------------------------------------------------------------------------------------------------
Fabric Performance Impact|1          |defALL_E_PORTSUTIL_95      |02/21/23 04:40:37|E-Port 8/34  |95.86 %
                  |
(30)                     |
                         |2          |defALL_PORTS_IO_PERF_IMPACT|02/21/23 04:35:02|E-Port 7/2   |
IO_PERF_IMPACT        |
                         |           |                           |                 |E-Port 8/32  |
IO_PERF_IMPACT        |
                         |2          |defALL_PORTS_IO_LATENCY_CLE|02/21/23 04:34:02|E-Port 7/2   |
IO_LATENCY_CLEAR      |
             ------------------------Output truncated-----------------------------

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

```
Stats(Units)        Current     02/20/23    --/--/--    --/--/--    --/--/--    --/--/--    --/--/--
-------------------------------------------------------------------------------------------------

CRC(CRCs)           -           -           -           -           -           -           -
ITW(ITWs)           -           -           -           -           -           -           -
LOSS_SYNC(SyncLoss) -           -           -           -           -           -           -
LF(LFs)             -           3/32(2)     -           -           -           -           -
                    -           3/33(1)     -           -           -           -           -
LOSS_SIGNAL(LOS)    -           3/32(1)     -           -           -           -           -
                    -           3/33(1)     -           -           -           -           -
PE(Errors)          -           -           -           -           -           -           -
STATE_CHG           -           3/32(2)     -           -           -           -           -
                    -           3/33(2)     -           -           -           -           -
LR(LRs)             -           3/32(3)     -           -           -           -           -
                    -           3/33(2)     -           -           -           -           -
C3TXTO(Timeouts)    -           -           -           -           -           -           -
RX(%)               7/2(69.85)  7/2(22.19)  -           -           -           -           -
                    8/35(53.78) 8/35(17.66) -           -           -           -           -
                    8/33(50.18) 8/33(16.15) -           -           -           -           -
                    ---------
TX(%)               8/34(81.24) 8/34(26.45) -           -           -           -           -
                    7/3(51.32)  7/3(14.82)  -           -           -           -           -
                    ---------
UTIL(%)             8/34(60.12) 8/34(19.31) -           -           -           -           -
                    7/3(41.18)  7/2(12.73)  -           -           -           -           -
                    -----------
BN_SECS(Seconds)    -           -           -           -           -           -           -

5 History Data for Backend ports:
=================================
Stats(Units)        Current     02/20/23    --/--/--    --/--/--    --/--/--    --/--/--    --/--/--
-------------------------------------------------------------------------------------------------

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

6 History Data for Gige Ethernet ports:
=======================================
Stats(Units)        Current     02/20/23    --/--/--    --/--/--    --/--/--    --/--/--    --/--/--
-------------------------------------------------------------------------------------------------

GE_CRC(CRCs)        -           -           -           -           -           -           -
GE_LOS_OF_SIG(LOS)  -           -           -           -           -           -           -
```

Enter the `mapsdb --show -category <db categories>` command to view the values for the db categories. The following example shows the data for the selected db category.

```
switch:admin> mapsdb --show -category FRU
```

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

DB start time:                  Tue Feb 14 16:38:55 2023
Active policy:                  dflt_conservative_policy
Configured Notifications:       RASLOG,EMAIL,FENCE,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,DECOM
Fenced Ports :                  None
Decommissioned Ports :          None
Fenced circuits :               None
Quarantined Ports :             None
Top Zoned PIDs <pid(it-flows)>: 0xbb3740(1) 0xbb1000(1) 0xbb3741(1) 0xbb8f00(1) 0xbb0000(1)

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

Current Switch Policy Status: HEALTHY


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

Category                |Today                      |Last 7 days              |
-------------------------------------------------------------------------------
Fru Health              |In operating range         |In operating range       |


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

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

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

Clearing MAPS Dashboard Data

To delete the stored data from the MAPS dashboard, enter the `mapsdb --clear` command. This command displays only the data that is logged after you changed a switch (or a rule). The dashboard is also cleared on a reboot or an HA failover.

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

1. Connect to the switch and log on using an account with admin permissions.
2. Enter the `mapsdb --clear` command and specify the level of data (`all, summary, history, or congestion`) you want to remove from the display. You can use the `-force` option to forcefully proceed with CLI execution for the `all, summary, history, or congestion` options without prompting for user input.

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

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

   The `mapsdb --clear summary` command clears the dashboard summary data.

   ```
   switch:admin> mapsdb --clear summary
   WARNING: This command will clear summary data
   ```

```
Do you want to continue? (yes, y, no, n): [no]
```

The following RASLog message shows the example of the dashboard summary data after clearing.

```
2023/02/24-16:56:05, [MAPS-1203], 47554, FID 128, WARNING, G610, Dashboard summary data has been
 cleared.
```

The following example shows the dashboard output after clearing the dashboard summary data.

```
switch:admin> mapsdb --show

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

DB start time:                  Thu Feb 23 19:25:46 2023
Active policy:                  dflt_conservative_policy
Configured Notifications:       RASLOG,EMAIL,FENCE,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,DECOM
Fenced Ports :                  None
Decommissioned Ports :          None
Fenced circuits :               None
Quarantined Ports :             None
Top Zoned PIDs <pid(it-flows)>: 0xbb3740(1) 0xbb1000(1) 0xbb3741(1) 0xbb8f00(1) 0xbb0000(1)

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

Current Switch Policy Status: HEALTHY


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

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

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

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

Port Monitoring Using MAPS

This section describes monitoring a group of ports, monitoring ports based on their assigned names, and ports that are connected to a device that has a device World Wide Name (WWN).

Monitoring a Group of Ports

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

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

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

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

   You must enable the policy even if it is an active policy. Adding a rule to the active policy does not take effect until you reenable 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:admin> logicalgroup --create unix_ports -type port -members "1,3,17,21"
switch:admin> mapsrule --create unixHiCrc -monitor crc -group unix_ports -timebase min -op g -value
 50 -action raslog -policy my_policy
switch:admin> mapspolicy --enable my_policy
```

Port Monitoring Using Port Names

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

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

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

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

Port Monitoring Using Device WWNs

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

You can use this feature 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 that is named WXYZ_Devs with a membership determined by having a WWN that begins with 30:08:00:05.

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

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

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

Adding a Port to an Existing Static Group

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

Add a port to a static group, to both user-defined and predefined.

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

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

1. Connect to the switch and log on using an account with admin permissions.
2. Enter the `--addmember <group_name> -member <member_list>` command.

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

   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 the `logicalgroup --show group_name` command to see the members of the named group.

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

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

   Listing this group displays the following output.

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

Adding Missing Ports to a Dynamic Group

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

Specify any of the following port information for dynamic groups:

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

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

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

The same restrictions apply as described in Adding a Port to an Existing Group.

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

   The following example shows how you might use the `logicalgroup` command to add port 5 to the ALL_HOST_PORTS group. First you see that port 5 is not part of the group, then the port is added to the group, and last you can see that the port was 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

You can remove ports from predefined groups or user-defined dynamic groups. This procedure is useful if you have devices that erroneously identify themselves as both host and target or you have one-off exceptions where you want to remove a port from a user-defined dynamic group.

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

1. Connect to the switch and log on using an account with admin permissions.
2. Enter the `logicalGroup --delmember <group_name> -members <member_list>` command.

   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 the `logicalGroup --show <group_name>` command to confirm that the named ports are no longer part of the group.

   The following example shows how you might use the `logicalgroup` command to delete port 5 to the ALL_TARGET_PORTS group. First you see that port 5 is part of the group, then the port is deleted from the group, and last you can see that the port was deleted to the group.

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

switch:admin> logicalgroup --delmember ALL_TARGET_PORTS -mem 5
```

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

Back-End Port Monitoring

A back-end port connects a core switching blade to a port or application blade (and the other way around). The primary task of back-end ports is to route packets passing through ASICs of a switch. Switch (and consequently fabric) performance degrades when there are errors in back-end ports. MAPS error notification allows you to take a corrective action earlier. In terms of monitoring system functionality, back-end ports can be connected to ports within a fixed-port switch or other blades within a chassis. Their functionality is different from front-end ports that are connected to devices outside the switch.

MAPS monitors port counter errors on back-end ports in chassis-based switches and the Brocade G630 and G730 Switches. When back-end port rules are triggered, you can then do SerDes (Serializer/Deserializer) tuning on those ports where errors exceed the desired threshold, which improves the packet-forwarding performance of these ports.

When a switch is initialized, all back-end ports are automatically brought online and stay online in any of the following conditions:

- The slot is powered off
- The blade is removed (for chassis-based switches)
- If the switch is down (for fixed-port switches)

This allows MAPS to monitor the platform for back-end port errors continuously. The monitoring systems are for a given switch or chassis, so all the monitoring systems are monitored only in the default switch.

MAPS monitors the port counter statistics for back-end ports through the group ALL_BE_PORTS, which identifies each port using a combination such as 3/31. In the case of fixed-port switches, the slot number is 0. The History data for the back-end ports are collected for seven days and is displayed in the History Data for the Backend ports section of the MAPS dashboard.

MAPS monitors the back-end port errors and keeps track of the connected port for every back-end port. When a RASLog is generated, the connected port information is added, as shown in the RASLog output. In RASLog, 1/14 is the port where the errors are seen, and port 5/182 is the connected port when errors are seen on any back-end port. This additional port information helps in debugging and fixing the issue.

The back-end port monitoring monitors the following back-end port errors:

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

The following example is a sample of a RASLog message that is generated for a back-end port when the threshold set for the number of CRC errors is reached.

```
2023/02/04-21:40:02, [MAPS-1003], 48, SLOT 2 | FID 1, WARNING, dcx_178, BE Port 1/14,
 Condition=ALL_BE_PORTS(CRC/5MIN>10), Current Value:[CRC,125 CRCs (Conn. port 5/182)],
 RuleName=defALL_BE_PORTSCRC_5M_10, Dashboard Category=BE Port Health.
```

For more information about back-end health monitoring, see BE Port Health.

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 arelisted on the line starting with BE Port Health (1) and in the section labeled *4.2 Backend port History Data*. The column headings in the example are edited to allow the example to display clearly.

```
      (output truncated)
Rules Affecting Health:
================================
Category(Rule Cnt)|Rpt Cnt|Rule Name                | Execution Time    |Object     |Triggered Value
--------------------------------------------------------------------------------------------------------
BE Port Health(1) |1      | defALL_BE_PORTSCRC_5M_10 | 01/21/23 01:30:60 |Port 6/8   | 50 CRCs


4.2 Backend port History Data:
================================
Stats(Units)        Current   01/21/23   01/14/23   --/--/--   --/--/--   --/--/--   --/--/--
                    Port(val) Port(val)  Port(val)
----------------------------------------------------------------------------------------------
CRC(CRCs)           6/8(50)     -          -          -          -          -          -
```

Front-End Encryption Port Monitoring

Encryption is supported on front-end ports. MAPS monitors error counters for these ports.

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

Note the following points when working with front-end encryption ports:

- Encryption for front-end ports is supported on fibre-channel ports in blades only on the Brocade X6-4 and Brocade X6-8 Directors, and Gen 7 systems.
- Encryption is supported only on E_Ports.

MAPS monitors the following error counters:

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

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

Example of Monitoring Violations on Encryption Ports

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

```
2023/01/11-14:48:12 (IST), [MAPS-1003], 689, FID 120 | PORT 0/43, WARNING, TYR_120, port43, E-Port 43,
 Condition=ALL_E_PORTS(CRC/min>21), Current Value:[CRC, 50 CRCs], RuleName=defALL_E_PORTSCRC_21, Dashboard
 Category=Port Health, Quiet Time=10 hour.
```

```
2023/01/11-14:48:12 (IST), [MAPS-1010], 690, FID 120 | PORT 0/43, ERROR, TYR_120, Port(s) fenced due to
 RuleName=defALL_E_PORTSCRC_40, Condition=ALL_E_PORTS(CRC/min>40), Obj:port43, E-Port 43 [CRC,50 CRCs]
```

The following example shows the dashboard output of the `mapsdb --show` command that is triggered for ENCR_BLK.

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

Category(Violation Count)|RepeatCount|Rule Name                  |Execution Time   |Object           |
Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------------------------------
Port Health(2)           |1          |defALL_E_PORTSCRC_40       |01/11/23 14:48:12|E-Port 43        |50 CRCs
               |
                         |1          |defALL_E_PORTSCRC_21       |01/11/23 14:48:12|E-Port 43        |50 CRCs
               |

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

Category(Violation Count)|RepeatCount|Rule Name                  |Execution Time   |Object     |Triggered
 Value(Units  |
-------------------------------------------------------------------------------------------------------------------------------
Port Health(4)           |1          |defALL_E_PORTSENCR_BLK     |01/26/23 12:21:00|E-Port 5/25|160 Timeouts
          |
                         |3          |defALL_TARGET_PORTSSTATE_CH|01/26/23 12:23:01|F-Port 5/17|1
          |
                         |           |G_0                        |                 |U-Port 5/17|1
          |
                         |           |                           |                 |F-Port 5/17|2
          |
Fabric Performance Impact|3          |defALL_TARGET_PORTSRX_95   |01/26/23 12:20:42|F-Port 5/17|95.96 %
          |
(9)                      |           |                           |                 |           |
          |
                         |           |                           |                 |F-Port 5/17|95.95 %
          |
                         |           |                           |                 |F-Port 5/17|95.54 %
          |
                         |3          |defALL_TARGET_PORTSUTIL_95 |01/26/23 12:20:36|F-Port 5/17|95.94 %
          |
                         |           |                           |                 |F-Port 5/17|95.11 %
          |
                         |           |                           |                 |F-Port 5/17|95.13 %
          |
                         |3          |defALL_TARGET_PORTSTX_95   |01/26/23 12:20:36|F-Port 5/17|95.91 %
          |
```

Gigabit Ethernet Port Monitoring

The Fabric OS software allows you to monitor GE ports in a switch and receive counter errors that are reported by ASIC drivers as RASLog, SNMP, and Email alerts. This reporting helps you identify the nature of FCIP and IP Extension traffic errors at the Level 2 (L2) link layer of the Fibre Channel protocol.

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

MAPS monitors the following error counters in GE ports using the ALL_EXT_GE_PORTS group:

- CRC – Frames received with a CRC error.
- Carrier – Frames aborted because of the carrier sense error, no carrier, or loss of carrier.
- Length – Frames received with the length error and the length type field that does not match the frame size.

The CRC and length error counters track receive errors, and the carrier error counter tracks transmission errors that are caused by signal loss.

Currently, the rules are created in the default policy for the minute time base, and there are two thresholds. The supported MAPS actions are RASLOG, SNMP, and EMAIL. MAPS sends the alerts with a RASLog message and then takes any other actions configured in a rule. The indexing scheme for addressing the members in this group is similar to the addressing scheme of the front-end ports.

The following RASLog message is generated when a GE port rule is triggered due to CRC errors.

```
2023/11/01-08:10:45 (GMT), [MAPS-2062], 99, FID 128 | PORT 0/GE6, WARNING, AWING_157-49-10,
 GE Port ge6, Condition=ALL_EXT_GE_PORTS(GE_CRC/min>0), Current Value:[GE_CRC, 100 CRCs],
 RuleName=defALL_EXT_GE_PORTSCRC_0, Dashboard Category=Extension GE Port Health, Quiet Time=None.
```

Creating a CRC Rule for GE Port Monitoring

The rule creation for Gigabit Ethernet port monitoring is similar to creating rules for other ports. You can enable default policies to monitor the counters or create custom rules and policies to monitor the GE ports.

The following example creates a CRC rule for a GE port.

```
switch:admin>mapsrule --create
test_backend_port_rule_1 -group ALL_EXT_GE_PORTS -monitor GE_CRC
-timebase min -op ge -value 35 -action raslog
```

When the CRC error for any GE port is greater than or equal to 35 during one minute, the rule is triggered, and a RASLog message is generated. GE port error counters are important in debugging problems. Rules are created to monitor the port errors when there is FCIP traffic between two switches connected across an IP WAN network.

Fabric Performance Impact Monitoring Using MAPS

MAPS allows you to monitor fabrics for performance impacts, including timeouts, congestion, oversubscription, and throughput.

There are many distinct elements and layers in a fabric (applications, servers, switches, targets, LUNs, and other elements) and consequently, multiple places can cause congestion and impact the fabric performance impact. Because the behavior of each application is unique, the impact of the congestion 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.

Fabric Performance Impact (FPI) monitors latency due to the congestion or oversubscription conditions, frame loss, and throughput on ports in fabric using the MAPS policies. The FPI detects the performance impacts automatically at the frequency of 10-seconds and provides alerts based on the severity of the impact caused by credit-stall devices and congested or oversubscribed ports. The remainder of this section provides information on identifying the congestion conditions that occur in a fabric using the FPI feature.

Congestion Detection

Congestion occurs when latencies are created by slow port due to the ingress/egress speed mismatch or the credit-stalled devices act as slow-drain devices. Such congestion increases the TX queue latency (the time from a frame is queued for transmission to when it is transmitted out on the link). Typical sources of congestion are links, hosts, or attached storage devices that are responding more slowly than expected. Early congestion detection is critical to maintaining a fabric's performance because the increased latency on a switch or port can propagate through a switch to the network as a whole. The cumulative effect of latency on many individual devices on a link can degrade the performance of an entire link or multiple links. While the latency for each device might not be a problem, the presence of too many flows with significant latency passing through a link could become a problem.

Congestion occurs when frames enter the fabric faster than they exit the fabric. As a result, frames build up in the fabric's switches while waiting for transmission. This causes traffic to move through the fabric to slow down or become congested. Congestion can occur on device links and inter-switch links (ISLs). The backpressure from a congested port in the fabric can cause traffic to slow down on upstream ISLs. This congestion spreading can cause traffic from unrelated flows that use the same ISL to slow down the victim flows.

Congestion is a condition that occurs when the persistent backpressure affects traffic flow(s). In congestion conditions, quality of service can erode with queuing delays, frame loss, or blocking of new connections.

MAPS FPI monitoring enables you to identify the congestion impacts using the following conditions:

- Credit-Stalled Devices
- Port Oversubscription

Credit-Stalled Devices

When traffic moves from device A to device B and device B fails to acknowledge and return buffer credits as required, the fabric holds frames for an excessive amount of time. The high-latency device B can cause application performance degradation, and it is referred to as a credit-stalled device.

Generally, a device can behave as a credit-stalled device under the following scenarios:

- The device is degraded, defective, or contains the aged hardware or firmware (slow-drain device).
- The device performs normally in a heavy traffic condition, but when it is overloaded with a high number of requests, the device might not be able to get credits equally to all requests.
- Whenever a fabric device or an ISL link processes more data than it can handle, backpressure occurs due to the excess data. Backpressure from a congested port can slow down peers in the fabric on upstream ISLs. This congestion can cause traffic from unrelated flows that use the same ISL to also slow down. For more information, see Port Oversubscription.

A credit-stalled device is one of the causes of fabric congestion. Fabric congestion is created when a receiving device stops issuing R_RDY buffer credits to the transmitting device for an abnormally long time and might not issue a credit until a frame loss or link reset occurs. As the transmitting device cannot increment its credit count without a R_RDY, it cannot move frames from the fabric. As a result, frames destined for the credit-stalled device back up immediately in the fabric, creating congestion in the network.

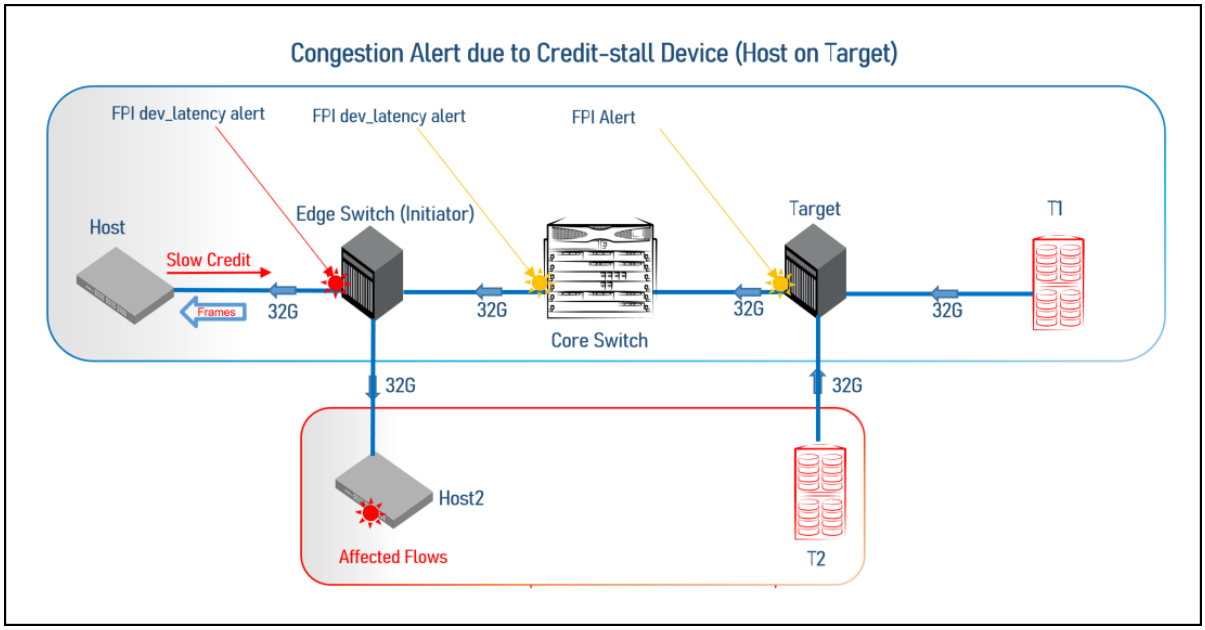

The previous figure shows the condition where a credit-stalled device (Host 1) results in a buffer backup on Host 1. Due to this, congestion occurs at the initiator of the edge switch upstream. Once all available credits are exhausted, the switch port that is connected to the device must hold additional outbound frames until a buffer credit is returned by the device. The network underperforms as a result.

When a device does not respond in time, the transmitting switch is forced to hold frames for longer periods of time, resulting in a high buffer occupancy. This in turn, results in the switch lowering the rate at which it returns buffer credits to other transmitting switches. This effect propagates through switches (and potentially multiple switches when devices attempt to send frames to devices attached to the switch with the credit-stalled device), ultimately affecting the fabric.

Also, latency on inter-switch links (ISLs) can be caused by backpressure from latency elsewhere in the fabric. The cumulative effect of latency on many individual devices can result in slowing down the link. The link itself might be producing latency if it is a long-distance link with distance delays or if too many flows are using the same ISL. Although each device might not appear to be a problem, the presence of too many flows with some level of latency across a single ISL or trunked ISL might become a problem. Latency on an ISL can ripple through other switches in the fabric and affect unrelated flows.

The following three conditions occur because of the above congestion scenarios:

- Transmit queue latency
- Buffer credit zero latency
- Frame loss

MAPS FPI monitors congestion over different time windows to determine the impact on the fabric. If congestion on these ports is severe enough to impact fabric performance, the port is classified into one of two states:

- **IO_PERF_IMPACT** - This state indicates a moderate level of latency. In this state, device-based latencies can negatively impact the overall network performance. This value is calculated using buffer credit zero or transient queue latency counters.
- **IO_FRAME_LOSS** - This state indicates a severe level of latency. In this state, frame timeouts either have occurred or are likely to occur. This value is calculated using transient queue latency only.

MAPS FPI monitors the latency on the following ports:

- F_Ports
- Trunked F_Ports
- E_Ports
- EX_Ports
- Trunked E_Port
- N_Ports (in AG mode)
- Trunked N_Port (in AG mode)

Fabric Performance Impact (FPI) monitoring uses frame delivery timeouts and credit latency to identify credit-stalled devices. MAPS generates FPI alerts when either the congestion levels or port latencies meet or exceed the thresholds that are specified in MAPS rules. Alerts might be generated through SNMP, RASLog, Email, or whatever you set as the action through MAPS for generating alerts.

The following example shows the credit-stall condition with the port details.

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

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

DB start time:                  Fri Apr 10 11:16:39 2021
Time Window:                    13:04 - 12:04
Total Credit-Stalled ports:     5
Total Oversubscribed ports:     0

2 Credit-Stall State Frequency Table:
=====================================

Port        |Current Minute State |Frame Loss   |Perf Impact  |Medium   |Low   |Info     |
-------------------------------------------------------------------------------------------
E-Port 8/35 |Perf Impact          |0            |60           |0        |0     |0        |
E-Port 7/3  |Perf Impact          |0            |60           |0        |0     |0        |
E-Port 3/9  |Perf Impact          |0            |60           |0        |0     |0        |
E-Port 8/32 |Perf Impact          |0            |51           |9        |0     |0        |
E-Port 3/31 |Medium St            |0            |0            |60       |0     |0        |

3 Oversubscription State Frequency Table:
=========================================

Port        |Current Minute State   |Frequency|
```

```
-----------------------------------------------
```

The MAPS dashboard continues to log events whether RASLog messages are set to on or off in the congestion configuration. Refer to the *Brocade Fabric OS Administration Guide* for specific command details and congestion monitoring parameters.

When the latency conditions are cleared, the **IO_LATENCY_CLEAR** state generates alerts. For more information, see Latency State Clearing.

The following example shows first the MAPS dashboard displaying the IO_PERF_IMPACT report and then the IO_LATENCY_CLEAR report. The dashboard is edited to show only Section 3.The backslash character (\) in the following example indicates a break that is inserted because the output is too long to display here as a single line.

```
switch:admin> mapsdb --show

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

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)|
\-----------------------------------------------------
\|02/19/21 22:48:01|F_Port 8/8 |IO_PERF_IMPACT        |
```

The following example shows that the MAPS dashboard display changes after the latency is cleared on F_Port 8/8.

```
switch:admin> mapsdb --show

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

```
Traffic Performance      |In operating range        |In operating range        |
Extension Health         |In operating range        |No Errors                 |
Fabric Performance Impact|In operating range        |In operating range        |

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

\|Execution Time   |Object     |Triggered Value(Units)|
\-----------------------------------------------------|
|02/19/2021:48:01  |F_Port 8/8 |IO_LATENCY_CLEAR      |
```

IO Performance Impact Monitoring

The IO_Perf_Impact state indicates a moderate level of latency. In this state, device-based latencies can negatively impact the overall network performance. This state is calculated using buffer credit zero or transient queue latency variations. The IO_Perf_Impact state involves the possible actions include RASLog, SNMP, FPIN, Email notifications, port toggling, and SDDQ (slow-drain device quarantining). See Port Toggling and Automatic VC Quarantining for specific information on these features. Default rules for these actions are included in all three default policies.

Credit Latency Monitoring

Buffer credit zero variations indicate credit latency. This variation indicates the number of times the available credit has become zero. Variations in credit zero are caused by the delay in receiving R_RDY and not transmitting the frames.

MAPS also monitors the total latency over multiple window sizes with a different threshold for each time window. As a result, MAPS can monitor various levels of latency severity simultaneously and at the VC level. When a violation occurs, the latency is reported as IO_PERF_IMPACT in the RASLog message; the message includes the bandwidth loss amount and the corresponding time window. In addition, the message specifies the actual increment of the counter as a percentage.

The following example displays a typical RASLog entry with the VC information for this condition. In this example, the bandwidth loss is 85%, and the time window is 1 second.

```
2021/03/01-13:00:00, [MAPS-2070], 2692, SLOT 1 | FID 1 | PORT 9/19, WARNING, Core_Switch, slot9 port19,
 E-Port 9/19, Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_PERF_IMPACT), Current Value:[DEV_LATENCY_IMPACT,
 IO_PERF_IMPACT, (85.1% of 5 secs in VC: 9) ], RuleName=defALL_PORTS_IO_PERF_IMPACT, Dashboard Category=Fabric
 Performance Impact, Quiet Time=15 min.
```

The monitoring window sizes and their corresponding thresholds are as follows:

- 70% of CRED_ZERO counter increment in 1 second.
- 50% of CRED_ZERO counter increment in 5 seconds.
- 30% of CRED_ZERO counter increment in 10 seconds.

The threshold values given above are default values and customizable using the FPI profiles. For more information, see FPI Profile Management.

Transmit Queue Latency Monitoring

MAPS monitors the port transmit queue latency (TXQ) for every port to identify ports that might be causing congestion in the SAN fabric so that you can take a corrective action before the congestion affects other parts of the fabric.

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. The TXQ latency is calculated at the VC level. MAPS monitors the maximum TXQ latency per VC at one-second interval. The latency of the VC having the highest latency time 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 given port.

The IO_PERF_IMPACT state is set for a port when the transmit queue latency is between the low threshold and high threshold values. If the time between successive credit returns by the device is between a few hundred microseconds to tens of milliseconds, then the device exhibits mild to moderate latency since this delay is typically not enough to cause the frame loss. Moderate device latency indicated by the IO_PERF_IMPACT state is defined as latency that is not severe enough to cause the frame loss. Moderate device latency state is set in the range of 10-ms to 80-ms queue latency at the device port and upstream ISL port. Moderate device latency causes a drop in application performance but typically does not cause frame drops or I/O failures.

The effect of moderate device latency on host applications may still be profound, based on the average disk service times expected by the application. Mission-critical applications that expect average disk service times of, for instance, 10-ms, are severely affected by storage latency more than the expected service times.

The following example displays typical RASLog with the VC information created when the IO_PERF_IMPACT state is set.

```
2023/04/12-16:21:00, [MAPS-1003], 8918, FID 128 | PORT 0/18, WARNING, SAN_FABRIC_1, port18, E-Port 18,
 Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_PERF_IMPACT), Current Value:[DEV_LATENCY_IMPACT, IO_PERF_IMPACT,
 (39 ms Frame Delay in VC: 3) ], RuleName=defALL_PORTS_IO_PERF_IMPACT, Dashboard Category=Fabric Performance
 Impact, Quiet Time=15 min.
```

The following example of the `mapsdb --show` command shows that port 18 has a latency problem. In this example, the output has been trimmed to focus on the explicit rule, and the backslash character (\) indicates a break inserted because the output is too long to display here as a single line.

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

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

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

Category(Violation Count)|RepeatCount|Rule Name                 |Execution Time   |Object           |
Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------------------------------------------
Fabric Performance Impact|1          |defALL_PORTS_IO_PERF_IMPACT|04/12/23 17:36:00|E-Port 18        |
IO_PERF_IMPACT(49)    |
```

Frame Loss Monitoring

The IO_FRAME_LOSS state indicates a severe level of latency. In this state, frame timeouts either have occurred or are likely to occur. MAPS monitors for the Frame Loss state that indicates and reports the following two types of frame loss events:

- Frame timeouts (C3TXTO)
- High Transmit Queue Latency (causes severe delays but without frame timeouts)

Frame Timeouts Monitoring

MAPS monitors frame loss, which indicates the count of discarded frames due to timeout of the transmission hold time. This is represented as Class 3 frame timeout errors (C3TXTO) on individual ports.

When a timeout is detected on a port, MAPS reports it 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.

Frame loss at an F_Port is explicitly tracked at the F_Port. However, frame loss at an E_Port might be caused by high R_RDY delays on F_Ports (which could be affected by various conditions, such as edge hold time, number of E_Ports, and other conditions). The delays might not be significant enough to cause timeouts on the F_Ports, but they could cause backups and timeouts at the E_Ports.

The following example displays a typical RASLog entry for this condition with 1150 C3 transmit timeouts that occurred in 10 seconds.

```
2021/04/30-20:15:59, [MAPS-1001], 2/2, SLOT 1 | FID 1 | PORT 1/19, 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, Quiet Time=15 min.
```

High Transmit Queue Latency Monitoring

MAPS FPI puts an F_Port in the IO_FRAME_LOSS state if an actual timeout occurs on the port. This timeout is indicated with the C3TX_TO counters on the port. In addition, FPI calculates the average R_RDY delay to predict the likely timeout on an F_Port. If the average R_RDY delay is greater than 80 milliseconds that is classified into the severe device latency, and the corresponding F_Port is set to the IO_FRAME_LOSS state.

Severe device latency results in frame loss, which triggers the host SCSI stack to detect failures and retry I/Os. This process can take tens of seconds (possibly as long as 30 to 60 seconds), which can cause a very noticeable application delay and potentially result in application errors. If the time between successive credit returns by the device is in excess of 80 ms, the device is exhibiting severe latency. When a device exhibits severe latency, the switch is forced to hold frames for excessively long periods of time (on the order of hundreds of milliseconds). When this time becomes greater than the established timeout threshold, the switch drops the frame according to Fibre Channel standards. Frame loss in switches is also known as Class3 (C3) discards or timeouts.

Since the effect of device latency often spreads through the fabric, frames can be dropped due to timeouts. In this case, frames are dropped not just on the F_Port to which the misbehaving device is connected but also on E_Ports that are carrying traffic to the F_Port. Dropped frames typically cause I/O errors that result in a host retry, resulting in significant decreases in application performance. The implications of this behavior are compounded and exacerbated by the fact that frame drops on the affected F_Port (device) result not only in I/O failures to the credit-stalled device (which are expected) but also on E_Ports, which might cause I/O failures for unrelated traffic flows involving other hosts (which typically are not expected).

The following example displays typical RASLogs that are created when the IO_FRAME_LOSS states is set.

```
2021/04/12-16:25:00, [MAPS-1001], 8919, FID 128 | PORT 0/18, CRITICAL, SAN_FABRIC_1, port18, E-Port 18,
 Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_FRAME_LOSS), Current Value:[DEV_LATENCY_IMPACT, IO_FRAME_LOSS,
 (98 ms Frame Delay in VC: 3) ], RuleName=defALL_PORTS_IO_FRAME_LOSS, Dashboard Category=Fabric Performance
 Impact, Quiet Time=15 min.
```

Latency State Clearing

An IO_LATENCY_CLEAR state indicates that backpressure caused by the congested port, or the credit-stall device, indicated either by IO_perf_impact state or by IO_Frame_Loss state, has been cleared. If there is a rule in the active policy to monitor the IO_LATENCY_CLEAR state, a RASLog message is posted. The following example shows an example of the RASLog message with the IO_LATENCY_CLEAR state.

```
2023/04/06-19:11:01:271770, [MAPS-1004], 279/118, FID 128 | PORT 0/43, INFO, Target_Edge_SW, port43, E-
Port 43, Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_LATENCY_CLEAR), Current Value:[DEV_LATENCY_IMPACT,
 IO_LATENCY_CLEAR, (Duration: 10s)], RuleName=defALL_PORTS_IO_LATENCY_CLEAR, Dashboard Category=Fabric
 Performance Impact, Quiet Time=15 min.
```

From the Fabric OS v9.2.0 release, the IO_LATENCY_CLEAR state notification is enhanced to display the latency duration. The port remains in the congestion state once it enters into a congested state. The RASLog alerts with IDs MAPS-1005 and MAPS-1006 are enhanced to display the Total Congestion Duration and the Maximum Congestion Duration after the Quiet Time expires.

Fabric OS v9.2.0 is enhanced to include configured port name details in RASLOG messages and other MAPS alerts for SFP rules monitoring.

Supported MAPS Actions

When a timeout is seen on a port, the state of that port is changed to IO_FRAME_LOSS. MAPS supports FENCE/ DECOM actions for default rules and RoR rules that are associated with the DEV_LATENCY_IMPACT monitoring system. When you configure the default rules on the DEV_LATENCY_IMPACT monitoring system with the disruptive actions, the mutual exclusion is validated. For example, if both SDDQ and FENCE actions are configured, an error message is displayed, and the rule is not configured.
For more information about supported MAPS actions, see MAPS Scale Numbers.

Troubleshooting Credit-Stalled Devices

After you determine that the cause of congestion is from credit-stalled devices, you can employ the mitigation techniques summarized in this section. However, to completely resolve the effect of congestion due to a credit-stalled device, the only solution is to stop credit-stalled behavior on all devices connected to the switch.

In summary, to reduce the effects of the credit-stalled device on fabric congestion and other device traffic until the issue with the credit-stalled device is resolved, perform the following tasks:

- Reset the credit-stalled device.
- Isolate the credit-stalled device.
- Remove the credit-stalled device.

Use the following Fabric OS features to mitigate the effects of the credit-stalled device until the problem can be remedied or the device is replaced:

- Isolate traffic from the credit-stalled device using the Brocade Slow-Drain Device Quarantine (SDDQ) feature. This action is automatically triggered when FPI detects an F_Port in either the IO_PERF_IMPACT state or the IO_FRAME_LOSS state.
- Isolate traffic from the credit-stalled device using QoS zoning. This requires that you manually configure QoS zones to manage traffic priority between specific host-target pairs.
- Enable the Brocade port toggle, fencing, or decommissioning features.

For more information about the congestion and troubleshooting, refer to the *Brocade Fabric Congestion Troubleshooting Guide*.

Monitoring Stuck Virtual Channels Using FPI

Stuck virtual channels (VC) is a condition in which virtual channels experience credit loss and do not have credits to transmit frames for an extended period. The stuck VC condition is monitored by FPI.

From Fabric OS v9.0.x release, the AN-1010 RASLog message is replaced by the MAPS FPI alerts that occur due to any of the following conditions:

- C3TXTO timeouts
- High TXQ latency
- Credit zero value of >90% in 1 sec

The following is an example of the RASLog alert that occurs due to the C3TXTO timeout:

```
2023/02/11-12:01:00:209538, [MAPS-1001], 1143/1081, FID 128 | PORT 0/45, CRITICAL, sw062080-G720, port45,
 E-Port 45, Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_FRAME_LOSS), Current Value:[DEV_LATENCY_IMPACT,
 IO_FRAME_LOSS, (3485 C3TX Timeouts) ], RuleName=defALL_PORTS_IO_FRAME_LOSS, Dashboard Category=Fabric
 Performance Impact, Quiet Time=15 min.
```

The following is an example of the RASLog alert that occurs due to the high TXQ latency:

```
2023/02/12-16:25:00, [MAPS-1001], 8919, FID 128 | PORT 0/18, CRITICAL, SAN_FABRIC_1, port18, E-Port 18,
 Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_FRAME_LOSS), Current Value:[DEV_LATENCY_IMPACT, IO_FRAME_LOSS,
 (98 ms Frame Delay in VC: 3) ], RuleName=defALL_PORTS_IO_FRAME_LOSS, Dashboard Category=Fabric Performance
 Impact, Quiet Time=15 min.
```

The following is an example of the RASLog alert that occurs due to the credit zero value of >90% in 1 second:

```
2023/02/06-13:00:00, [MAPS-2070], 2692, SLOT 1 | FID 1 | PORT 9/19, WARNING, Core_Switch, slot9 port19,
 E-Port 9/19, Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_PERF_IMPACT), Current Value:[DEV_LATENCY_IMPACT,
 IO_PERF_IMPACT, (91.1% of 1 sec in VC: 9) ], RuleName=defALL_PORTS_IO_PERF_IMPACT, Dashboard Category=Fabric
 Performance Impact, Quiet Time=15 min.
```

Port Oversubscription

Oversubscription is a condition that occurs when the aggregated ingress throughput rate for a port or link is greater than the egress traffic going out of the switch through an egress port.

If the ingress throughput is 16G, but the aggregate bandwidth of the egress links is less, for example, 8G, there will be congestion due to oversubscription on 16G links.

MAPS FPI monitoring enables you to identify the impacted ports due to oversubscription. The FPI monitors the port oversubscription on the following ports:

- F_Port
- Trunked F_Port
- E_Port
- Trunked E_Port
- EX_Port
- N_Port (in AG mode)
- Trunked N_Port (in AG mode)

MAPS FPI detects the port oversubscription using the following three parameters across a 10-second window:

- **TXQ Latency** – The time a frame is queued for transmission to the time that it is transmitted out on the link. The TXQ latency for a frame will be impacted by the number of frames that are already queued in front of it.
- **TX Bandwidth Utilization** – The percentage of available port (or trunk) bandwidth being used for transmitted traffic.
- **Buffer Credit Zero Counter** – The number of times the available credit has become zero due to the delay in receiving the R_RDY and not being able to transmit the frames.

The high TXQ latency and high TX utilization without the buffer credit zero (credit delay) issue indicates that there is a port oversubscription condition on the port. In some fabric configurations, IO_PERF_IMPACT and IO_FRAME_LOSS alerts might occur on multiple upstream ports due to oversubscription on a port. While the impact on the upstream ports is similar to the congestion spreading caused by a credit-stalled device, with port oversubscription, the downstream ports do not show the abnormal latency due to credit-stall. These ports exhibit only high bandwidth utilization. The periods of high bandwidth utilization on a downstream port can be used to identify devices that are suspected of being oversubscribed.

You can use the following default FPI rules to monitor the oversubscription in the fabric and clear the oversubscription state:

- defALL_PORTS_OVERSUBSCRIBED
- defALL_PORTS_OVERSUBSCRIPTION_CLEAR

The default Quiet Time for the FPI rules is 15-minutes. These rules are supported by all of the following default MAPS policies:

- dflt_conservative_policy
- dflt_aggressive_policy
- dflt_moderate_policy
- dflt_base_policy

The oversubscription conditions generally occur in the SAN configurations as a result of the following SAN design-specific issues:

- Link Speed Mismatch
- Fan-In/Fan-Out Ratio

Link Speed Mismatch

This section describes the link speed mismatch between initiator and target ports.

An improperly designed SAN fabric where the design considerations do not match the actual link throughput can cause oversubscription issues. Whenever the fabric has insufficient link capacity and the link is no longer capable of transmitting more frames, the resulting oversubscription can cause congestion-related issues, for example, quality of service (QoS) and I/O failures.

The following illustration describes the congestion from the read oversubscription that is caused by a link speed mismatch between a 32G storage device sending to a 8G host.

Figure 2: Oversubscription due to a Device Port Link Speed Mismatch

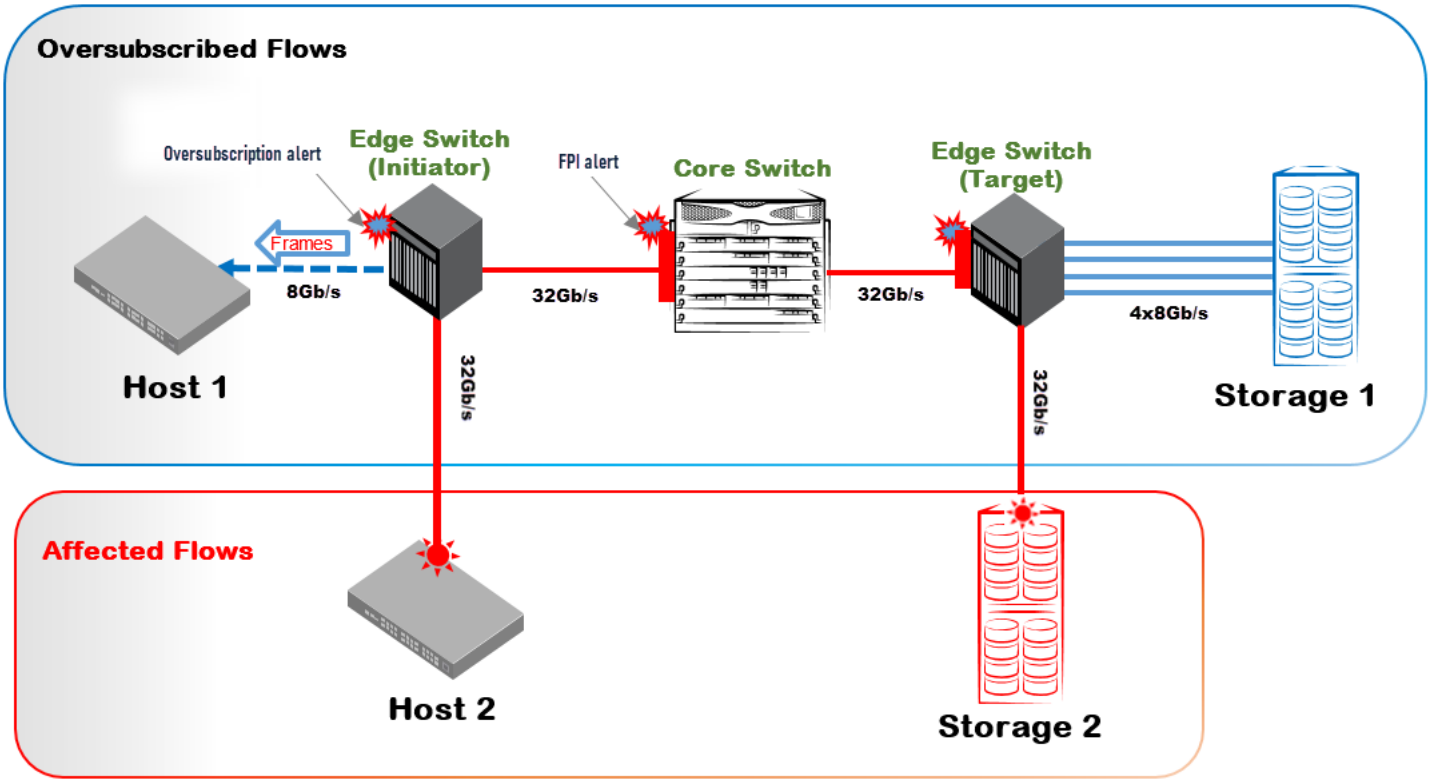

In the preceding diagram, Host 1 is connected to the fabric through an 8G HBA, and Host 1 requests a large read I/O from Storage 1. Similarly, Host 2 with a 32G HBA demands a large read I/O from Storage 2. Both Storage 1 and Storage 2 are 32G storage devices. The traffic flow from Storage 1 overwhelms the slower 8G link from the edge initiator switch to Host 1 and results in oversubscription. The edge initiator switch slows down the incoming traffic from the core switch, and this creates the traffic performance due to congestion. It creates the IO_PERF_IMPACT condition on the link that connects these two switches. Subsequently, the backpressure on the core switch is propagated to the E_Port of the target edge switch.

Edge Host Switch Alert Examples

The following example shows the RASLog alert that is received for the F_Port 7/35 on the initiator edge switch (Host) to indicate the oversubscription condition.

```
2023/02/06-06:58:00:747206, [MAPS-3038], 5671/3133, SLOT 1 | FID 1 | PORT 7/35, WARNING, Intiator_Edge_SW,
 slot7 port35, F-Port 7/35, Condition=ALL_PORTS(PORT_BANDWIDTH==OVERSUBSCRIBED), Current Value:
[PORT_BANDWIDTH, OVERSUBSCRIBED, (TXQL=464 us, TX=96.5%) ], RuleName=defALL_PORTS_OVERSUBSCRIBED, Dashboard
 Category=Fabric Performance Impact, Quiet Time=15 min.
```

The following example shows the oversubscription condition on the host port that is displayed on the corresponding MAPS dashboard output.

```
Fabric Performance Impact|1          |defALL_PORTS_OVERSUBSCRIBED|02/06/23 06:58:00|F-Port 7/35      |
OVERSUBSCRIBED        |
(7)                       |          |                           |                  |                 |
```

The following example shows the oversubscription condition that is displayed on the corresponding MAPS congestion dashboard output.

```
Congestion DB:

3 Oversubscription State Frequency Table:
=========================================

Port        |Current Minute State   |Frequency|
-----------------------------------------------
F-Port 7/35 |Oversubscribed         |5        |
```

Core Switch Alert Examples

The following example shows the MAPS RASLog alert that is received for E_Port 9/19 on the core switch that indicates the traffic performance due to congestion, that is, the IO_PERF_IMPACT condition.

```
2023/02/06-13:00:00, [MAPS-2070], 2692, SLOT 1 | FID 1 | PORT 9/19, WARNING, Core_Switch, slot9 port19,
 E-Port 9/19, Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_PERF_IMPACT), Current Value:[DEV_LATENCY_IMPACT,
 IO_PERF_IMPACT, (53.1% of 5 secs in VC: 9) ], RuleName=defALL_PORTS_IO_PERF_IMPACT, Dashboard Category=Fabric
 Performance Impact, Quiet Time=15 min.
```

The following example shows the IO_PERF_IMPACT condition that is displayed in the corresponding MAPS dashboard output for the core switch.

```
Fabric Performance Impact|1          |defALL_PORTS_IO_PERF_IMPACT|02/06/23 13:00:00|E-Port 9/19      |
IO_PERF_IMPACT        |
(1)                      |           |                           |                  |                 |
```

The following example shows the credit-stall condition that is displayed on the corresponding MAPS congestion dashboard output of the core switch.

```
2 Credit-Stall State Frequency Table:
=====================================

Port        |Current Minute State |Frame Loss   |Perf Impact  |Medium   |Low      |Info     |
-------------------------------------------------------------------------------------------
E-Port 9/19 |Perf Impact          |0            |2            |0        |0
```

Edge Target Switch Alert Examples

The following example shows the MAPS RASLog alert that is received for the E_Port 43 on the IO_PERF_IMPACT condition. As a result, backpressure has been propagated to the E_Port of the target edge switch.

```
2023/02/06-19:00:01:284580, [MAPS-1003], 275/114, FID 128 | PORT 0/43, WARNING, Target_Edge_SW, port43,
 E-Port 43, Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_PERF_IMPACT), Current Value:[DEV_LATENCY_IMPACT,
 IO_PERF_IMPACT, (57.0% of 10 secs in VC: 9) ], RuleName=defALL_PORTS_IO_PERF_IMPACT, Dashboard
 Category=Fabric Performance Impact, Quiet Time=15 min.
```

The following example shows the IO_PERF_IMPACT condition that is displayed in the corresponding MAPS dashboard output for the target edge switch.

```
Fabric Performance Impact|1          |defALL_PORTS_IO_PERF_IMPACT|02/06/23 19:00:01|E-Port 43        |
IO_PERF_IMPACT        |
(1)                      |           |                           |                  |                 |
               |
```

The following example shows the credit-stall condition that is displayed on the corresponding MAPS congestion dashboard output of the target edge switch.

```
2 Credit-Stall State Frequency Table:
```

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

Port        |Current Minute State |Frame Loss   |Perf Impact  |Medium   |Low      |Info     |
-------------------------------------------------------------------------------------------
E-Port 43   |Perf Impact          |0            |1            |0        |0        |0        |
```

The following example shows the RASLog alert when the oversubscription state is cleared for the F_Port 7/35 of the host.

```
2023/02/06-07:17:00:769621, [MAPS-3039], 5689/3151, SLOT 1 | FID 1 | PORT 7/35, INFO, Intiator_Edge_SW,
 slot7 port35, F-Port 7/35, Condition=ALL_PORTS(PORT_BANDWIDTH==OVERSUBSCRIPTION_CLEAR), Current Value:
[PORT_BANDWIDTH, OVERSUBSCRIPTION_CLEAR], RuleName=defALL_PORTS_OVERSUBSCRIPTION_CLEAR, Dashboard
 Category=Fabric Performance Impact, Quiet Time=15 min.
```

The following example shows the corresponding MAPS dashboard output when the oversubscription state is cleared on the host.

```
Fabric Performance Impact|1          |defALL_PORTS_OVERSUBSCRIPTI|02/06/23 07:17:00|F-Port 7/35      |
OVERSUBSCRIPTION_CLEAR|(56)  |           |ON_CLEAR                   |                 |             |
```

The following example shows the RASLog alert for the E_Port 9/19 when the credit-stall condition is cleared on the core switch.

```
2023/02/06-13:21:00, [MAPS-2071], 2708, SLOT 1 | FID 1 | PORT 9/19, INFO, Core_Switch, slot9 port19, E-
Port 9/19, Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_LATENCY_CLEAR), Current Value:[DEV_LATENCY_IMPACT,
 IO_LATENCY_CLEAR, (Duration: 10s)], RuleName=defALL_PORTS_IO_LATENCY_CLEAR, Dashboard Category=Fabric
 Performance Impact, Quiet Time=15 min.
```

The following example shows the corresponding MAPS dashboard output when the credit-stall condition is cleared on the core switch.

```
Fabric Performance Impact|1          |defALL_PORTS_IO_LATENCY_CLE|02/06/23 13:21:00|E-Port 9/19      |
IO_LATENCY_CLEAR |(2)        |           |AR
```

The following example shows the RASLog alert for the E_Port 43 when the credit-stall condition is cleared on the target edge switch.

```
2023/02/06-19:11:01:271770, [MAPS-1004], 279/118, FID 128 | PORT 0/43, INFO, Target_Edge_SW, port43, E-
Port 43, Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_LATENCY_CLEAR), Current Value:[DEV_LATENCY_IMPACT,
 IO_LATENCY_CLEAR, (Duration: 10s)], RuleName=defALL_PORTS_IO_LATENCY_CLEAR, Dashboard Category=Fabric
 Performance Impact, Quiet Time=15 min.
```

The following example shows the corresponding MAPS dashboard output when the credit-stall condition is cleared on the target edge switch.

```
Fabric Performance Impact|1          |defALL_PORTS_IO_LATENCY_CLE|02/06/23 19:11:01|E-Port 43        |
IO_LATENCY_CLEAR           |(2)        |AR
```

Figure 3: Oversubscription due to an ISL Speed Mismatch

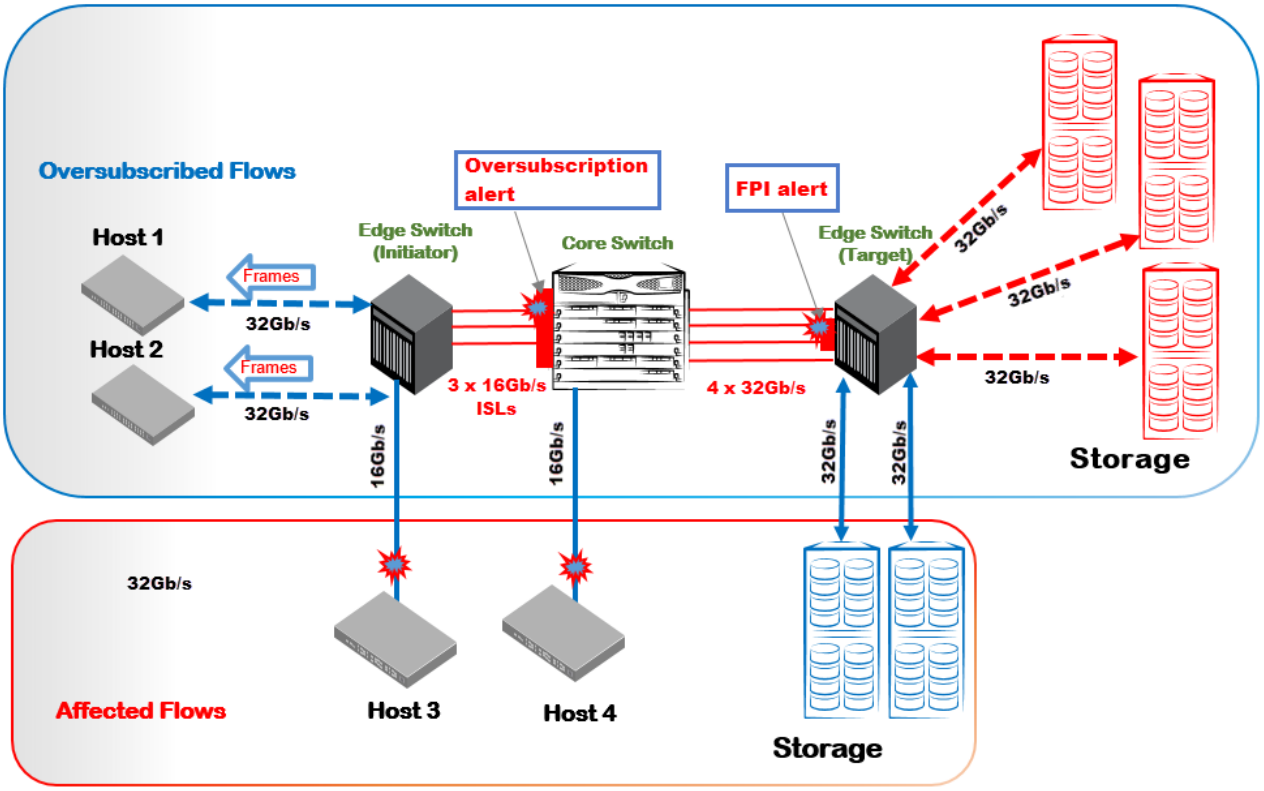

In the previous example, the 16G link speed between the edge the initiator switch and the core switch is less than the link speed between the core switch and the edge target switch, which is 32G. When the edge target switch is transmitting frames at the link speed of 32G, and the read I/O traffic of the initiator side core switch is only16G, the result is the accumulation of frames to the edge target switch. This accumulation results in the oversubscription on the initiator side core switch, which is E_Port. The oversubscription creates backpressure to the edge target switch, because the downstream links have the lower bandwidth. The Host-to-edge initiator switch needs more time to transmit the frames. When the initiator side core switch makes more read I/O requests to multiple targets simultaneously, the oversubscription condition exists.

Core Switch Alert Examples

The following example shows the RASLog alert that is received for the E_Port 11/40 on the core switch. This alert indicates that the oversubscription condition that occurred due to a bottleneck in the ISL between the core switch and the edge initiator switch.

```
2023/02/07-11:35:01:000640, [MAPS-3038], 5189/3863, SLOT 1 | FID 1 | PORT 11/40, WARNING, Core_Switch, slot11
 port40, E-Port 11/40, Condition=ALL_PORTS(PORT_BANDWIDTH==OVERSUBSCRIBED), Current Value:[PORT_BANDWIDTH,
 OVERSUBSCRIBED, (TXQL=389 us, TX=100.0%) ], RuleName=defALL_PORTS_OVERSUBSCRIBED, Dashboard Category=Fabric
 Performance Impact, Quiet Time=15 min.
```

The following example shows the oversubscription condition that is displayed on the corresponding MAPS dashboard output of the core switch.

```
Fabric Performance Impact|1          |defALL_PORTS_OVERSUBSCRIBED|02/07/23 11:35:00|E-Port 11/40     |
OVERSUBSCRIBED        |
(3)                      |           |                           |                  |                 |
```

The following example shows the oversubscription condition that is displayed on the corresponding MAPS congestiondashboard output of the core switch.

```
3 Oversubscription State Frequency Table:
=========================================

Port        |Current Minute State   |Frequency|
-----------------------------------------------
E-Port 11/40|Oversubscribed         |2        |
```

Target Edge Switch Alert Examples

The following example shows the RASLog alert that is received for the E_Port 3/49. This alert indicates that the traffic performance due to congestion which is an IO_PERF_IMPACT condition.

```
2023/02/07-05:41:00:735675, [MAPS-2070], 6880/4300, SLOT 1 | FID 1 | PORT 3/49, WARNING, Target_Edge_SW,
 slot3 port49, E-Port 3/49, Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_PERF_IMPACT), Current Value:
[DEV_LATENCY_IMPACT, IO_PERF_IMPACT, (72.4% of 1 secs in VC: 3) ], RuleName=defALL_PORTS_IO_PERF_IMPACT,
 Dashboard Category=Fabric Performance Impact, Quiet Time=15 min.
```

The following example shows the oversubscription condition that is displayed on the corresponding MAPS dashboard output of the target edge switch.

```
Fabric Performance Impact|1          | defALL_PORTS_IO_PERF_IMPACT|02/07/23 05:37:00|E-Port 3/49     |
 IO_PERF_IMPACT       |
```

The following example shows the oversubscription condition that is displayed on the corresponding MAPS congestion dashboard output of the target edge switch.

```
2 Credit-Stall State Frequency Table:
=====================================

Port        |Current Minute State |Frame Loss   |Perf Impact  |Medium   |Low      |Info     |
-------------------------------------------------------------------------------------------
E-Port 3/49 |Medium St            |0            |2            |10       |1        |0        |
```

Fan-In and Fan-Out Ratio

Fan-in and fan-out are important in SAN design, in terms of source ports to target ports and target ports to source ports. The ratio is the number of device ports that share a single target. This is always stated from the point of view of a single entity, such as 7:1 for seven hosts utilizing a single target port.

Determining how many hosts to connect to a particular storage port can be narrowed down to three considerations: port queue depth, I/O operations per second (IOPS), and throughput. Of these three, throughput is the only network component. The total bandwidth of the connected hosts should not exceed the supported bandwidth of the target port.

The following illustration describes the 1:1 oversubscription for hosts to target ports.

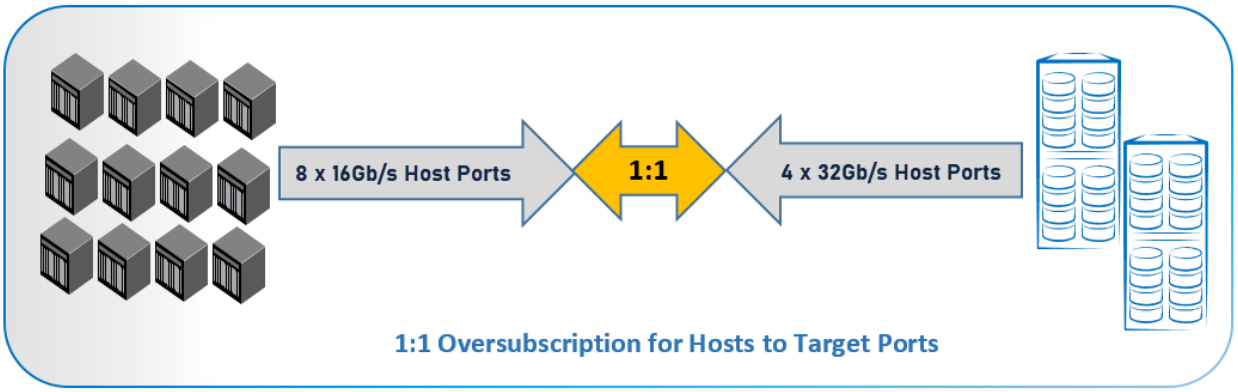

In practice, however, it is highly unlikely that all hosts perform at their maximum level at any one time. With the traditional application-per-server deployment, the host bus adapter (HBA) bandwidth is overprovisioned. However, with virtual servers, the game can change radically. Network oversubscription is built into the virtual server concept. When servers use virtualization technologies, you should reduce the network-based oversubscription proportionally. It might therefore be prudent to oversubscribe ports to ensure a balance between cost and performance.

An example of 3:1 oversubscription is shown in the following illustration.

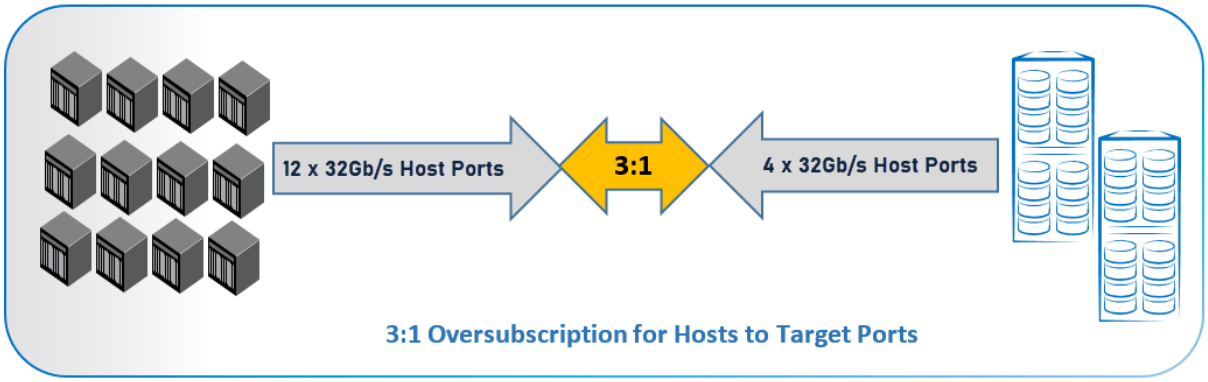

Another aspect of data flow is the fan-in ratio, in terms of a device to ISLs. This is also referred to as the fan-out ratio if viewed from the storage array perspective. The ratio is the number of device ports that share a single port, whether ISL or ICL. This is always expressed from the single entity point of view, such as 7:1 for seven hosts utilizing a single ISL or storage port. Congestion might also occur on ISLs when data flows from multiple devices exceed the capacity of the ISL, even though no individual device has requested more data than it can handle.

For determining the device-to-ISL fan-in ratio, a simple calculation is that the storage port should not be oversubscribed into the core. For example, a 32G storage port should have a 32G pipe into the core.

The following diagram shows an example of the improper fan-in ratio, where the congestion condition caused by oversubscription occurring in a fabric consisting of multiple storage targets communicating with a single host.

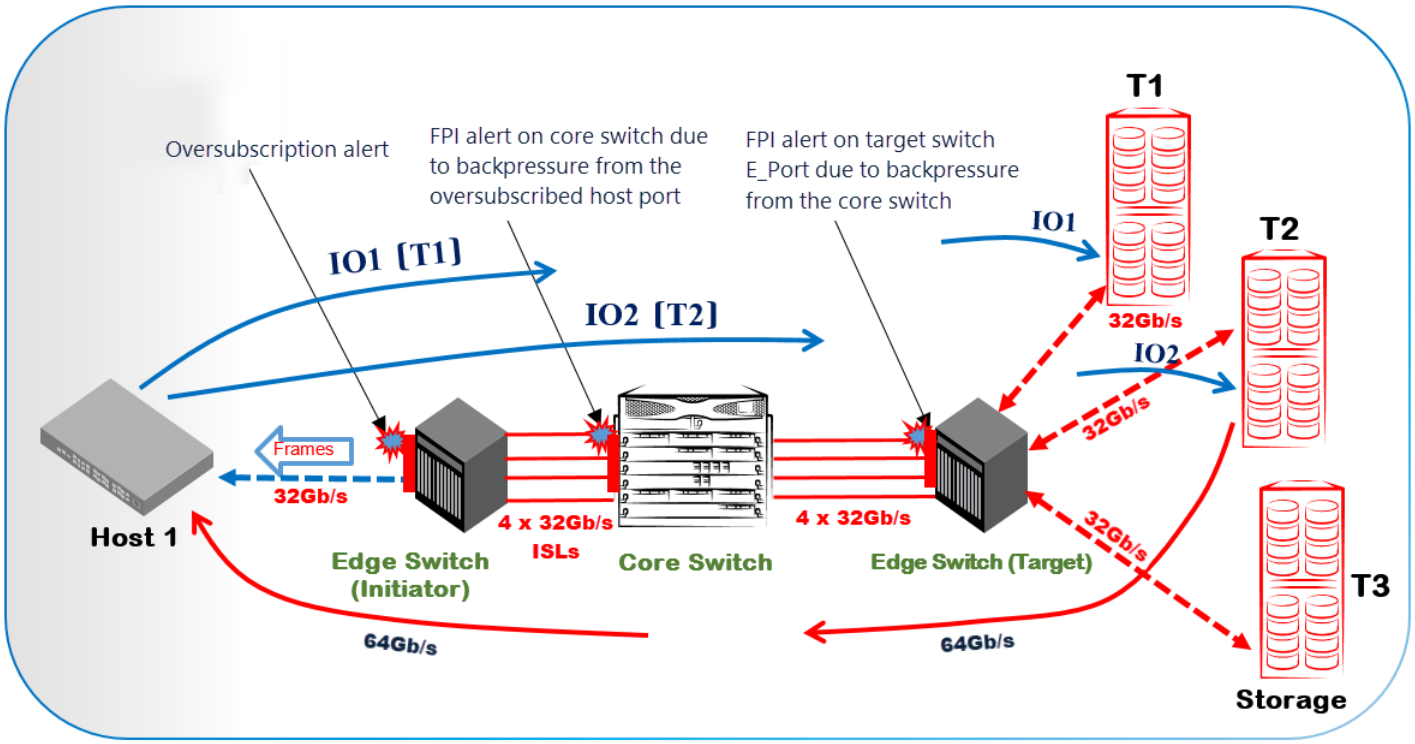

The fabric consists of a host with a 32G link that is connected to an edge initiator to target switches and multiple storage devices on 32G links. In a common scenario, when the host is demanding the read I/O from a single target, there will not be any oversubscription as both the host and targets have the same bandwidth. Let us assume that the host is performing the read I/O requests to multiple storage devices on 32G links. The host sends the read I/O requests IO1 to target 1, and IO2 to target 2. In response, the traffic is flowing from the targets to the host, and targets 1 and 2 are returning 32G data each, that is, collectively 64G data. The traffic from the faster storage devices overwhelms the slower 32G link, causing the backpressure in the oversubscribed host port, and this causes the congestion condition on the E_Port of the core switch. The congestion condition that occurred due to the backpressure that is created in the host port is propagated to the edge target switch. The examples in this section describe the MAPS alerts on the congestion conditions due to the oversubscription that is triggered on various devices in the fabric as a result of an improper fan-in and fan-out ratio.

Edge Host Switch Alert Examples

The following example shows the RASLog alert that is received for F_Port 7/35 on the initiator edge switch (Host) to indicate the oversubscription condition.

```
2023/02/06-06:58:00:747206, [MAPS-3038], 5671/3133, SLOT 1 | FID 1 | PORT 7/35, WARNING, Intiator_Edge_SW,
 slot7 port35, F-Port 7/35, Condition=ALL_PORTS(PORT_BANDWIDTH==OVERSUBSCRIBED), Current Value:
[PORT_BANDWIDTH, OVERSUBSCRIBED, (TXQL=464 us, TX=96.5%) ], RuleName=defALL_PORTS_OVERSUBSCRIBED, Dashboard
 Category=Fabric Performance Impact, Quiet Time=15 min.
```

The following example shows the oversubscription condition on the host port that is displayed on the corresponding MAPS dashboard output.

```
Fabric Performance Impact|1          |defALL_PORTS_OVERSUBSCRIBED|06/02/23 06:58:00|F-Port 7/35      |
OVERSUBSCRIBED        |
(7)                       |          |                           |                  |                 |
```

The following example shows the oversubscription condition that is displayed on the corresponding MAPS congestion dashboard output.

```
Congestion DB:

3 Oversubscription State Frequency Table:
=========================================

Port        |Current Minute State   |Frequency|
-----------------------------------------------
F-Port 7/35 |Oversubscribed         |5        |
```

Core Switch Alert Examples

The following example shows the MAPS RASLog alert received for E_Port 9/19 on the core switch that indicates the traffic performance due to congestion, that is, the IO_PERF_IMPACT condition.

```
2023/02/06-13:00:00, [MAPS-2070], 2692, SLOT 1 | FID 1 | PORT 9/19, WARNING, Core_Switch, slot9 port19,
 E-Port 9/19, Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_PERF_IMPACT), Current Value:[DEV_LATENCY_IMPACT,
 IO_PERF_IMPACT, (53.1% of 5 secs in VC: 9) ], RuleName=defALL_PORTS_IO_PERF_IMPACT, Dashboard Category=Fabric
 Performance Impact, Quiet Time=15 min.
```

The following example shows the IO_PERF_IMPACT condition that is displayed in the corresponding MAPS dashboard output for the core switch.

```
Fabric Performance Impact|1          |defALL_PORTS_IO_PERF_IMPACT|06/02/23 13:00:00|E-Port 9/19      |
IO_PERF_IMPACT        |
(1)                      |           |                           |                  |                 |
```

The following example shows the credit-stall condition that is displayed on the corresponding MAPS congestion dashboard output of the core switch.

```
2 Credit-Stall State Frequency Table:
=====================================

Port        |Current Minute State |Frame Loss   |Perf Impact  |Medium   |Low      |Info     |
-------------------------------------------------------------------------------------------
E-Port 9/19 |Perf Impact          |0            |2            |0        |0
```

Edge Target Switch Alert Examples

The following example shows the MAPS RASLog alert that is received for the E_Port 43 on the IO_PERF_IMPACT condition. This indicates that the backpressure is propagated to the E_Port of the target edge switch.

```
2023/04/06-19:00:01:284580, [MAPS-1003], 275/114, FID 128 | PORT 0/43, WARNING, Target_Edge_SW, port43,
 E-Port 43, Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_PERF_IMPACT), Current Value:[DEV_LATENCY_IMPACT,
 IO_PERF_IMPACT, (57.0% of 10 secs in VC: 9) ], RuleName=defALL_PORTS_IO_PERF_IMPACT, Dashboard
 Category=Fabric Performance Impact, Quiet Time=15 min.
```

The following example shows the IO_PERF_IMPACT condition that is displayed in the corresponding MAPS dashboard output for the target edge switch.

```
Fabric Performance Impact|1          |defALL_PORTS_IO_PERF_IMPACT|06/02/23 19:00:01|E-Port 43        |
IO_PERF_IMPACT        |
(1)                      |           |                           |                  |                 |
               |
```

The following example shows the credit-stall condition that is displayed on the corresponding MAPS congestion dashboard output of the target edge switch.

```
2 Credit-Stall State Frequency Table:
```

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

Port        |Current Minute State |Frame Loss   |Perf Impact  |Medium   |Low      |Info     |
-------------------------------------------------------------------------------------------
E-Port 43   |Perf Impact          |0            |1            |0        |0        |0        |
```

The following example shows the RASLog alert when the oversubscription state is cleared for the F_Port 7/35 of the host.

```
2023/02/06-07:17:00:769621, [MAPS-3039], 5689/3151, SLOT 1 | FID 1 | PORT 7/35, INFO, Intiator_Edge_SW,
 slot7 port35, F-Port 7/35, Condition=ALL_PORTS(PORT_BANDWIDTH==OVERSUBSCRIPTION_CLEAR), Current Value:
[PORT_BANDWIDTH, OVERSUBSCRIPTION_CLEAR], RuleName=defALL_PORTS_OVERSUBSCRIPTION_CLEAR, Dashboard
 Category=Fabric Performance Impact, Quiet Time=15 min.
```

The following example shows the corresponding MAPS dashboard output when the oversubscription state is cleared on the host.

```
Fabric Performance Impact|1          |defALL_PORTS_OVERSUBSCRIPTI|06/02/23 07:17:00|F-Port 7/35      |
OVERSUBSCRIPTION_CLEAR|(56)  |           |ON_CLEAR                   |                 |             |
```

The following example shows the RASLog alert for the E_Port 9/19 when the credit-stall condition is cleared on the core switch.

```
2023/02/06-13:21:00, [MAPS-2071], 2708, SLOT 1 | FID 1 | PORT 9/19, INFO, Core_Switch, slot9 port19, E-
Port 9/19, Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_LATENCY_CLEAR), Current Value:[DEV_LATENCY_IMPACT,
 IO_LATENCY_CLEAR, (Duration: 10s)], RuleName=defALL_PORTS_IO_LATENCY_CLEAR, Dashboard Category=Fabric
 Performance Impact, Quiet Time=15 min.
```

The following example shows the corresponding MAPS dashboard output when the credit-stall condition is cleared on the core switch.

```
Fabric Performance Impact|1          |defALL_PORTS_IO_LATENCY_CLE|06/02/23 13:21:00|E-Port 9/19      |
IO_LATENCY_CLEAR |(2)        |           |AR
```

The following example shows the RASLog alert for the E_Port 43 when the credit-stall condition is cleared on the target edge switch.

```
2023/02/06-19:11:01:271770, [MAPS-1004], 279/118, FID 128 | PORT 0/43, INFO, Target_Edge_SW, port43, E-
Port 43, Condition=ALL_PORTS(DEV_LATENCY_IMPACT==IO_LATENCY_CLEAR), Current Value:[DEV_LATENCY_IMPACT,
 IO_LATENCY_CLEAR, (Duration: 10s)], RuleName=defALL_PORTS_IO_LATENCY_CLEAR, Dashboard Category=Fabric
 Performance Impact, Quiet Time=15 min.
```

The following example shows the corresponding MAPS dashboard output when the credit-stall condition is cleared on the target edge switch.

```
Fabric Performance Impact|1          |defALL_PORTS_IO_LATENCY_CLE|06/02/23 19:11:01|E-Port 43        |
IO_LATENCY_CLEAR           |(2)        |AR
```

Troubleshooting Oversubscription

Congestion due to oversubscription is primarily a result of improper SAN design. When oversubscription is identified as the source of congestion, the problem should be addressed by correcting the design.

To mitigate the effects of congestion due to oversubscription, consider the following recommendations:

- Reduce link speed mismatch ratios between the target and the host end device.
- Reduce fan-in.
- Reduce queue depths.
- Add ISLs.

For more information about the oversubscription and troubleshooting, refer to the *Brocade Fabric Congestion Troubleshooting Guide*.

Slow-Drain Device Quarantining

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

The following figure illustrates how the slow-draining device functions. In the figure, the inability of the slow-draining device (1) to clear frames quickly enough is causing backpressure not just to the edge switch or AG device (2), but also to the core switch (3) and the target (4) that is trying to send data to the slow-draining device.

Figure 4: Slow Drain Dataflow

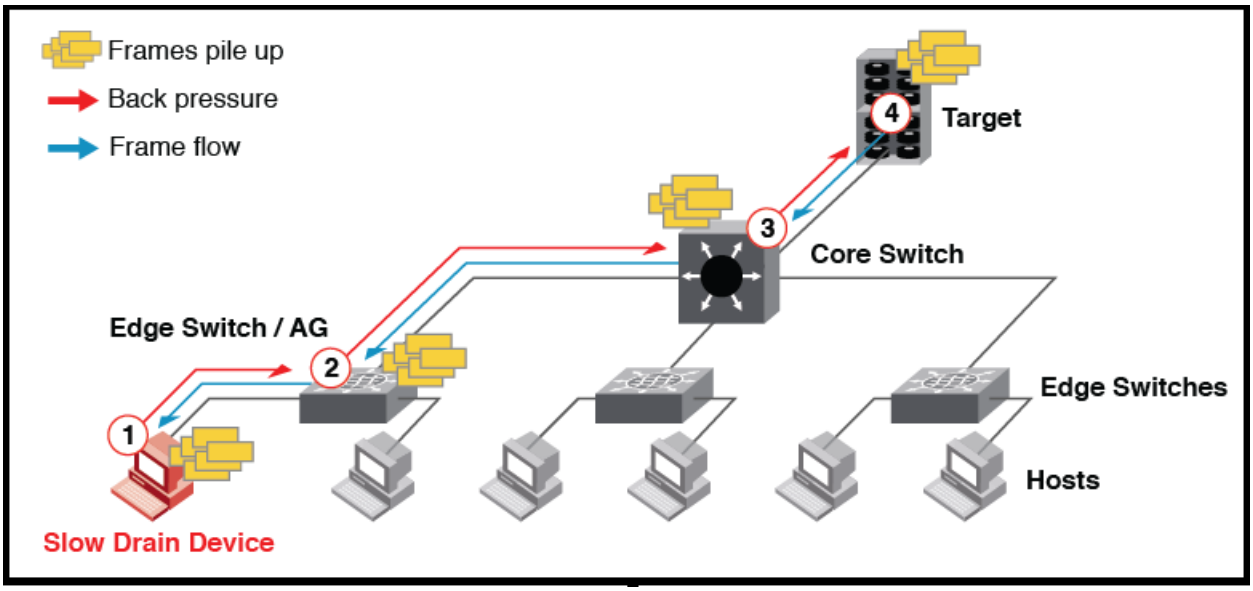

To isolate the traffic destined to slow drain the device or port, the Slow-Drain Device Quarantine (SDDQ) feature of Fabric Performance Impact (FPI) quarantines the port. SDDQ moves all the traffic that is destined for the slow-drain devices to a low-priority VC and adds the port to the predefined ALL_QUARANTINED_PORTS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

The following example displays the RASLog messages for SDDQ:

```
[NS-1026], 372, FID 128, INFO, G620, Local domain 1, port index 40: A zoned device 0x012800 has been
 quarantined.
[NS-1024], 373, FID 128, INFO, G620, Flow from device 0x012200 to quarantined device 0x012800 of Local domain
 1 has been successfully quarantined.
[NS-1024], 374, FID 128, INFO, G620, Flow from device 0x012300 to quarantined device 0x012800 of Local domain
 1 has been successfully quarantined.
[MAPS-1022], 375, FID 128 | PORT 0/40, WARNING, G620, Port 40 (Port index 40) has been marked as Slow Drain
 Device.
[AG-1086], 513, FID 128 | PORT 0/8, INFO, G610, Slow drain device 0x011205 connected to F-port (8) have been
 quarantined.
[AG-1087], 528, FID 128 | PORT 0/8, INFO, G610, Slow drain device 0x011205 connected to F-port (8) have been
 unquarantined.
```

For more information about RASLog messages, refer to the *Brocade Fabric OS Message Reference Guide*.

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

Prerequisites

You must meet the following requirements before you use the SDDQ feature:

- Quality of Service (QoS) must be enabled to use the SDDQ feature.
- The Fabric Vision® license must be installed on the switch to use the SDDQ feature. Intermediate switches do not need the Fabric Vision license for SDDQ, but they must be enabled with QoS on all inter-switch links (ISLs).

Important Points on Slow-Drain Device Quarantining

Consider the following points when you use the SDDQ action:

- SDDQ is disabled by default on installation.
- SDDQ is supported on N_Port ID Virtualization (NPIV) devices that are connected to F_Ports.
- SDDQ is not supported in the Fibre Channel Routing (FCR) backbone fabrics. Although SDDQ can be supported on devices within an FCR edge fabric, the edge fabric will not apply SDDQ to the flows for any FCR-imported devices.
- SDDQ is supported on F_Ports and F_Port trunks that are directly connected to an HBA or enabled with the AG mode.
- The SDDQ and UNQUAR actions are associated with the DEV_LATENCY_IMPACT monitoring system but not with the state. In this case, SDDQ and UNQUAR actions are configured with frame_loss, perf_impact, and latency clear.
- Although MAPS allows you to enable and disable SDDQ on switches, you should enable SDDQ on the entire fabric.
- No default rules are defined for the ALL_QUARANTINED_PORTS group, but you can create custom rules to monitor the ports in this group.
- When a port is marked as slow-drain, only the flows that are destined to it are shifted to the low-priority VC. Flows in the reverse direction are not affected.
- Quarantine action takes precedence over the QoS zone configuration. The traffic is moved to the low-priority VC irrespective of the zone QoS priority (high, medium, or low).
- If the device latency is due to Class 3 timeouts (C3TXTO) and the active C3TXTO rule has port fencing (port disabling) as an action, then port fencing may be performed first and quarantining will not occur.
- If there are switches in the fabric that do not have the QoS feature support, end-to-end slow-drain flow isolation might not be possible.
- The maximum number of ports that can be isolated per unit (chassis or fixed-port switch) is 32. The default value is 10. The maximum number of ports that can be isolated per unit is controlled by the chassis SDDQ limit user-configurable key, which is set using the `configurechassis` command. Refer to the *Brocade Fabric OS Command Reference Manual* for more information.
- To restore the default value to 10 for the chassis SDDQ limit, you must enter the `switchdisable` command and then enter the `chassisdefault -all` command to restore the value to 10.
- The SDDQ action is blocked in the following scenarios to prevent the quarantine of many source ports:
  - The total number of ports that are zoned to the slow-drain port are set to 33 or more. The total number of ports does not include slow-drain ports.
  - The defzone is labeled as all access.

NOTE
Do not use the Port Toggling action with SDDQ because it can result an unpredictable behavior.

Unquarantine Timeout Behavior Across HA Failover

After an HA failover completes, MAPS continues to monitor for quarantined ports for the remainder of the time specified in the rule. Also, MAPS continues to perform an automatic unquarantine if the `UNQUAR` action is specified in the FPI rule

and the UNQUAR action is globally enabled in `mapsconfig` , when the congestion state is clear. For example, if a rule specified an unquarantine action for two hours and an HA failover event occurs after 30 minutes, MAPS will monitor the state for an additional 1.5 hours after the HA failover completes. All the SDDQ information is persistent across HA failover, which includes quarantine ports configuration, unquarantine timeout and monitoring window, and rules.

Enabling Slow-Drain Device Quarantining

SDDQ is enabled by having a valid Fabric Vision license, and it is activated by including the SDDQ action in the configured MAPS action list.

To enable SDDQ, perform the following steps:

1. Connect to the device and log on using an account with admin permissions.
2. Use the `mapsconfig --actions` command, and include SDDQ as one of the actions.
3. Enable a policy with rules that have the SDDQ action.

   The following example enables SDDQ as an available action:

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

Scalability Limits of SDDQ

SDDQ has the following scalability limits:

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

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

Usually, you should not enable the Slow-Drain Device Quarantining (SDDQ) feature in a FICON environment, because a FICON environment typically has a single zone for all devices. However, if you are using 1-to-1 zoning, SDDQ can help with slow-draining device issues.

Typically, in a FICON setup or other setup with a single zone, do not use the SDDQ feature.

Considerations When Using the UNQUAR Action

When using the UNQUAR action to unquarantine ports automatically, you must note the following information:

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

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

Cases of Using the UNQUAR Action

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

Case 1:

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

Case 2:

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

Disabling Slow-Drain Device Quarantining

To disable the SDDQ feature, exclude SDDQ from the configured list of MAPS actions, as described in the following procedure:

1. Connect to the device and log on 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 SDDQ from the list of available actions:

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

Confirming the Slow-Drain Status of a Device

You can use the `nsshow` , `nscamshow` , and `nodefind` commands to verify that a device is slow-draining.

The following example shows the output of the `nsshow` command where there is a slow-draining device. The identifying line is called out in this example. If there is no slow-draining device, the line does not appear.

```
switch:admin> nsshow
{
Type Pid    COS     PortName                NodeName                 TTL(sec)
N    010000;    2,3;20:00:00:05:1e:92:e8:00;20:00:00:05:1e:92:e8:00; na
    Fabric Port Name: 20:00:00:05:1e:92:e8:00
    Permanent Port Name: 20:00:00:05:1e:92:e8:00
    Port Index: 0
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
    Device link speed: 8G
    Connected through AG: Yes
    Real device behind AG: Yes
    Port Properties: SIM Port
```

```
N    014a00;    2,3;30:0a:00:05:1e:84:b5:c3;10:00:00:05:1e:84:b5:c3; na
    FC4s: FCP
    NodeSymb: [42] "dsim:ns_host:sw_5300edsim[172.26.26.188]"
    Fabric Port Name: 20:4a:00:05:1e:92:e8:00
    Permanent Port Name: 30:0a:00:05:1e:84:b5:c3
    Port Index: 74
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
    Device link speed: 8G
    Connected through AG: Yes
    Real device behind AG: Yes
N    015000;      3;10:00:00:00:00:01:00:01;10:00:00:00:00:00:01:01; na
    Fabric Port Name: 20:50:00:05:1e:92:e8:00
    Permanent Port Name: 10:00:00:00:00:01:00:01
    Port Index: 80
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
    Device link speed: 8G
    Connected through AG: Yes
    Real device behind AG: Yes
    Slow Drain Device: Yes      <== Slow-draining device identified
N    015100;      3;10:00:00:00:00:02:00:01;10:00:00:00:00:00:02:01; na
    Fabric Port Name: 20:51:00:05:1e:92:e8:00
    Permanent Port Name: 10:00:00:00:00:02:00:01
    Port Index: 81
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
    Device link speed: 8G
    Connected through AG: Yes
    Real device behind AG: Yes
    The Local Name Server has 4 entries }
```

The following example shows the output of the `nscamshow` command where there is a slow-draining device. The identifying line is called out in this example. If there was no slow-draining device, the line does not appear.

```
switch:admin> nscamshow
nscam show for remote switches:
Switch entry for 2
  state  rev    owner  cap_available
  known   v740  0xfffc01 1
  Device list: count 7
    Type Pid    COS     PortName                NodeName
    N    020a00;      3;10:00:00:00:00:04:00:01;10:00:00:00:00:00:04:01;
        Fabric Port Name: 20:0a:00:05:1e:84:98:c7
        Permanent Port Name: 10:00:00:00:00:04:00:01
```

```
        Port Index: 10
        Share Area: No
        Device Shared in Other AD: No
        Redirect: No
        Partial: No
    N    020b00;      3;10:00:00:00:00:0a:00:01;10:00:00:00:00:00:0a:01;
        Fabric Port Name: 20:0b:00:05:1e:84:98:c7
        Permanent Port Name: 10:00:00:00:00:0a:00:01
        Port Index: 11
        Share Area: No
        Device Shared in Other AD: No
        Redirect: No
        Partial: No
        Slow Drain Device: Yes      <== Slow-draining device identified
    N    021000;      3;10:00:00:00:00:05:00:01;10:00:00:00:00:00:05:01;
        Fabric Port Name: 20:10:00:05:1e:84:98:c7
        Permanent Port Name: 10:00:00:00:00:05:00:01
        Port Index: 16
        Share Area: No
        Device Shared in Other AD: No
        Redirect: No
        Partial: No
```

The following example shows the output of the `nodefind` command where there is a slow-draining device. The identifying line is called out in this example. If there was no slow-draining device, the line does not appear.

```
switch:admin> nodefind 015000
Local:
 Type Pid COS PortName NodeName SCR
 N 015000; 3;10:00:00:00:00:01:00:01;10:00:00:00:00:00:01:01; 0x00000003
    Fabric Port Name: 20:50:00:05:1e:92:e8:00
    Permanent Port Name: 10:00:00:00:00:01:00:01
    Device type: Physical Unknown(initiator/target)
    Port Index: 80
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    LSAN: No
    Slow Drain Device: Yes      <== Slow-draining device identified
    Aliases:
```

Displaying Quarantined Ports

MAPS allows you to display a list of ports that are quarantined in the ALL_QUARANTINED_GROUP.

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

1. Connect to the switch and log on using an account with admin permissions.
2. Enter the `sddquarantine --show` command 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: None
-----------------------------------------------------------

Online Quarantined Devices across the fabric
-------------------------------------------------------------------------------------
Port|PID   |PWWN                   |Locality  |Port Name      |Zone Aliases        |
-------------------------------------------------------------------------------------
480 |27e040|20:01:00:11:0d:69:d7:00|REMOTE    |-              |SB_124082_P1vport0  |
480 |27e041|20:01:00:11:0d:69:d7:01|REMOTE    |-              |SB_124082_P1vport1  |
480 |27e042|20:01:00:11:0d:69:d7:02|REMOTE    |-              |SB_124082_P1vport2  |
```

Clearing Quarantined Ports

MAPS allows you to remove ports which are quarantined from the quarantine group (ALL_QUARANTINED_GROUP). If the slow-draining behavior is not detected for a period, you can configure MAPS to automatically remove the quarantined ports from the quarantine group.

Manually Clearing Quarantined Ports

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

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

   A port is not allowed to be cleared from quarantine while the latency or frame loss condition that caused it to be quarantined persists, because the port will not be flagged again until the condition is cleared and then retriggered. However, you can override this restriction by using the `-force` keyword as part of the `sddquarantine` command. The `Ports marked as Slow-Drain Quarantined but not enforced` message appears only if the number of devices that are zoned to slow-drain port exceeds to zoned limits of 32.

   The first part of the following example uses the `sddquarantine --show` command to display all quarantined local ports and the online quarantined device information across the fabric. The second part shows the use of the `sddquarantine --clear` command to remove the ports from quarantine. Notice that port 3 was not able to be cleared using the `all` option, because it was still in either the IO_FRAME_LOSS or the IO_PERF_IMPACT condition. This port was cleared using the `-force` option. These two options can also be used together.

```
switch:admin> sddquarantine --show
-----------------------------------------------------------
Ports marked as Slow Drain Quarantined in the Local Switch:   2/1, 1/3
-----------------------------------------------------------
Ports marked as Slow Drain Quarantined but not enforced:   1/3
-----------------------------------------------------------
Online Quarantined Devices across the fabric
-------------------------------------------------------------------------------------------------------------
 Port Index |   PID  |             PWWN          |  Locality      |Port Name      |Zone Aliases        |
-------------------------------------------------------------------------------------------------------------
   17       | 051100  | 30:10:00:05:33:ac:c6:13   |  LOCAL         |-              |SB_124082_P1vport0  |
   17       | 051101  | 30:10:01:05:33:ac:c6:13   |  LOCAL         |-              |SB_124082_P1vport1  |
-------------------------------------------------------------------------------------------------------------

switch:admin> sddquarantine --clear 2/1, 1/3
```

```
Initiated clearing port from quarantined state

switch:admin> sddquarantine --clear 1/3
Clear failed since port is still in latency state

switch:admin> sddquarantine --clear all
The clear action was not initiated for the following port(s). Try with individual ports
               3
Initiated quarantine action on other ports

switch:admin> sddquarantine --clear 1/3 -force
Initiated clearing port from quarantined state
```

Automatically Clearing Quarantined Ports

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

Perform the following steps:

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

Example of Clearing Quarantined Ports

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

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

Example of Removing Automatic Unquarantine Action

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

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

Delaying the Quarantining of Ports

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

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

1. Create a policy.
   ```
   switch:admin> mapspolicy --create test_policy1
   ```
2. Create a base FPI rule, but do **not** set an SDDQ action.
   ```
   switch:admin> mapsrule --create test_PORTS_IO_FRAME_LOSS -group ALL_PORTS
                     -monitor DEV_LATENCY_IMPACT -timebase none -op eq
                     -value IO_FRAME_LOSS -action raslog,snmp,email -policy test_policy1
   ```

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

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

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

Port Toggling

MAPS supports port toggling, which allows Fabric OS software to recover a port that is congested by a target device. While there are many reasons why the target device could be congested, 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 reenables it, allowing it to reset and recover from the issue. To enable recovering ports using port toggling, MAPS assumes a redundant path to all ports. It does not check to see if there is one, nor can it check if traffic to or from the port is switched over to a redundant path. MAPS also assumes that while the port is being toggled, the operational state of the port is not changed by any other mechanism such as an administrator disabling or moving the port or a port fencing operation. Port toggling cannot be used with automatic VC quarantining as this might result in the unpredictable behavior.

Port toggling is enabled in MAPS by including the toggle action in the rule and specifying a value from 2 through 3600 seconds as the length of time that the port is to be disabled.

The following example defines a rule that toggles a port offline for 180 seconds (3 minutes) when the number of CRC errors in a minute on the port is greater than 0.

```
switch:admin> mapsrule --config toggle_rule -group DB_PORTS -monitor DEV_LATENCY_IMPACT -timebase none -op eq
 -value IO_PERF_IMPACT -action TOGGLE -tt 180
```

When a port is toggled by a MAPS rule, TOGGLE appears as a notification action in the output of the `mapsconfig` and `mapsdb` commands. The following example displays a sample of the `mapsdb --show` command that illustrates this result.

```
switch:admin> mapsdb --show
1 Dashboard Information:
=======================
DB start time:                 Fri Aug 11 18:38:12 2021
Active policy:                 test_xy1
Configured Notifications:      RASLOG,FENCE,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL
Fenced Ports :                 None
Decommissioned Ports :         None
Fenced circuits :              38/0,38/1,38/2,38/3,38/4,38/5,38/6,38/7
Quarantined Ports :            None

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

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

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

```
switch:admin> switchshow
444  8 28   ------   cu   8G    No_Sync   FC
```

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

FPIN Action

Whenever congestion-related FPI rules are triggered, a Fabric Notification (specifically an FPIN having a congestion descriptor) is sent periodically to participating devices until MAPS clears the congestion criteria.

For more information about fabric notifications, refer to the *Brocade Fabric OS Administration Guide*.

Congestion Dashboard

Congestion dashboard displays the list of potentially congested ports that are determined by monitoring the following aspects:

- **Transient Queue Latency (TQL)** – TQL is the maximum time that a frame spent in the transmit queue. The following thresholds are used if the default FPI profile is enabled. Otherwise, custom FPI profile thresholds are used.
  - If the TQL is 80 or more milliseconds, it is considered as frame loss due to congestion (IO_FRAME_LOSS state).
  - If the TQL is 10 or more milliseconds, it is considered as a traffic performance impact due to congestion (IO_PERF_IMPACT state).
- **Credit Zero Statistics** – Credit zero statistics indicate the time that is taken between a frame ready to be transmitted and the port made buffer credits available for the transmission. The following are the allowed latency per time intervals in the order of checking done if the default FPI profile is used. If a port exceeds any of these latency values, then the port is considered to be the impacting traffic performance (IO_PERF_IMPACT state). You can configure different values in the custom FPI profiles.
  a. 700 milliseconds in 1 second.
  b. 2500 milliseconds in 5 seconds.
  c. 3000 milliseconds in 10 seconds.
- **Class 3 Discard Errors** – If a port drops frames, it is considered to be frame loss due to congestion (IO_FRAME_LOSS state).

Congested Port States

In the congestion dashboard, the ports are categorized into the following states based on the specified criteria:

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

  If you use the custom FPI profile, 50% of the performance impact threshold is considered as the medium congestion state.
- **Low** – A port is marked in a low congestion state if any or both of the following conditions are met:
  - If the TQL is 3 or more milliseconds but less than 5 milliseconds.
  - If the credit zero statistics indicate a latency of 50 or more milliseconds but less than 100 milliseconds in 1 second.

If you use the custom FPI profile, >30% of the performance impact threshold is considered as the low congestion state.

- **Info** – A port is marked in informative congestion state if any or both of the following conditions are met:
  – If the TQL is 1 or more milliseconds but less than 3 milliseconds.
  – If the credit zero statistics indicate a latency of 10 or more milliseconds but less than 50 milliseconds in 1 second.

  If you use the custom FPI profile, >10% of the performance impact threshold is considered as the informative congestion state.

Congestion State Calculation

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

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

Table 38: Congestion State Calculation

Duration	State	Threshold Range Based on Area Weight
Per Second	Frame Loss	5
	I/O Perf Impact	4
	Medium	3
	Low	2
	Info	1
Per Minute	Frame Loss	>240 and <=300
	I/O Perf Impact	>180 and <=240
	Medium	>120 and <=180
	Low	>60 and <=120
	Info	>0 and <=60
Per Hour	Frame Loss	>14400 and <=18000
	I/O Perf Impact	>10800 and <=14400
	Medium	>7200 and <=10800
	Low	>3600 and <=7200
	Info	>0 and <=3600

Clearing MAPS Dashboard Data

To delete the stored data from the MAPS dashboard, enter the `mapsdb --clear` command. This command displays only the data that is logged after you changed a switch (or a rule). The dashboard is also cleared on a reboot or an HA failover.

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

1. Connect to the switch and log on using an account with admin permissions.
2. Enter the `mapsdb --clear` command and specify the level of data (`all, summary, history, or congestion`) you want to remove from the display. You can use the `-force` option to forcefully proceed with CLI execution for the `all, summary, history, or congestion` options without prompting for user input.

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

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

   The `mapsdb --clear summary` command clears the dashboard summary data.

```
switch:admin> mapsdb --clear summary
WARNING: This command will clear summary data
Do you want to continue? (yes, y, no, n): [no]
```

   The following RASLog message shows the example of the dashboard summary data after clearing.

```
2023/02/24-16:56:05, [MAPS-1203], 47554, FID 128, WARNING, G610, Dashboard summary data has been
 cleared.
```

   The following example shows the dashboard output after clearing the dashboard summary data.

```
switch:admin> mapsdb --show

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

DB start time:                  Thu Feb 23 19:25:46 2023
Active policy:                  dflt_conservative_policy
Configured Notifications:       RASLOG,EMAIL,FENCE,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL,DECOM
Fenced Ports :                  None
Decommissioned Ports :          None
Fenced circuits :               None
Quarantined Ports :             None
Top Zoned PIDs <pid(it-flows)>: 0xbb3740(1) 0xbb1000(1) 0xbb3741(1) 0xbb8f00(1) 0xbb0000(1)

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

Current Switch Policy Status: HEALTHY


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

Category                |Today                      |Last 7 days               |
-------------------------------------------------------------------------------
Port Health             |No Errors                  |No Errors                 |
BE Port Health          |No Errors                  |No Errors                 |
Extension GE Port Health |No Errors                  |No Errors                 |
```

```
    Fru Health               |In operating range        |In operating range        |
    Security Violations      |In operating range        |No Errors                 |
    Fabric State Changes     |No Errors                 |No Errors                 |
    Switch Resource          |In operating range        |In operating range        |
    Extension Health         |No Errors                 |No Errors                 |
    Fabric Performance Impact|In operating range        |In operating range        |
    IO Health                |In operating range        |In operating range        |
    IO Latency               |In operating range        |In operating range        |

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

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

Displaying the Congestion Dashboard

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

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

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

DB start time: Thu Nov 11 03:23:26 2021
Time Window: 03:28 - 03:23
Total Credit-Stalled ports: 0
Total Oversubscribed ports: 1

2 Credit-Stall State Frequency Table:
=====================================

Port      |Current Minute State |Frame Loss   |Perf Impact   |Medium   |Low       |Info     |
-------------------------------------------------------------------------------------------

3 Oversubscription State Frequency Table:
=========================================

Port      |Current Minute State |Frequency    |
-------------------------------------------------
E-Port 42 |Oversubscribed       |2            |
```

The following example shows the congestion condition that occurred due to the credit-stalled devices.

```
switch:admin> mapsdb --show congestion -credit-stall

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

DB start time:                  Fri Feb 21 15:54:08 2021
Time Window:                    08:35 - 09:35
Total Credit-Stalled ports:     5
```

```
2 Credit-Stall State Frequency Table:
=====================================

Port        |Current Min State |Frame Loss   |Perf Impact  |Medium   |Low      |Info     |
-------------------------------------------------------------------------------------------
E-Port 3/9  |Perf Impact       |0            |60           |0        |0        |0        |
E-Port 8/32 |Perf Impact       |0            |60           |0        |0        |0        |
E-Port 3/31 |Medium St         |0            |0            |60       |0        |0        |
E-Port 7/2  |Perf Impact       |0            |60           |0        |0        |0        |
E-Port 7/0  |Perf Impact       |0            |33           |27       |0        |0        |
```

The following example shows the congestion dashboard at the specified hour using the `-hr` option.

```
mapsdb --show congestion -credit-stall -hr 06

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

DB start time:          Sun Jun  7 05:18:52 2021
Time Window:            06:00:00 - 07:00:00
Total Credit-Stalled ports:    1

2 Credit-Stall State Frequency Table:
=====================================

Port        |Current Minute State |Frame Loss   |Perf Impact  |Medium   |Low      |Info     |
-------------------------------------------------------------------------------------------
E-Port 3/49 |Medium St            |0            |4            |56       |0        |0        |
```

To display the congestion dashboard with frequency details, use the `-freq` option. The frequency table displays the frequency of the congestion for the past 10 hours. It is sorted based on the congestion frequency count for a port in a given hour. MAPS samples port state every second. This indicates the number of seconds in the particular hour bucket the port was in a congestion state. The first column represents the current hour. So if the current switch time is 06:30:00, then the first column in the following output will have 30-minutes frequency value (06:00:00 to 06:30:00). All the other columns represent complete one-hour data.

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

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

DB start time: Thu Nov 11 03:23:26 2021
Time Window: 03:28 - 03:23
Total Credit-Stalled ports: 0
Total Oversubscribed ports: 1

2 Credit-Stall State Frequency Table:
=====================================

Port      |Current Minute State |Frame Loss   |Perf Impact   |Medium   |Low       |Info     |
-----------------------------------------------------------------------------------------------

3 Oversubscription State Frequency Table:
```

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

Port      |Current Minute State |Frequency    |
-----------------------------------------------
E-Port 42 |Oversubscribed       |2            |
```

The following example shows the frequency of the credit-stall conditions. The output is edited to display clearly.

```
switch:admin> mapsdb --show congestion -credit-stall -freq

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

DB start time:                  Tue Apr 14 16:39:00 2021

2 Credit-Stall Frequency Table:
===============================

-------------------------------------------------------------------------------------------------------
05:00:00           |04:00:00           |03:00:00           |02:00:00           |01:00:00           |
-------------------------------------------------------------------------------------------------------
Total Ports: 4     |Total Ports: 4     |Total Ports: 4     |Total Ports: 4     |Total Ports: 4     |
-------------------------------------------------------------------------------------------------------
E-Port 3/9,(3156)  |E-Port 3/31,(3567) |E-Port 3/31,(3567) |E-Port 3/9,(3571)  |E-Port 3/31,(3567) |
E-Port 3/31,(3150) |E-Port 3/9,(3567)  |E-Port 3/9,(3557)  |E-Port 3/31,(3562) |E-Port 3/9,(3553)  |
F-Port 6/23,(36)   |E-Port 3/38,(39)   |E-Port 3/38,(41)   |F-Port 6/23,(36)   |F-Port 6/23,(37)   |
E-Port 3/38,(22)   |F-Port 6/23,(36)   |F-Port 6/23,(33)   |E-Port 3/38,(28)   |E-Port 3/38,(18)   |
```

FPI Profile Management

The Fabric Performance Impact (FPI) profile is the MAPS feature used to manage and customize the FPI congestion thresholds. By default, the FPI profile is enabled for E, F, and trunking ports with the static thresholds. If the static thresholds are suitable for your monitoring requirements and you can continue to use them. In addition to the default FPI profile, MAPS allows you to create two additional custom profiles with the average TXQ latency and CRED-ZERO values for IO_PERF_IMPACT and IO_FRAME_LOSS to suit your custom requirements. In Fabric OS v9.2.x, you can customize the transmit utilization percentage of OVERSUBSCRIPTION state.

You can use the custom FPI profiles under the following cases to suit your stricter monitoring requirements.

- Static thresholds applied to an environment might not be suitable for other environments. For example, an aggressive threshold in an SSD environment might be conservative in the non-volatile memory express (NVMe) environment.
- You might want to configure different thresholds depends on the port types, for example, custom thresholds for F_Ports and E_Ports.

> **NOTE**
> If you are upgrading from the Fabric OS v9.1.x version, the FPI monitoring works seamlessly by automatically using the default FPI profile with static thresholds. You can create two custom profiles that suit your monitoring requirements.

The following command shows the example of the default FPI profile, and you can use this to manage the FPI profile threshold configurations.

```
switch:admin> fpiprofile --show
-------------------------------------------------------------------------------------------------------
FPI Profile         |Profile Type |IO_PERF_IMPACT      |IO_FRAME_LOSS       |OVERSUBSCRIBED        |
```

```
-------------------------------------------------------------------------------------------------------
dflt_fpi_profile    |F & E Ports  |TXQL          = 10ms|TXQL        = 80ms  |TX          = 70%     |
                    |             |CRED-ZERO-1s  = 70% |                    |                      |
                    |             |CRED-ZERO-5s  = 50% |                    |                      |
                    |             |CRED-ZERO-10s = 30% |                    |                      |
-------------------------------------------------------------------------------------------------------
```

The following command creates a custom FPI profile with different thresholds for F_Port and E_Port. You can edit the default profile to create custom profiles.

```
fpiprofile --create Custom_FPI -perfimpact -txql 5 -cred-zero-1 60 -cred-zero-5 30 -cred-zero-10 20 -frameloss
 -txql 70 -oversubscription -tx 80
```

The following example shows the output after creating a custom policy named Custom_FPI.

```
switch:admin> fpiprofile --show
-------------------------------------------------------------------------------------------------------
FPI Profile         |Profile Type |IO_PERF_IMPACT      |IO_FRAME_LOSS       |OVERSUBSCRIBED        |
-------------------------------------------------------------------------------------------------------
dflt_fpi_profile    |F & E Ports  |TXQL          = 10ms|TXQL        = 80ms  |TX          = 70%     |
                    |             |CRED-ZERO-1s  = 70% |                    |                      |
                    |             |CRED-ZERO-5s  = 50% |                    |                      |
                    |             |CRED-ZERO-10s = 30% |                    |                      |
-------------------------------------------------------------------------------------------------------
Custom_FPI          |Not Active   |TXQL = 5ms          |TXQL = 70ms         |TX = 70%              |
                    |             |CRED-ZERO-1s = 60%  |                    |                      |
                    |             |CRED-ZERO-5s = 30%  |                    |                      |
                    |             |CRED-ZERO-10s = 20% |                    |                      |
-------------------------------------------------------------------------------------------------------
```

The following example shows the RASLog message that depicts that the custom FPI profile is created.

```
2023/02/24-02:36:51, [MAPS-1140], 300, SLOT 2 | FID 128, INFO, Gen7-8, FPI Profile Custom_FPI is created
```

The following command clones the existing FPI profile with the modified thresholds.

```
switch:admin> fpiprofile --clone custom_FPI -name MyCustom_FPI
```

The following example shows the RASLog message that depicts that the FPI profile is cloned.

```
2023/02/24-04:06:24, [MAPS-1142], 301, SLOT 2 | FID 128, INFO, Gen7-8, FPI Profile custom_FPI is cloned to
 MyCustom_FPI
```

FPI Profile Thresholds

The impact of the congestion on the fabric varies based on the severity of the latency and frame loss that occurred. The following summarizes congestion severity levels based on the credit latency, queue latency, and frame loss that are derived from the performance impact state (which are the minimum thresholds among the two FPI states). The medium, low, and info states are derived as 70%, 50%, and 10% of performance impact thresholds, respectively.

The following table defines the FPI thresholds and congestion severity based on various combinations of oversubscription, credit-stalled device, and lost credit metrics.

Table 39: FPI Profile Thresholds

FPI Monitoring States	Thresholds	Units	Minimum	Maximum	Customizable
Performance Impact	TXQL	Milliseconds	1	50	Yes
	CRED-ZERO in 1 second	Percentage	50%	100%	Yes
	CRED-ZERO in 5 seconds	Percentage	30%	100%	Yes
	CRED-ZERO in 10 seconds	Percentage	10%	100%	Yes
Frame Loss	C3TX Timeouts	Timeouts	0	NA	No
	TXQL	Milliseconds	50	200	Yes
Oversubscribed	TX Utilization	Percentage	70%	100%	Yes

Zoned Device Ratio Monitoring

MAPS allows monitoring of the zoned device ratio per port.

Devices can be zoned to other devices to allow communications. When MAPS finds that a port has more number of devices than the limit configured in the rule, then it alerts the administrator. Zone configuration can cause backpressure in the following situations:

1. If a device is zoned with a disproportionate number of devices.
2. If a port is allowed to communicate with a disproportionate number of ports.

MAPS uses the following rules to support this monitoring.

```
Rule name                   |Condition                       |Actions          |Policy
-------------------------------------------------------------------------------------------|
defALL_LOCAL_PIDSIT_FLOW_8  |ALL_LOCAL_PIDS(IT_FLOW/NONE>8) |RASLOG,SNMP,EMAIL|Aggressive  |
-------------------------------------------------------------------------------------------|
defALL_LOCAL_PIDSIT_FLOW_16 |ALL_LOCAL_PIDS(IT_FLOW/NONE>16)|RASLOG,SNMP,EMAIL|Moderate    |
-------------------------------------------------------------------------------------------|
defALL_LOCAL_PIDSIT_FLOW_32 |ALL_LOCAL_PIDS(IT_FLOW/NONE>32)|RASLOG,SNMP,EMAIL|Conservative|
-------------------------------------------------------------------------------------------|
```

Example of Dashboard Output for Zoned Device Ratio Monitoring

```
switch:admin> mapsdb --show

Dashboard Information:
======================
Ports with highest Zoned device ratio:            0x20000,0x200200
```

Examples of Using the Logical Group Supporting Zoned Device Ratio Monitoring

The ALL_LOCAL_PIDS logical group helps you to manage all the PIDs. The following are examples of using the ALL_LOCAL_PIDS logical group.

```
switch:admin> logicalgroup --show all_local_pids details
```

```
-------------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type              |Member Count |Members
-------------------------------------------------------------------------------------------------
ALL_LOCAL_PIDS                  |Yes        |Pid               |6480         |All Pids monitored


switch:admin> logicalgroup --show | grep ALL_LOCAL_PIDS
ALL_LOCAL_PIDS                  |Yes        |Pid               |6480         |All Pids monitored
```

NOTE
The ALL_LOCAL_PIDS logical group is not supported in FICON environments because they use flat zoning.

Considerations for Zoned Device Ratio Monitoring

The following information should be considered for zoned device ratio monitoring:

- MAPS also shows the top five PIDs in the dashboard output.
- QT is not supported by zoned device ratio monitoring; therefore, creating a rule that specifies QT fails.
- The monitoring of zoned device ratios starts an hour after a system reboot, provided the fabric has been formed. Duplicate alerts for the same PID might occur if the system comes up a few minutes before midnight. However, after the system has been up for more than a day, the rule is executed, and alerts are generated once per day.

Monitoring UCS Uplink Distribution to the Brocade SAN

MAPS monitors the distribution of the Cisco Unified Computing System (UCS) server vHBA logons on the physical FC uplinks from the Fabric Interconnect (FI) to the Brocade SAN switch.

Feature Overview

With multiple parallel FC uplinks from the UCS FI to the Brocade SAN switch, the FI evenly distributes the server vHBA logons (NPIV) among the available uplinks and keeps the links in a balanced state. If there is any failure on any of the uplinks, the server logons failover to the remaining available uplinks; but once the failed uplink is restored, the server logons do not failback. This results in an oversubscription of some of the uplinks and causes an imbalanced state.

Figure 5: UCS Rebalancing

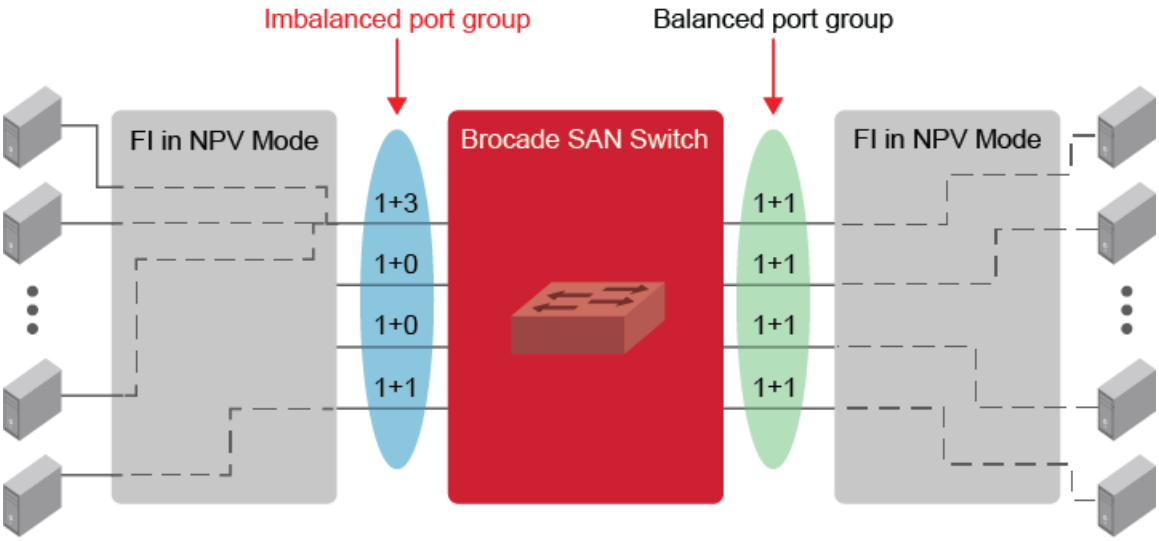

The MAPS Monitoring UCS Uplink Distribution feature automatically detects and group the parallel UCS FI uplinks to the Brocade switch and monitor the UCS server logons over these links for oversubscription conditions. On detecting an imbalance, MAPS generates alerts to notify the administrators. Administrators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

Formation of the Dynamic Port Group

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

Verifying the Port Group State

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

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

Detecting Imbalance

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

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

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

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

NOTE

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

Configuring UCS Uplink Distribution Monitoring

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

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

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

```
switch:admin> devicelogin --show
-------------------------------------------------------------------------------
Node WWN               |State      |Ports Count|Ports (Number of devices)      |
-------------------------------------------------------------------------------
```

```
20:07:00:de:fb:a2:87:01|BALANCED   |4          |7/39(9),7/36(9),7/37(9),7/38(9)|
```

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

```
switch:admin> mapsdb --show
1 Dashboard Information:
------------------------
DB start time:              Tue Sep 26 07:12:16 2021
Active policy:              dflt_conservative_policy
Configured Notifications:   SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL

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

Category(Violation Count)         | RepeatCount|Rule Name                                |Execution Time
         |Object      |Triggered Value(Units)|
-----------------------------------------------------------------------------------------------------------------
Fabric Performance Impact(2)      |1           |defALL_F_PORTSDEV_LOGIN_DIST_BALANCED   |09/26/21
 09:54:04      |F-Port 7/36 |BALANCED              |
                                  |1           |defALL_F_PORTSDEV_LOGIN_DIST_IMBALANCED |09/26/21
 09:44:04      |F-Port 7/36 |IMBALANCED              |
switch:admin> mapsdb --show
1 Dashboard Information:
------------------------
DB start time:              Tue Sep 26 07:12:16 2021
Active policy:              dflt_conservative_policy
Configured Notifications:   SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL

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

Category(Violation Count)         | RepeatCount|Rule Name                                |Execution Time
         |Object      |Triggered Value(Units)|
-----------------------------------------------------------------------------------------------------------------
Fabric Performance Impact(2)      |1           |defALL_F_PORTSDEV_LOGIN_DIST_BALANCED   |09/26/21
 09:54:04      |F-Port 7/36 |BALANCED          |
                                  |1           |defALL_F_PORTSDEV_LOGIN_DIST_IMBALANCED |09/26/21
 09:44:04      |F-Port 7/36 |IMBALANCED          |
```

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

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

switch:admin> errshow
[MAPS-1003], 1173, SLOT 1 | FID 1 | PORT 7/36, WARNING, X6-4_140_061, slot7 port36, F-Port 7/36,
 Condition=ALL_F_PORTS(DEV_LOGIN_DIST==IMBALANCED), Current Value:[DEV_LOGIN_DIST,IMBALANCED],
RuleName=defALL_F_PORTSDEV_LOGIN_DIST_IMBALANCED, Dashboard Category=Fabric Performance Impact.
```

```
[MAPS-1003], 1174, SLOT 1 | FID 1 | PORT 7/36, WARNING, X6-4_140_061, slot7 port36, F-Port 7/36,
 Condition=ALL_F_PORTS(DEV_LOGIN_DIST==BALANCED), Current Value:[DEV_LOGIN_DIST,BALANCED],
RuleName=defALL_F_PORTSDEV_LOGIN_DIST_BALANCED, Dashboard Category=Fabric Performance Impact.
```

3. Enable the rebalancing operation using one of the following methods.
   - Enable the MAPS auto-rebalancing operation by using `mapsconfig --actions RE_BALANCE` command.

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

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

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

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

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

Congestion Due to Lost Credits

This section describes the identification and causes of congestion conditions due to lost credits. The lost credit is independent of the MAPS FPI feature.

Lost Credits

Link errors, such as cyclic redundancy check (CRC) and invalid transmission word (ITW) errors, can affect frame integrity and can cause valid R_RDYs to stop being returned to the sending switch. This ultimately causes lost credits on device-to-switch links, inter-switch links (ISLs), and back-end switch ports. Lost credits are caused by the signal being degraded by the transmitter, the cable system, or the receiver, such that the transmitted data cannot be successfully recovered.

MAPS monitors and provides alerts for various link errors and statistics through the Port Health category to notify you about the following link errors occurring on switch ports:

- CRC (cyclic redundancy check)
- ITW (invalid transmission word) errors
- Other signal-related issues

Link errors that lead to lost credits are caused by faulty cables, faulty or improperly seated transceivers, and faulty connections. Connections between internal back-end links between ASICs can also suffer from physical errors.

Perform the following tasks to mitigate the effects of congestion due to lost credits:

- Maintain cables, transceivers, connections, and network infrastructure.
- Enable credit recovery.
- Launch ClearLink® (D_Port) Diagnostics.

For more information about congestion conditions, refer to the *Brocade Fabric Congestion Troubleshooting Guide*.

The most common link errors that occur with lost credits are CRC or ITW errors. Other link errors can include link reset, loss of sync, and loss of signal. Check for MAPS alerts warning that these errors have exceeded thresholds set in MAPS rules, such as the number of errors occurring per minute. MAPS alerts display through RASLog messages, email, SNMP, or whatever you set as action through MAPS for generating alerts.

You can use the 3.1 Summary Report and 3.2 Rules Affecting Health sections in the MAPS dashboard to display current port health and counts of triggered MAPS rules on all switch ports. To display the MAPS dashboard, use the `mapsdb --show` command.

The following example shows the link errors that cause the lost credits.

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

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


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

Category(Violation Count)|RepeatCount|Rule Name                   |Execution Time   |Object           |
Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------------------------------
Port Health(1)           |1          |defALL_PORTSLOSS_SIGNAL_0   |09/29/19 23:29:38|U-Port 3/48      |1 LOS
               |
BE Port Health(1)        |1          |defALL_BE_PORTSCRC_5M_10    |09/30/19 02:36:02|BE Port 11/88    |26 CRCs
               |
Fru Health(31)           |4          |defALL_PORTSSFP_STATE_IN    |09/30/19 00:08:20|U-Port 10/45     |IN
               |
                         |           |                            |                 |U-Port 10/47     |IN
               |
                         |           |                            |                 |U-Port 10/46     |IN
               |
                         |           |                            |                 |U-Port 10/44     |IN
               |
                         |4          |defALL_PORTSSFP_STATE_OUT   |09/29/19 23:29:34|U-Port 3/49      |OUT
               |
           ----------------------------Output truncated--------------------------------

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

Stats(Units)        Current     09/29/19    --/--/--    --/--/--    --/--/--    --/--/--    --/--/--
----------------------------------------------------------------------------------------------------

CRC(CRCs)           -           -           -           -           -           -           -
ITW(ITWs)           -           -           -           -           -           -           -
LOSS_SYNC(SyncLoss) -           -           -           -           -           -           -
```

```
LF(LFs)              -           -           -           -           -           -           -
LOSS_SIGNAL(LOS)     -           3/48(1)     -           -           -           -           -
                     -           3/49(1)     -           -           -           -           -
                     -           3/50(1)     -           -           -           -           -
                     -           3/51(1)     -           -           -           -           -
PE(Errors)           -           -           -           -           -           -           -
STATE_CHG            -           3/49(2)     -           -           -           -           -
                     -           3/50(2)     -           -           -           -           -
                     -           3/51(2)     -           -           -           -           -
                     -           3/31(1)     -           -           -           -           -
                     -           3/48(1)     -           -           -           -           -
LR(LRs)              -           -           -           -           -           -           -
C3TXTO(Timeouts)     -           -           -           -           -           -           -
RX(%)                -           -           -           -           -           -           -
TX(%)                -           -           -           -           -           -           -
UTIL(%)              -           -           -           -           -           -           -
BN_SECS(Seconds)     -           -           -           -           -           -           -

5 History Data for Backend ports:
=================================
Stats(Units)         Current     09/29/19    --/--/--    --/--/--    --/--/--    --/--/--    --/--/--
-------------------------------------------------------------------------------------------------------
CRC(CRCs)            11/88(26)   -           -           -           -           -           -
ITW(ITWs)            -           -           -           -           -           -           -
LR(LRs)              -           7/23(2)     -           -           -           -           -
                     -           7/25(2)     -           -           -           -           -
                     -           7/28(2)     -           -           -           -           -
                     -           7/87(2)     -           -           -           -           -
                     -           7/91(2)     -           -           -           -           -
BAD_OS(Errors)       -           -           -           -           -           -           -
FRM_LONG(Errors)     -           -           -           -           -           -           -
FRM_TRUNC(Errors)    -           -           -           -           -           -           -
```

Scalability Limit Monitoring

MAPS monitors changes in fabric-level monitoring systems. These systems have all scalability limits, which MAPS can monitor and send alerts using RASLog entries, SNMP messages, or emails.

The monitoring results are captured in the MAPS dashboard under the Fabric State Changes category. Although there are default rules that monitor these values, MAPS allows you to define new rules with different thresholds and actions. MAPS monitors the following scalability limits:

- The number of logged-in device connections in a pure Layer 2 fabric.
- The size of the zone configuration resource that is used.
- The number of Fiber Channel Router configurations.
- The number of Logical SAN (LSAN) device connections (includes both edge fabric and backbone fabric device connections).

  NOTE
  MAPS monitors the device count per FCR switch.

When a rule is triggered, the corresponding RASLog appears in the summary section of the dashboard. The following example shows two rules (LSAN_DEVCNT_PER and L2_DEVCNT_PER) that are triggered. The column headings in the example are edited to allow the example to display clearly.

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

3.2 Rules Affecting Health:
===========================
Category(Rule Count)   |RptCnt|Rule Name      |Execution Time  |Object   |Triggered Value (Units)|
-------------------------------------------------------------------------------------------------
Fabric State Changes(2)|1     |LSAN_DEVCNT_PER|03/21/21 00:30:6|D_Port 23|12 %                   |
                       |1     |L2_DEVCNT_PER  |03/21/21 01:04:6|D_Port 23|12 %                   |
```

Layer 2 Fabric Device Connection Monitoring

A pure Layer 2 fabric is a collection of fibre switches (FC) and devices that do not 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 8280 devices logged in is using 92 percent of the maximum limit of 9000 devices for a Layer 2 fabric. So if the user configured a rule to trigger an alert at 90 percent or greater, then MAPS triggers the action configured for that rule and sends the data to the dashboard.

LSAN Device Connection Monitoring in a MetaSAN

The collection of all devices, switches, edge and fabrics, LSANs, and routers that make up a physically connected but logically partitioned storage network are called a metaSAN. Using MAPS, the total number of LSAN device connections (including the total number of devicesfrom all edge fabrics) in a metaSAN can be monitored for a scalability limit.

MAPS rules for monitoring imported LSAN device connections in a metaSAN can be configured only on switches that are a part of the 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 the fabric contains four switches in the fabric and four switches each in four edge fabrics, the total number of LSAN devices in this metaSAN (including imported devices from all edge fabrics) is 1200. Given a maximum of 10000 devices, this is 12 percent. If you configured a rule to trigger at 10 percent or greater, then MAPS triggers the action configured for the rule, but only on those switches that are part of the fabric and caches the data in the dashboard.

Fabric Fibre Channel Router Count Monitoring

In a fabric, there can be a maximum number of 16 fibre channel routers (FCRs). MAPS rules can be configured to monitor the number of FC routers in the fabric as an absolute value. If the number of FC routers reaches the configured threshold; MAPS triggers the action configured for the rule and caches the data in the dashboard. See Default Rules for Scalability Limit Monitoring for these values.

The following example shows a typical RASLog entry for exceeding the threshold for the number of FCRs in the fabric.

```
2021/05/27-17:02:00, [MAPS-1003], 14816, SLOT 4 | FID 20, WARNING, switch_20, Switch,
```

```
Condition=SWITCH(BB_FCR_CNT>16), Current Value:[BB_FCR_CNT,17], RuleName= defSWITCHBB_FCR_CNT_MAX,
Dashboard Category=Fabric State Changes.
```

The following example shows a typical MAPS dashboard entry for exceeding the threshold for the number of FCRs in the fabric.

```
3.1 Summary Report:
===================
Category                 |Today                     |Last 7 days         |
-------------------------------------------------------------------------
Port Health              |No Errors                 |No Errors           |
BE Port Health           |No Errors                 |No Errors           |
Fru Health               |In operating range        |In operating range  |
Security Violations      |No Errors                 |No Errors           |
Fabric State Changes     |Out of operating range    |No Errors           |
Switch Resource          |In operating range        |In operating range  |
Traffic Performance      |In operating range        |In operating range  |
Extension Health         |No Errors                 |No Errors           |
Fabric Performance Impact|In operating range        |In operating range  |
---------------------------Output truncated----------------------------------
3.2 Rules Affecting Health:
===========================
Category(RuleCount) |RptCount|Rule Name              |Execution Time   |Object|Triggered Value(Units)|
-------------------------------------------------------------------------------------------------------
Fabric State Changes|1       |defSWITCHBB_FCR_CNT_MAX|05/27/14 17:02:00|Switch|17                    |
(1)                 |        |                       |                 |      |                      |
```

Zone Configuration Size Monitoring

MAPS monitors the zone configuration database size across the fabric in terms of percentage of the maximum support zone size. The maximum size is determined by the Fabric OS version and system types as follows:

- The maximum zone configuration database size for the systems running Fabric OS v9.2.x is 16 MB per fabric that includes 4 MB per virtual fabric (VF).
- The maximum zone configuration database size for the chassis systems running pre-Fabric OS v9.2.x is 2 MB.
- The maximum zone configuration database size for the fixed-port systems running pre-Fabric OS v9.2.x is 1 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 that is configured for the rule and caches the data in the dashboard. See Default Rules for Scalability Limit Monitoring for these limit values.

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 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 has scalability limits themselves, and these cannot be monitored using MAPS.
- The LSAN-imported device metric is only monitored in switches that are a part of a backbone fabric.

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 conditions for each rule.

Table 40: Scalability Monitoring Default Rules

Policy Name	Rule Name	Rule Condition	Rule Action
dflt_conservative_policy	defSWITCHL2_DEVCNT _PER_90 defSWITCHLSAN_DEVCNT _PER_90 defSWITCHZONE_ CFGSZ_PER_90 defSWITCHBB_FCR_ CNT_MAX	L2_DEVCNT_PER greater than 90 LSAN_DEVCNT_PER greater than 90 ZONE_CFGSZ_PER greater than 90 BB_FCR_CNT greater than 16	RASLOG, SNMP, EMAIL, FMS
dflt_moderate_policy	defSWITCHL2_DEVCNT_ PER_75 defSWITCHLSAN_DEVCNT_ PER_75 defSWITCHZONE_CFGSZ_ PER_80 defSWITCHBB_FCR _CNT_MAX	L2_DEVCNT_PER greater than 75 LSAN_DEVCNT_PER greater than 75 ZONE_CFGSZ_PER greater than 80 BB_FCR_CNT greater than 16	RASLOG, SNMP, EMAIL, FMS
dflt_aggressive_policy	defSWITCHL2_DEVCNT _PER_60 defSWITCHLSAN_DEVCNT _PER_60 defSWITCHZONE_CFGSZ _PER_70 defSWITCHBB_FCR _CNT_MAX	L2_DEVCNT_PER greater than 60 LSAN_DEVCNT_PER greater than 60 ZONE_CFGSZ_PER greater than 70 BB_FCR_CNT greater than 16	RASLOG, SNMP, EMAIL, FMS

For more information about default rules, see Default Monitoring Rules.

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.

```
switch: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.

```
switch: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 a 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 the 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 the ZONE_CFGSZ_PER policy, the default time base is none, and the system performs a daily check.

Switch-Level and Chassis-Level Features

This section describes the configuration specific to logical switches and chassis systems.

Most of the MAPS configurations are specific to logical switches (LSs), and few are specific to chassis. We can categorize the MAPS functions into the following two categories:

- Configuration
- Monitoring

Configuration

All the MAPS configurations are LS-specific except the RASLog mode and notification mode is a chassis-wide configuration. The LS-specific configurations include the structural elements of MAPS, such as policies, rules, groups, and actions. However, when you set the RASLog mode or notification mode, it applies to all the logical switches in the fabric.

For example, when you enter the `mapsconfig --raslog custom` command, the custom RASLog mode is set on all the logical switches in the fabric. The following example shows that the RASLog alert depicts the RASLog mode set to custom.

```
2023/02/16-11:09:12, [MAPS-1208], 851, SLOT 2 | FID 128, INFO, Gen7-8, Raslog mode is set to Custom.
```

Monitoring

Most of the MAPS monitoring features are LS-specific; however, some monitoring features are specific to the chassis. For example, BE Port Health monitoring. The MAPS monitoring systems indicate whether the monitoring system is specific to an LS or a chassis. For more information about the monitoring systems specific to an LS or a chassis, see Monitoring Systems Support Matrix.

The Default-Switch flag in the monitoring system indicates that the monitoring is chassis-wide and being monitored only in the default switch. If the monitoring system already contains the default rules, they are marked invalid in non-default logical switches. For more information about the constraints associated with each rule parameter, see Monitoring Systems Support Matrix.

The following example shows the Default-Switch flag that indicates the monitoring is chassis-wide and monitored in the default switch.

Table 41: Default-Switch Flags

MS	Name	Flags	Type	TimeBase	Unit	DB Category	Actions	Min QT	Threshold	RASLog ID
57	IT_COUNT	Lic, N-ROR, All, Default-Switch	Chassis	Y, N, N, N, N, N, N, N	%	Switch Resource	R, SN, E, F	60	float: 0.000000-100.000000	3016

If the monitoring system does not have the Default-Switch flag, it depicts that it is monitored in all the logical switches. For example, the CRC monitoring system provides alerts in all the logical switches.

Pausing and Restarting MAPS

This section describes the MAPS support for pausing and restarting MAPS monitoring.

MAPS provides early fault detection, fault isolation, and performance measurements. The MAPS-enabled switch can constantly monitor itself for potential faults. When actions are configured using the `mapsconfig` command, real-time alerts will be sent automatically if any problems are detected. When you want to perform the maintenance operations such as device or server upgrades, MAPS allows you to pause and restart the MAPS monitoring.

Pausing and restarting MAPS monitoring is supported on the following logical group types:

- Ports
- FCIP circuits
- SFPs
- Switch

Pausing MAPS Monitoring

To pause the monitoring an element in MAPS, enter the following command with the group type and members where you want to pause the monitoring.

```
mapsconfig --config pause -type <memberType> -members <member-string>
```

The following example command pauses monitoring for the port group type and all the members of the group.

```
switch:admin> mapsconfig --config pause -type port -members all
```

You must specify both the group type and the member information when you use the `–members all` keyword in the command to pause monitoring the ports, FCIP circuits, and SFP group types. The `–members all` keyword is not required for the switch. You can specify multiple members by separating them with a comma for individual members or a hyphen for a range of members.

You can validate the group types where you paused monitoring using the `mapsconfig –show` command as shown in the following example.

```
switch:admin> mapsconfig --show


---------------------Output truncated---------------------------
Paused members :
================
PORT : all
CIRCUIT : all
SFP : all
SWITCH : paused

Note: "all/paused" indicates complete group is paused with monitoring
```

MAPs actively monitors the paused objIfects. If any paused objects are detected, an alert is sent in a day. The following example shows the output of the `mapsconfig --show` command with paused individual group members.

```
2023/02/01-08:28:00, [MAPS-1211], 9920, FID 128, INFO, sw0, Monitoring on members 3,9-12 of type PORT, SFP is
 paused.
```

The following command pauses monitoring for a specific group type and specific group members:

```
mapsconfig --config pause -type <memberType> -members <member-string>
```

The following example shows that the monitoring is paused for port slot 4 and member 1:

```
switch:admin> mapsconfig --config pause -type port -members 4/1
```

When a port is moved from a logical switch to another with a paused group monitoring, ensure that the port to be moved is in the paused state before moving.

You can use the `mapsconfig --config pause -type all` command to pause monitoring for all the supported group types.

The following output shows an example of the RASLog message with all the group types that paused monitoring.

```
2023/02/07-10:54:06, [MAPS-1213], 45579, FID 128, INFO, G610, Monitoring is paused for all the supported
 group-types
```

Restarting MAPS Monitoring

You can restart the monitoring of the paused group types and members at any time. To restart monitoring for a paused port or another element in MAPS, enter the following command with both the element type and the group.

```
mapsConfig --config restart -type <memberType> -members <member-string>
```

Use the following command to restart monitoring for the specific group type and all the members of the group.

```
switch:admin> mapsconfig --config restart -type port -members all
```

The following output shows an example of the RASLog message with the restarted group types and members.

```
2023/02/07-10:56:36, [MAPS-1131], 45581, FID 128, INFO, G610,  Monitoring on members 4 of type PORT is paused.
2021/02/07-11:05:07, [MAPS-1132], 45584, FID 128, INFO, G610,  Monitoring on members 4 of type PORT is
 restarted.
```

The following output shows an example of all the port group types are paused for monitoring.

```
switch:admin> mapsconfig --show
Configured Notifications:       RASLOG,SNMP,EMAIL,SW_CRITICAL,SW_MARGINAL,SFP_MARGINAL
Mail Recipient:                 Not Configured
Mail From Address:              Not Configured
Raslog Mode:                    Default
Decom Action Config:            With Disable
Global Quiet Time:              2 Hours
Notification Behavior:          Adaptive
Paused members :
================
PORT : all
CIRCUIT :
SFP :
SWITCH :

Note: "all/paused" indicates complete group is paused with monitoring
```

You cannot restart monitoring for if a specific monitoring has already been stopped monitoring. The following output shows an example of the message that is shown when you restart a specific group type with the associated group is already stopped.

```
switch:admin> mapsconfig --config restart -type port -members 4/1 Monitoring is stopped for specified group.
 Individual Members cannot be restarted with monitoring.
```

On every restart, the previous MAPS data is deleted, and monitoring starts with the fresh data.

MAPS Support Matrix and Default Rules

This section provides information on various MAPS thresholds, the monitoring systems support matrix, and the list of default monitoring rules.

MAPS Default Groups Support Matrix

The following table lists the MAPS default groups, group types, supported Brocade platforms, and whether the MAPS default group is customizable and supported in AG mode.

Table 42: MAPS Default Groups Support Matrix

Default Group Name	Type	Customizable	Supported Platforms	Access Gateway Support
ALL_100M_16GSWL_QSFP	SFP	No	G620, X6-4, X6-8, G630, X7-4, X7-8, G720, G730, MXG610s, G648G620, X6-4, X6-8, G630, X7-4, X7-8, G720, G730, MXG610s, G648	True
ALL_10GLWL_SFP	SFP	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_10GSWL_SFP	SFP	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_16GLWL_SFP	SFP	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_16GSWL_SFP	SFP	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_25Km_16GLWL_SFP	SFP	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_25Km_32GELWL_JBA_SFP	SFP	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_25Km_32GELWL_SFP	SFP	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_25Km_64GELWL_SFP	SFP	No	X7-4, X7-8, G720, G730	True
ALL_2K_QSFP	SFP	No	X6-4, X6-8, X7-4, X7-8	False
ALL_2Km_32GLWL_QSFP	SFP	No	X6-4, X6-8, X7-4, X7-8	False
ALL_2Km_GEN7LWL_QSFP	SFP	No	X7-4, X7-8	False
ALL_32GLWL_SFP	SFP	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_32GLWL_JDB_SFP	SFP	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_32GSWL_QSFP	SFP	No	G620, X6-4, X6-8, G630, X7-4, X7-8, G720, G730, MXG610s, G648	True
ALL_32GSWL_SFP	SFP	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_40Km_16GELWL_SFP	SFP	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True

Default Group Name	Type	Customizable	Supported Platforms	Access Gateway Support
ALL_40Km_32GELWL_SFP	SFP	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_64GLWL_SFP	SFP	No	X7-4, X7-8, G720, G730	True
ALL_64GSWL_SFP	SFP	No	X7-4, X7-8, G720, G730	True
ALL_64GSWL_SFP_DD	SFP	No	G720, G730	True
ALL_80Km_8GELWL_SFP	SFP	No	G620, X6-4, X6-8, X7-4, X7-8	True
ALL_ASICS	ASIC	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	False
ALL_BE_PORTS	BE Ports	No	X6-4, X6-8, G630, X7-4, X7-8, G730	False
ALL_CERTS	Certificate	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_CIRCUITS	Circuit	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_CIRCUIT_F_QOS	Circuit QoS	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_CIRCUIT_HIGH_QOS	Circuit QoS	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_CIRCUIT_IP_HIGH_QOS	Circuit QoS	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_CIRCUIT_IP_LOW_QOS	Circuit QoS	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_CIRCUIT_IP_MED_QOS	Circuit QoS	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_CIRCUIT_LOW_QOS	Circuit QoS	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_CIRCUIT_MED_QOS	Circuit QoS	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_CORE_BLADES	Blade	No	X6-4, X6-8, X7-4, X7-8	False
ALL_DP	DP	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_D_PORTS	Port	Yes	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_ETH_PORTS	Eth Port	No	X6-4, X6-8, G630, X7-4, X7-8	False
ALL_EXT_GE_PORTS	GE Port	No	X6-4, X6-8, G630, X7-4, X7-8	False
ALL_E_PORTS	Port	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	False
ALL_FAN	Fan	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_FCOE_100G_SR4_QSFP	SFP	No	X6-4, X6-8, X7-4, X7-8	False
ALL_FCOE_40G_QSFP	SFP	No	X6-4, X6-8, X7-4, X7-8	False
ALL_FLASH	Flash	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_F_PORTS	Port	Yes	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_GEN7SWL_QSFP	SFP	No	X7-4, X7-8	False
ALL_HOST_PORTS	Port	Yes	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	False
ALL_LOCAL_PIDS	PID	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	False
ALL_N_PORTS	Port	Yes	G620, G720, G730, G610, MXG610s, G648	True

Default Group Name	Type	Customizable	Supported Platforms	Access Gateway Support
ALL_OTHER_F_PORTS	Port	Yes	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_OTHER_SFP	SFP	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_PORTS	Port	Yes	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_PS	Power Supply	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_QUARANTINED_PORTS	Port	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_SFP	SFP	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_SLOTS	Blade	No	X6-4, X6-8, X7-4, X7-8	False
ALL_SW_BLADES	Blade	No	X6-4, X6-8, X7-4, X7-8	False
ALL_TARGET_PORTS	Port	Yes	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	False
ALL_TS	Temperature sensor	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
ALL_TUNNELS	Tunnel	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_TUNNEL_F_QOS	Tunnel QoS	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_TUNNEL_HIGH_QOS	Tunnel QoS	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_TUNNEL_LOW_QOS	Tunnel QoS	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_TUNNEL_MED_QOS	Tunnel QoS	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_TUNNEL_IP_HIGH_QOS	Tunnel QoS	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_TUNNEL_IP_LOW_QOS	Tunnel QoS	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_TUNNEL_IP_MED_QOS	Tunnel QoS	No	7810, X6-4, X6-8, X7-4, X7-8	False
ALL_WWN	WWN	No	7810, X6-4, X6-8, X7-4, X7-8	False
CHASSIS	—	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
NON_E_F_PORTS	Port	Yes	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
SWITCH	—	No	7810, G620, X6-4, X6-8, G630, X7-4, X7-8, G610, G720, G730, MXG610s, G648	True
sys_flow_monitor	Flow	No	G620, X6-4, X6-8, G630, X7-4, X7-8, G720, G730, MXG610s, G648	False
sys_flow_monitor_nvme	Flow	No	X7-4, X7-8, G720, G730	False
sys_flow_monitor_scsi	Flow	No	X7-4, X7-8, G720, G730, G620, X6-8, X6-4	False

MAPS Action Support Matrix

The following table describes the MAPS actions and their support on Brocade platforms, Access Gateways, and port types, and whether a Fabric Vision license is required.

Actions	Gen 6 Support	Gen 7 Support	Access Gateway Support	Port Type Support	Fabric Vision License Required
DECOM	Yes	Yes	Yes	All	Yes
EMAIL	Yes	Yes	Yes	All	No
FENCE	Yes	Yes	Yes	All	Yes
FMS	Yes (only in the FICON setup)	Yes (only in the FICON setup)	No	All	No
FPIN	Yes	Yes	Yes	F_Port	No
RASLog	Yes	Yes	Yes	All	No
RE_BALANCE	Yes	Yes	No	F_Port	Yes
SDDQ	Yes	Yes	Yes	F_Port	Yes
SFP_MARGINAL	Yes	Yes	Yes	All	Yes
SNMP	Yes	Yes	Yes	All	No
SW_CRITICAL	Yes	Yes	Yes	N/A	No
SW_MARGINAL	Yes	Yes	Yes	N/A	No
Toggle	Yes	Yes	Yes	F_Port	Yes
UNQUAR	Yes	Yes	Yes	F_Port	Yes
VTAP	Yes	No	No	AE	Yes

Monitoring Systems Support Matrix

The Monitoring Systems Support Matrix (Primary) and Monitoring Systems Support Matrix (Secondary) tables contain the list of MAPS monitoring systems, the monitoring threshold policy values, the actions triggered, and so on. You can use the following sections to understand how the tabular data can be interpreted.

The flag values in the following table are denoted as given below:

- **DS** – Applicable only in default switch
- **NAG** – Not supported by Access Gateway
- **NR** – Rule on Rule is not supported
- **NQ** – Quiet Time is not supported
- **R** – Read only (Create, Edit, Delete or Customize operations are not supported)
- **L** – Monitoring system is active if the switch has the Fabric Vision license
- **states** – Monitoring system monitors the state. For example, IN, OUT, FAULTY, and OFF state values that are applicable to the BLADE_STATE monitoring system. For more information, refer to the Monitoring Systems and State-Based Threshold Values table.
- **Blank** – None of the flags are applicable

The action values in the following table are denoted as given below:

- **R** – RASLog
- **SN** – SNMP
- **E** – EMAIL
- **FN** – FENCE
- **D** – DECOM
- **F** – FMS
- **FPN** – FP
- **S_C** – SW_CRITICAL
- **S_M** – SW_MARGINAL
- **SF_M** – SFP_MARGINAL
- **T** – TOGGLE
- **SD** – SDDQ
- **U** – UNQUAR
- **U_V** – UNINSTALL_VTAP
- **R_B** – RE_BALANCE

The threshold values in the following table are denoted as given below:

- **max_u** – Data with the format and value of unsigned integer 0-999999999
- **max_f** – Data with the format and value in the range between 0.0 to 100.0
- **max_u64** – Data with the format and value in the range between 0.0 to 18446744073709551615

The system values in the following table are denoted as given below:

- **ALL** – Supported in all systems
- **M_ASICs** – Supported in multiple ASIC systems including the Brocade G630, G730, X6-4, X6-8, X7-4, and X7-8 platforms.
- **CHASSIS** – Supported only in the Brocade X6-4, X6-8, X7-4, and X7-8 Chassis systems.
- **LEG_THER** – Legacy thermal policy that is supported in all the Brocade Gen 6 systems except Brocade G620 (Switch Type 183) and G630 (Switch Type 184).
- **NEW_THER** – New thermal policy that is supported in all the Brocade Chassis systems including Brocade G620 (Switch Type 183) and G630 (Switch Type 184).
- **EXTN** – Supported in the Brocade Chassis systems and Brocade 7810 Switch.
- **TRUFOS** – All systems that support Brocade TruFOS Certificates. From Fabric OS v9.2.0, Brocade TruFOS Certificate monitoring is supported on all Brocade Chassis systems and switches excluding Brocade embedded platforms.
- **FCOE** – Supported in only Brocade Chassis systems.
- **G7_FLOW** – Supported in all the Brocade Gen 7 systems such as G720 and G730 Switches and X7-4 and X7-8 Chassis systems.
- **G6_FLOW** – Supported in all the Brocade Gen 6 Switches and Chassis systems.
- **ENCR** – Supported by the Brocade G620 Switches and X6-4 and X6-8 Chassis systems.

The following table describes the monitoring systems that have states as threshold values:

Table 43: State-Based Monitoring Systems and Their Threshold Values

Monitoring System	Group Type	Threshold Values
DEV_LATENCY_IMPACT	Port	IO_LATENCY_CLEAR, IO_PERF_IMPACT, IO_FRAME_LOSS
TEMP	Temperature sensor	IN_RANGE, OUT_OF_RANGE
PS_STATE	Power Supply	OUT, ON, FAULTY

Monitoring System	Group Type	Threshold Values
FAN_STATE	Fan	OUT, ON, FAULTY
SFP_STATE	Port	OUT, IN, FAULTY
BLADE_STATE	Blade	OUT, ON, OFF, FAULTY
WWN	WWN	OUT, ON, FAULTY
FAN_AIRFLOW_MISMATCH	Chassis	T
ETH_MGMT_PORT_STATE	Chassis	DOWN, UP
BE_LATENCY_IMPACT	BE Port	IO_LATENCY_CLEAR, IO_PERF_IMPACT, IO_FRAME_LOSS
HA_SYNC	Chassis	0-1
SFP_STATE	Eth Port	OUT, IN, FAULTY
DEV_LOGIN_DIST	Port	BALANCED, IMBALANCED, BALANCE_FAILED
SYSTEM_TEMP	Chassis	HEALTHY, MARG_OUT_OF_RANGE, CRIT_OUT_OF_RANGE
PORT_BANDWIDTH	Port	OVERSUBSCRIPTION_CLEAR, OVERSUBSCRIBED
TRUFOS_CERT_INSTALLED	Chassis	F
TRUFOS_CERT_EXPIRED	Chassis	T

Monitoring Systems Support Matrix (Primary)

The following table provides information that is required to clone an existing default policy and add to your custom rule. When you create a custom rule, you can use this information to account for the constraints associated with each parameter of that rule.

Type	Monitoring System	Flags	TimeBase	Actions	Thresholds	Unit	System
ASIC	VTAP_IOPS	DS, NR, L	sec	E, F, R, SN, U_V	0 - max_u64	IOPS	ALL
BE Port	BAD_OS	DS, R, L	min, hour, day	E, F, R, SN	0 - max_u	Errors	M_ASICs
BE Port	BE_LATENCY_IMPACT	DS, L	none	E, F, R, SN	*states	N/A	M_ASICs
BE Port	CRC	DS, R, L	min, hour, day	E, F, R, SN	0 - max_u	CRCs	M_ASICs
BE Port	FRM_LONG	DS, R, L	min, hour, day	E, F, R, SN	0 - max_u	Errors	M_ASICs
BE Port	FRM_TRUNC	DS, R, L	min, hour, day	E, F, R, SN	0 - max_u	Errors	M_ASICs
BE Port	ITW	DS, R, L	min, hour, day	E, F, R, SN	0 - max_u	ITWs	M_ASICs
BE Port	LR	DS, R, L	min, hour, day	E, F, R, SN	0 - max_u	LRs	M_ASICs
Blade	BLADE_STATE	NQ	none	E, F, R, SN	*states	N/A	CHASSIS
Certificate	DAYS_TO_EXPIRE	DS, NR, L	none	E, F, R, SN	0 - 99999	days	ALL
Chassis	AMP_RX_IOPS	DS, NAG, NR, R, L	sec	D, E, FN, F, R, SN	0 - max_u	IOPS	ALL

Type	Monitoring System	Flags	TimeBase	Actions	Thresholds	Unit	System
Chassis	ASC_UPLOAD_FAILURE	NR	week	E, F, R, SN	0 - max_u	N/A	ALL
Chassis	BAD_FAN	NR	none	E, F, R, SN, S_C, S_M	0 - max_u	N/A	ALL
Chassis	BAD_PWR	NR	none	E, F, R, SN, S_C, S_M	0 - max_u	N/A	ALL
Chassis	BAD_TEMP	NR	none	E, F, R, SN, S_C, S_M	0 - max_u	N/A	LEG_THER
Chassis	CPU		none	E, F, R, SN	0 - max_f	%	ALL
Chassis	DOWN_CORE	NR	none	E, F, R, SN, S_C, S_M	0 - max_u	N/A	CHASSIS
Chassis	ETH_MGMT_PORT_STATE	DS, NR, NQ	none	E, F, R, SN	*states	N/A	ALL
Chassis	EXPIRED_CERTS	DS, NR	none	E, F, R, SN, S_C, S_M	0 - max_u	certs	ALL
Chassis	FAN_AIRFLOW_MISMATCH	NR	none	E, F, R, SN, S_C, S_M	*states	N/A	ALL
Chassis	FAULTY_BLADE	NR	none	E, F, R, SN, S_C, S_M	0 - max_u	N/A	CHASSIS
Chassis	FLASH_USAGE	NR	none	E, F, R, SN, S_C, S_M	0 - max_u	%	ALL
Chassis	HA_SYNC	NR	none	E, F, R, SN, S_C, S_M	0 - max_u	N/A	CHASSIS
Chassis	ITL_RES_USAGE	DS, NAG, NR, L	none	E, F, R, SN	0 - max_f	%	G7_FLOW
Chassis	IT_RES_USAGE	DS, NAG, NR, L	none	E, F, R, SN	0 - max_f	%	G7_FLOW
Chassis	MEMORY_USAGE		none	E, F, R, SN	0 - max_f	%	ALL
Chassis	MEMORY_USAGE_STATE		none	E, F, R, SN	*states	N/A	ALL
Chassis	SYSTEM_TEMP	NR	none	E, F, R, SN, S_C, S_M	*states	N/A	NEW_THER
Chassis	TRUFOS_CERT_DAYS_TO_EXPIRE	DS, NR	none	E, F, R, SN	0 - max_u	days	ALL
Chassis	TRUFOS_CERT_EXPIRED	DS, NR	none	E, F, R, SN	*states	N/A	ALL
Chassis	TRUFOS_CERT_INSTALLED	DS, NR	none	E, F, R, SN	*states	N/A	ALL
Chassis	WWN_DOWN	NR	none	E, F, R, SN, S_C, S_M	0 - max_u	N/A	CHASSIS
Chassis	ZONE_CFGSZ_PER	DS, NAG, NR, L	none	E, F, R, SN	0 - max_f	%	ALL
Circuit	CIR_PKTLOSS	L	min, hour, day	E, F, R, SN	0 - max_f	%	EXTN

Type	Monitoring System	Flags	TimeBase	Actions	Thresholds	Unit	System
Circuit	CIR_STATE	L	min, hour, day	E, FN, F, R, SN	0 - max_u	N/A	EXTN
Circuit	CIR_UTIL	L	min, hour, day	E, F, R, SN	0 - max_f	%	EXTN
Circuit	IP_JITTER	L	none	E, F, R, SN	0 - max_f	%	EXTN
Circuit	IP_PKTLOSS	L	min, hour, day	E, F, R, SN	0 - max_f	%	EXTN
Circuit	IP_RTT	L	none	E, F, R, SN	0 - max_u	Milliseconds	EXTN
Circuit	IP_UTIL	L	min, hour, day	E, F, R, SN	0 - max_f	%	EXTN
Circuit	JITTER	L	none	E, F, R, SN	0 - max_f	%	EXTN
Circuit	RTT	L	none	E, F, R, SN	0 - max_u	Milliseconds	EXTN
Circuit QoS	PKTLOSS	L	min, hour, day	E, F, R, SN	0 - max_f	%	EXTN
Circuit QoS	STATE_CHG	L	min, hour, day	E, FN, F, R, SN	0 - max_u	N/A	EXTN
Circuit QoS	UTIL	L	min, hour, day	E, F, R, SN	0 - max_f	%	EXTN
DP	IP_EXTN_FLOW	DS, L	none	E, F, R, SN	0 - max_u	%	EXTN
Fan	FAN_STATE	NQ	none	E, F, R, SN	*states	N/A	ALL
Flow	IO_ABORT	NAG, NR, L	10sec	E, F, R, SN	0 - max_u	Aborts	G7_FLOW
Flow	IO_ERROR	NAG, NR, L	10sec	E, F, R, SN	0 - max_u	IO-ERRORS	G7_FLOW
Flow	IO_TIMEOUT	NAG, NR, L	10sec	E, F, R, SN	0 - max_u	Timeouts	G7_FLOW
Flow	MAX_RD_PENDING_IO	NAG, NR, L	10sec	E, F, R, SN	0 - max_u	IOs	G7_FLOW
Flow	MAX_WR_PENDING_IO	NAG, NR, L	10sec	E, F, R, SN	0 - max_u	IOs	G7_FLOW
Flow	RD_1stDATA_TIME	NAG, NR, L	IO, 10sec	E, F, R, SN	0 - max_u64	Microseconds	G6_FLOW
Flow	RD_PENDING_IO	NAG, NR, L	10sec	E, F, R, SN	0 - max_u64	IOs	G6_FLOW
Flow	RD_STATUS_TIME	NAG, NR, L	IO, 10sec	E, F, R, SN	0 - max_u64	Microseconds	G6_FLOW
Flow	RESERVE	NAG, NR, L	10sec	E, F, R, SN	0 - max_u	Reserves	G7_FLOW
Flow	WR_1stXFER_RDY	NAG, NR, L	IO, 10sec	E, F, R, SN	0 - max_u64	Microseconds	G6_FLOW
Flow	WR_PENDING_IO	NAG, NR, L	10sec	E, F, R, SN	0 - max_u64	IOs	G6_FLOW
Flow	WR_STATUS_TIME	NAG, NR, L	IO, 10sec	E, F, R, SN	0 - max_u64	Microseconds	G6_FLOW
GE Port	GE_CRC	L	min, hour, day	E, F, R, SN	0 - max_u	CRCs	EXTN
GE Port	GE_LOS_OF_SIG	L	min, hour, day	E, F, R, SN	0 - max_u	LOS	EXTN

Type	Monitoring System	Flags	TimeBase	Actions	Thresholds	Unit	System
PID	IT_FLOW	NAG, NR, NQ, L	none	E, F, R, SN	0 - max_u	IT-Flow(s)	ALL
Port	C3TXTO	L	min, hour, day	D, E, FN, F, R, SN	0 - max_u	Timeouts	ALL
Port	CRC	L	min, hour, day	D, E, FN, F, FP, R, SN	0 - max_u	CRCs	ALL
Port	DEV_LATENCY_IMPACT		none	D, E, FN, F, FP, R, SD, SN, T, U	*states	N/A	ALL
Port	DEV_LOGIN_DIST	NAG, L	none	E, F, R, R_B, SN	*states	N/A	ALL
Port	DEV_NPIV_LOGINS	L	none	E, F, R, SN	0 - max_f	%	ALL
Port	ENCR_BLK	R, L	min, hour, day	D, E, FN, F, R, SN	0 - max_u	Errors	ENCR
Port	ENCR_DISC	R, L	none	D, E, FN, F, R, SN	0 - max_u	Errors	ENCR
Port	ENCR_SHRT_FRM	R, L	min, hour, day	E, F, R, SN	0 - max_u	Errors	ENCR
Port	ITW	L	min, hour, day	D, E, FN, F, FP, R, SN	0 - max_u	ITWs	ALL
Port	LF	L	min, hour, day	E, F, R, SN	0 - max_u	LFs	ALL
Port	LOSS_SIGNAL	L	min, hour, day	E, F, R, SN	0 - max_u	LOS	ALL
Port	LOSS_SYNC	L	min, hour, day	E, F, R, SN	0 - max_u	SyncLoss	ALL
Port	LR	L	min, hour, day	D, E, FN, F, R, SN	0 - max_u	LRs	ALL
Port	MAX_RD_PENDING_IO	NAG, NR, L	10sec	E, F, R, SN	0 - max_u	IOs	G7_FLOW
Port	MAX_WR_PENDING_IO	NAG, NR, L	10sec	E, F, R, SN	0 - max_u	IOs	G7_FLOW
Port	PE	L	min, hour, day	D, E, FN, F, R, SN	0 - max_u	Errors	ALL
Port	PORT_BANDWIDTH		none	E, F, FP, R, SD, SN, U	*states	N/A	ALL
Port	RD_1stDATA_TIME_VIOL	NAG, NR, L	10sec	E, F, R, SN	0 - max_f	%	G7_FLOW
Port	RD_STATUS_TIME_VIOL	NAG, NR, L	10sec	E, F, R, SN	0 - max_f	%	G7_FLOW
Port	ROUTING_ERR	NAG	sec, min, hour, day	E, F, R, SN	0 - max_u	REs	ALL
Port	RX		min, hour, day	E, F, R, SN	0 - max_f	%	ALL
Port	RX_IOPS	NAG, NR, L	sec	E, F, R, SN	0 - max_u	IOPS	ALL

Type	Monitoring System	Flags	TimeBase	Actions	Thresholds	Unit	System
Port	SFP_STATE	NQ	none	E, F, R, SN	*states	N/A	ALL
Port	STATE_CHG	L	min, hour, day	D, E, FN, F, R, SN	0 - max_u	N/A	ALL
Port	TX		sec, min, hour, day	E, F, R, SN	0 - max_f	%	ALL
Port	UTIL		min, hour, day	E, F, R, SN	0 - max_f	%	ALL
Port	WR_1stXFER_RDY_VIOL	NAG, NR, L	10sec	E, F, R, SN	0 - max_f	%	G7_FLOW
Port	WR_STATUS_TIME_VIOL	NAG, NR, L	10sec	E, F, R, SN	0 - max_f	%	G7_FLOW
Power supply	PS_STATE	NQ	none	E, F, R, SN	*states	N/A	ALL
SFP	CURRENT	L	none	E, F, R, SF_M, SN	0 - max_u	mAmps	ALL
SFP	PWR_HRS	L	none	E, F, R, SN	0 - max_u	Hours	ALL
SFP	RXP	L	none	D, E, FN, F, R, SF_M, SN	0 - max_u	uW	ALL
SFP	SFP_TEMP	L	none	E, F, R, SF_M, SN	-40 - 100	Centigrade	ALL
SFP	TXP	L	none	D, E, FN, F, R, SF_M, SN	0 - max_u	uW	ALL
SFP	VOLTAGE	L	none	E, F, R, SF_M, SN	0 - max_u	mVolts	ALL
Switch	BB_FCR_CNT	NAG, NR, L	none	E, F, R, SN	0 - max_u	N/A	ALL
Switch	DID_CHG	NAG, NR, L	min, hour, day	E, F, R, SN	0 - max_u	Changes	ALL
Switch	EPORT_DOWN	NAG, NR, L	min, hour, day	E, F, R, SN	0 - max_u	Events	ALL
Switch	ERR_PORTS	NR, L	none	E, F, R, SN, S_C, S_M	0 - max_f	%	ALL
Switch	FAB_CFG	NAG, NR, L	min, hour, day	E, F, R, SN	0 - max_u	Changes	ALL
Switch	FAB_SEG	NAG, NR, L	min, hour, day	E, F, R, SN	0 - max_u	Segmentations	ALL
Switch	FAULTY_PORTS	NR, L	none	E, F, R, SN, S_C, S_M	0 - max_f	%	ALL
Switch	FLOGI	NR, L	min, hour, day	E, F, R, SN	0 - max_u	Logins	ALL
Switch	L2_DEVCNT_PER	NAG, NR, L	none	E, F, R, SN	0 - max_f	%	ALL
Switch	LSAN_DEVCNT_PER	NAG, NR, L	none	E, F, R, SN	0 - max_f	%	ALL

Type	Monitoring System	Flags	TimeBase	Actions	Thresholds	Unit	System
Switch	MARG_PORTS	NR, L	none	E, F, R, SN, S_C, S_M	0 - max_f	%	ALL
Switch	MARG_SFPS	NR, L	none	E, F, R, SN, S_C, S_M	0 - max_f	%	ALL
Switch	MISSING_SFP	NR	none	E, F, R, SN, S_C, S_M	0 - max_f	%	ALL
Switch	SEC_AUTH_FAIL	NAG, L	min, hour, day	E, F, R, SN	0 - max_u	Violations	ALL
Switch	SEC_CMD	NAG, L	min, hour, day	E, F, R, SN	0 - max_u	Violations	ALL
Switch	SEC_DCC	NAG, L	min, hour, day	E, F, R, SN	0 - max_u	Violations	ALL
Switch	SEC_HTTP	DS, L	min, hour, day	E, F, R, SN	0 - max_u	Violations	ALL
Switch	SEC_LV	DS, L	min, hour, day	E, F, R, SN	0 - max_u	Violations	ALL
Switch	SEC_SCC	NAG, L	min, hour, day	E, F, R, SN	0 - max_u	Violations	ALL
Switch	SEC_TELNET	DS, L	min, hour, day	E, F, R, SN	0 - max_u	Violations	ALL
Switch	SEC_TS	DS, NAG, L	min, hour, day	E, F, R, SN	0 - max_u	Violations	ALL
Switch	ZONE_CFGSZ_PER	NAG, NR, L	none	E, F, R, SN	0 - max_f	%	ALL
Switch	ZONE_CHG	NAG, NR, L	min, hour, day	E, F, R, SN	0 - max_u	Changes	ALL
Temperature sensor	TEMP	NR	none	E, F, R, SN	*states	N/A	ALL
Tunnel	IP_UTIL	L	min, hour, day	E, F, R, SN	0 - max_f	%	EXTN
Tunnel	STATE_CHG	L	min, hour, day	E, F, R, SN	0 - max_u	N/A	EXTN
Tunnel	UTIL	L	min, hour, day	E, F, R, SN	0 - max_f	%	EXTN
Tunnel QoS	PKTLOSS	L	min, hour, day	E, F, R, SN	0 - max_f	%	EXTN
Tunnel QoS	UTIL	L	min, hour, day	E, F, R, SN	0 - max_f	%	EXTN
VTAP_PORT	AVG_PENDING_IOS	NAG, NR, L	none	E, F, R, SN	0 - max_u64	IOs	ALL
VTAP_PORT	AVG_ROS	NAG, NR, L	none	E, F, R, SN	0 - max_f	%	ALL
VTAP_PORT	MAX_PENDING_IOS	NAG, NR, L	none	E, F, R, SN	0 - max_u64	IOs	ALL
VTAP_PORT	MAX_ROS	NAG, NR, L	none	E, F, R, SN	0 - max_f	%	ALL
WWN	WWN	NQ	none	E, F, R, SN	*states	N/A	CHASSIS

Monitoring Systems Support Matrix (Secondary)

The following table provides secondary information that is optional to clone an existing default policy and add to your custom rule. When you create a custom rule, you can use this information to account for the constraints associated with each parameter of that rule.

Type	Monitoring System	Min QT (Seconds)	Timebase (RoR)	Actions (RoR)	DB
ASIC	VTAP_IOPS	1	N/A	N/A	Switch Resource
BE Port	BAD_OS	60	hour, day, week	E, R, SN	BE Port Health
BE Port	BE_LATENCY_IMPACT	60	hour, day, week	E, R, SN	Fabric Performance Impact
BE Port	CRC	60	hour, day, week	E, R, SN	BE Port Health
BE Port	FRM_LONG	60	hour, day, week	E, R, SN	BE Port Health
BE Port	FRM_TRUNC	60	hour, day, week	E, R, SN	BE Port Health
BE Port	ITW	60	hour, day, week	E, R, SN	BE Port Health
BE Port	LR	60	hour, day, week	E, R, SN	BE Port Health
Blade	BLADE_STATE	N/S	hour, day, week	E, R, SN	FRU Health
Certificate	DAYS_TO_EXPIRE	86400	N/A	N/A	Security Violations
Chassis	AMP_RX_IOPS	1	N/A	N/A	AMP Health
Chassis	ASC_UPLOAD_FAILURE	604800	N/A	N/A	—
Chassis	BAD_FAN	60	N/A	N/A	—
Chassis	BAD_PWR	60	N/A	N/A	—
Chassis	BAD_TEMP	60	N/A	N/A	—
Chassis	CPU	120	hour, day, week	E, R, SN	Switch Resource
Chassis	DOWN_CORE	60	N/A	N/A	—
Chassis	ETH_MGMT_PORT_STATE	N/S	N/A	N/A	Switch Resource
Chassis	EXPIRED_CERTS	86400	N/A	N/A	Security Violations
Chassis	FAN_AIRFLOW_MISMATCH	6	N/A	N/A	—
Chassis	FAULTY_BLADE	60	N/A	N/A	—
Chassis	FLASH_USAGE	60	N/A	N/A	Switch Resource
Chassis	HA_SYNC	60	N/A	N/A	—
Chassis	ITL_RES_USAGE	60	N/A	N/A	Switch Resource
Chassis	IT_RES_USAGE	60	N/A	N/A	Switch Resource
Chassis	MEMORY_USAGE	120	hour, day, week	E, R, SN	Switch Resource
Chassis	MEMORY_USAGE_STATE	120	hour, day, week	E, R, SN	Switch Resource
Chassis	SYSTEM_TEMP	6	N/A	N/A	—
Chassis	TRUFOS_CERT_DAYS_TO_EXPIRE	86400	N/A	N/A	—
Chassis	TRUFOS_CERT_EXPIRED	60	N/A	N/A	—
Chassis	TRUFOS_CERT_INSTALLED	86400	N/A	N/A	—
Chassis	WWN_DOWN	60	N/A	N/A	—

Type	Monitoring System	Min QT (Seconds)	Timebase (RoR)	Actions (RoR)	DB
Chassis	ZONE_CFGSZ_PER	86400	N/A	N/A	Fabric State Changes
Circuit	CIR_PKTLOSS	60	hour, day, week	E, R, SN	Extension Health
Circuit	CIR_STATE	60	hour, day, week	E, FN, R, SN	Extension Health
Circuit	CIR_UTIL	60	hour, day, week	E, R, SN	Extension Health
Circuit	IP_JITTER	60	hour, day, week	E, R, SN	Extension Health
Circuit	IP_PKTLOSS	60	hour, day, week	E, R, SN	Extension Health
Circuit	IP_RTT	60	hour, day, week	E, R, SN	Extension Health
Circuit	IP_UTIL	60	hour, day, week	E, R, SN	Extension Health
Circuit	JITTER	60	hour, day, week	E, R, SN	Extension Health
Circuit	RTT	60	hour, day, week	E, R, SN	Extension Health
Circuit QoS	PKTLOSS	60	hour, day, week	E, R, SN	Extension Health
Circuit QoS	STATE_CHG	60	hour, day, week	E, FN, R, SN	Extension Health
Circuit QoS	UTIL	60	hour, day, week	E, R, SN	Extension Health
DP	IP_EXTN_FLOW	10	hour, day, week	E, R, SN	Extension Health
Fan	FAN_STATE	N/S	hour, day, week	E, R, SN	FRU Health
Flow	IO_ABORT	10	N/A	N/A	IO Health
Flow	IO_ERROR	10	N/A	N/A	IO Health
Flow	IO_TIMEOUT	10	N/A	N/A	IO Health
Flow	MAX_RD_PENDING_IO	10	N/A	N/A	IO Latency
Flow	MAX_WR_PENDING_IO	10	N/A	N/A	IO Latency
Flow	RD_1stDATA_TIME	10	N/A	N/A	IO Latency
Flow	RD_PENDING_IO	10	N/A	N/A	IO Latency
Flow	RD_STATUS_TIME	10	N/A	N/A	IO Latency
Flow	RESERVE	10	N/A	N/A	IO Health
Flow	WR_1stXFER_RDY	10	N/A	N/A	IO Latency
Flow	WR_PENDING_IO	10	N/A	N/A	IO Latency
Flow	WR_STATUS_TIME	10	N/A	N/A	IO Latency
GE Port	GE_CRC	60	hour, day, week	E, R, SN	Extension GE Port Health
GE Port	GE_LOS_OF_SIG	60	hour, day, week	E, R, SN	Extension GE Port Health
PID	IT_FLOW	N/S	N/A	N/A	Fabric Performance Impact
Port	C3TXTO	60	hour, day, week	D, E, FN, R, SN	Port Health
Port	CRC	60	hour, day, week	D, E, FN, FP, R, SN	Port Health

Type	Monitoring System	Min QT (Seconds)	Timebase (RoR)	Actions (RoR)	DB
Port	DEV_LATENCY_IMPACT	10	hour, day, week	D, E, FN, FP, R, SD, SN, T, U	Fabric Performance Impact
Port	DEV_LOGIN_DIST	60	hour, day, week	E, R, R_B, SN	Fabric Performance Impact
Port	DEV_NPIV_LOGINS	3600	day, week	E, R, SN	Port Health
Port	ENCR_BLK	60	hour, day, week	D, E, FN, R, SN	Port Health
Port	ENCR_DISC	60	hour, day, week	D, E, FN, R, SN	Port Health
Port	ENCR_SHRT_FRM	60	hour, day, week	E, R, SN	Port Health
Port	ITW	60	hour, day, week	D, E, FN, FP, R, SN	Port Health
Port	LF	60	hour, day, week	E, R, SN	Port Health
Port	LOSS_SIGNAL	60	hour, day, week	E, R, SN	Port Health
Port	LOSS_SYNC	60	hour, day, week	E, R, SN	Port Health
Port	LR	60	hour, day, week	D, E, FN, R, SN	Port Health
Port	MAX_RD_PENDING_IO	10	N/A	N/A	IO Latency
Port	MAX_WR_PENDING_IO	10	N/A	N/A	IO Latency
Port	PE	60	hour, day, week	D, E, FN, R, SN	Port Health
Port	PORT_BANDWIDTH	10	hour, day, week	E, FP, R, SD, SN, U	Fabric Performance Impact
Port	RD_1stDATA_TIME_VIOL	10	N/A	N/A	IO Latency
Port	RD_STATUS_TIME_VIOL	10	N/A	N/A	IO Latency
Port	ROUTING_ERR	1	hour, day, week	E, R, SN	Port Health
Port	RX	60	hour, day, week	E, R, SN	Fabric Performance Impact
Port	RX_IOPS	1	N/A	N/A	AMP Health
Port	SFP_STATE	N/S	hour, day, week	E, R, SN	FRU Health
Port	STATE_CHG	60	hour, day, week	D, E, FN, R, SN	Port Health
Port	TX	1	hour, day, week	E, R, SN	Fabric Performance Impact
Port	UTIL	60	hour, day, week	E, R, SN	Fabric Performance Impact
Port	WR_1stXFER_RDY_VIOL	10	N/A	N/A	IO Latency
Port	WR_STATUS_TIME_VIOL	10	N/A	N/A	IO Latency
Power supply	PS_STATE	N/S	hour, day, week	E, R, SN	FRU Health
SFP	CURRENT	720	hour, day, week	E, R, SN	Port Health
SFP	PWR_HRS	720	hour, day, week	E, R, SN	Port Health
SFP	RXP	720	hour, day, week	E, R, SN	Port Health
SFP	SFP_TEMP	720	hour, day, week	E, R, SN	Port Health

Type	Monitoring System	Min QT (Seconds)	Timebase (RoR)	Actions (RoR)	DB
SFP	TXP	720	hour, day, week	E, R, SN	Port Health
SFP	VOLTAGE	720	hour, day, week	E, R, SN	Port Health
Switch	BB_FCR_CNT	86400	N/A	N/A	Fabric State Changes
Switch	DID_CHG	60	N/A	N/A	Fabric State Changes
Switch	EPORT_DOWN	60	N/A	N/A	Fabric State Changes
Switch	ERR_PORTS	60	N/A	N/A	—
Switch	FAB_CFG	60	N/A	N/A	Fabric State Changes
Switch	FAB_SEG	60	N/A	N/A	Fabric State Changes
Switch	FAULTY_PORTS	60	N/A	N/A	—
Switch	FLOGI	60	N/A	N/A	Fabric State Changes
Switch	L2_DEVCNT_PER	86400	N/A	N/A	Fabric State Changes
Switch	LSAN_DEVCNT_PER	86400	N/A	N/A	Fabric State Changes
Switch	MARG_PORTS	60	N/A	N/A	—
Switch	MARG_SFPS	60	N/A	N/A	—
Switch	MISSING_SFP	60	N/A	N/A	—
Switch	SEC_AUTH_FAIL	60	hour, day, week	E, R, SN	Security Violations
Switch	SEC_CMD	60	hour, day, week	E, R, SN	Security Violations
Switch	SEC_DCC	60	hour, day, week	E, R, SN	Security Violations
Switch	SEC_HTTP	60	hour, day, week	E, R, SN	Security Violations
Switch	SEC_LV	60	hour, day, week	E, R, SN	Security Violations
Switch	SEC_SCC	60	hour, day, week	E, R, SN	Security Violations
Switch	SEC_TELNET	60	hour, day, week	E, R, SN	Security Violations
Switch	SEC_TS	60	hour, day, week	E, R, SN	Security Violations
Switch	ZONE_CFGSZ_PER	86400	N/A	N/A	Fabric State Changes
Switch	ZONE_CHG	60	N/A	N/A	Fabric State Changes
Temperature sensor	TEMP	60	N/A	N/A	Switch Resource
Tunnel	IP_UTIL	60	hour, day, week	E, R, SN	Extension Health
Tunnel	STATE_CHG	60	hour, day, week	E, R, SN	Extension Health
Tunnel	UTIL	60	hour, day, week	E, R, SN	Extension Health

Type	Monitoring System	Min QT (Seconds)	Timebase (RoR)	Actions (RoR)	DB
Tunnel QoS	PKTLOSS	60	hour, day, week	E, R, SN	Extension Health
Tunnel QoS	UTIL	60	hour, day, week	E, R, SN	Extension Health
VTAP_PORT	AVG_PENDING_IOS	10	N/A	N/A	—
VTAP_PORT	AVG_ROS	10	N/A	N/A	—
VTAP_PORT	MAX_PENDING_IOS	10	N/A	N/A	—
VTAP_PORT	MAX_ROS	10	N/A	N/A	—
WWN	WWN	N/S	hour, day, week	E, R, SN	FRU Health

Default Monitoring Rules Overview

This section describes the default monitoring rules available in MAPS. You can use default rules as a template and create custom rules for your setup or fabric.

A monitoring rule name has the following structure:

defALL_HOST_PORTSTX_60 | **ALL_HOST_PORTS**(TX/HOUR>60) RASLOG,SNMP,EMAIL

Consider the following notes when you interpret the structure of a default monitoring rule:

- The prefix **def** denotes that the rule is a default rule.
- **ALL_HOST_PORTS** denotes the name of the group.
- **TX** denotes the name of the monitoring system.
- **60** denotes the configured threshold value for when potential problems might occur and alerts are generated.

The following example shows a sample rule that explains how to interpret the contents of the default monitoring rules table.

Rule	Group	Monitoring System	Threshold	Timebase	Policy
defALL_10GLWL _SFPCURRENT _95	ALL_10GLWL_ SFP	CURRENT	>=95	NONE	M, A, C

Consider the following notes when you interpret the structure of a default monitoring rules table:

- The Rule column lists the name of the default MAPS rule.
- The Group column lists the name of the default MAPS group.
- The Monitoring System column lists the name of the monitoring system.
- The Threshold column lists the threshold value for the selected rule.
- The Timebase column lists the timebase. The timebase is the time interval between samples and the user-defined threshold values.
- The Policy column lists the rule that is associated with the policy. The options are base (B), conservative (C), moderate (M), and aggressive (A).

The following items are the timebase values:

- **SEC** – Seconds
- **10SEC** – 10 seconds
- **HOUR** – Hours
- **DAY** – Days
- **MIN** – Minutes
- **5MIN** – 5 minutes
- **NONE** – Not Applicable

Default Monitoring Rules

The following table lists the default rules, monitoring threshold policy value, actions triggered, and conditions.

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_100G_QSFPCURRENT_10	ALL_FCOE_100G_SR4_QSFP	CURRENT	>=10	NONE	M, A, C
defALL_100G_QSFPCURRENT_2	ALL_FCOE_100G_SR4_QSFP	CURRENT	<=2	NONE	M, A, C
defALL_100G_QSFPRXP_2187	ALL_FCOE_100G_SR4_QSFP	RXP	>=2187	NONE	M, A, C
defALL_100G_QSFPRXP_60	ALL_FCOE_100G_SR4_QSFP	RXP	<=60	NONE	M, A, C
defALL_100G_QSFPSFP_TEMP_75	ALL_FCOE_100G_SR4_QSFP	SFP_TEMP	>=75	NONE	M, A, C
defALL_100G_QSFPSFP_TEMP_n5	ALL_FCOE_100G_SR4_QSFP	SFP_TEMP	<=	NONE	M, A, C
defALL_100G_QSFPTXP_3467	ALL_FCOE_100G_SR4_QSFP	TXP	>=3467	NONE	—
defALL_100G_QSFPTXP_48	ALL_FCOE_100G_SR4_QSFP	TXP	<=48	NONE	—
defALL_100G_QSFPVOLTAGE_2970	ALL_FCOE_100G_SR4_QSFP	VOLTAGE	<=2970	NONE	M, A, C
defALL_100G_QSFPVOLTAGE_3630	ALL_FCOE_100G_SR4_QSFP	VOLTAGE	>=3630	NONE	M, A, C
defALL_100M_16GSWL_QSFPCURRENT_1	ALL_100M_16GSWL_QSFP	CURRENT	<=1	NONE	M, A, C
defALL_100M_16GSWL_QSFPCURRENT_10	ALL_100M_16GSWL_QSFP	CURRENT	>=10	NONE	M, A, C
defALL_100M_16GSWL_QSFPRXP_2187	ALL_100M_16GSWL_QSFP	RXP	>=2187	NONE	M, A, C
defALL_100M_16GSWL_QSFPRXP_44	ALL_100M_16GSWL_QSFP	RXP	<=44	NONE	M, A, C
defALL_100M_16GSWL_QSFPSFP_TEMP_85	ALL_100M_16GSWL_QSFP	SFP_TEMP	>=85	NONE	M, A, C
defALL_100M_16GSWL_QSFPSFP_TEMP_n5	ALL_100M_16GSWL_QSFP	SFP_TEMP	<=	NONE	M, A, C
defALL_100M_16GSWL_QSFPVOLTAGE_2970	ALL_100M_16GSWL_QSFP	VOLTAGE	<=2970	NONE	M, A, C
defALL_100M_16GSWL_QSFPVOLTAGE_3630	ALL_100M_16GSWL_QSFP	VOLTAGE	>=3630	NONE	M, A, C
defALL_10GLWL_SFPCURRENT_8	ALL_10GLWL_SFP	CURRENT	<=8	NONE	M, A, C
defALL_10GLWL_SFPCURRENT_95	ALL_10GLWL_SFP	CURRENT	>=95	NONE	M, A, C
defALL_10GLWL_SFPRXP_14	ALL_10GLWL_SFP	RXP	<=14	NONE	M, A, C
defALL_10GLWL_SFPRXP_2230	ALL_10GLWL_SFP	RXP	>=2230	NONE	M, A, C

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_10GLWL_SFPSFP_TEMP_90	ALL_10GLWL_SFP	SFP_TEMP	>=90	NONE	M, A, C
defALL_10GLWL_SFPSFP_TEMP_n5	ALL_10GLWL_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_10GLWL_SFPTXP_2230	ALL_10GLWL_SFP	TXP	>=2230	NONE	M, A, C
defALL_10GLWL_SFPTXP_60	ALL_10GLWL_SFP	TXP	<=60	NONE	M, A, C
defALL_10GLWL_SFPVOLTAGE_2970	ALL_10GLWL_SFP	VOLTAGE	<=2970	NONE	M, A, C
defALL_10GLWL_SFPVOLTAGE_3600	ALL_10GLWL_SFP	VOLTAGE	>=3600	NONE	M, A, C
defALL_10GSWL_SFPCURRENT_10	ALL_10GSWL_SFP	CURRENT	>=10	NONE	M, A, C
defALL_10GSWL_SFPCURRENT_2	ALL_10GSWL_SFP	CURRENT	<=2	NONE	M, A, C
defALL_10GSWL_SFPRXP_1999	ALL_10GSWL_SFP	RXP	>=1999	NONE	M, A, C
defALL_10GSWL_SFPRXP_31	ALL_10GSWL_SFP	RXP	<=31	NONE	M, A, C
defALL_10GSWL_SFPSFP_TEMP_90	ALL_10GSWL_SFP	SFP_TEMP	>=90	NONE	M, A, C
defALL_10GSWL_SFPSFP_TEMP_n5	ALL_10GSWL_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_10GSWL_SFPTXP_125	ALL_10GSWL_SFP	TXP	<=125	NONE	M, A, C
defALL_10GSWL_SFPTXP_1999	ALL_10GSWL_SFP	TXP	>=1999	NONE	M, A, C
defALL_10GSWL_SFPVOLTAGE_3000	ALL_10GSWL_SFP	VOLTAGE	<=3000	NONE	M, A, C
defALL_10GSWL_SFPVOLTAGE_3600	ALL_10GSWL_SFP	VOLTAGE	>=3600	NONE	M, A, C
defALL_16GLWL_SFPCURRENT_1	ALL_16GLWL_SFP	CURRENT	<=1	NONE	M, A, C
defALL_16GLWL_SFPCURRENT_70	ALL_16GLWL_SFP	CURRENT	>=70	NONE	
defALL_16GLWL_SFPCURRENT_90	ALL_16GLWL_SFP	CURRENT	>=90	NONE	M, A, C
defALL_16GLWL_SFPRXP_1995	ALL_16GLWL_SFP	RXP	>=1995	NONE	M, A, C
defALL_16GLWL_SFPRXP_28	ALL_16GLWL_SFP	RXP	<=28	NONE	M, A, C
defALL_16GLWL_SFPSFP_TEMP_90	ALL_16GLWL_SFP	SFP_TEMP	>=90	NONE	M, A, C
defALL_16GLWL_SFPSFP_TEMP_n5	ALL_16GLWL_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_16GLWL_SFPTXP_125	ALL_16GLWL_SFP	TXP	<=126	NONE	—
defALL_16GLWL_SFPTXP_126	ALL_16GLWL_SFP	TXP	<=126	NONE	M, A, C
defALL_16GLWL_SFPTXP_1995	ALL_16GLWL_SFP	TXP	>=1995	NONE	M, A, C
defALL_16GLWL_SFPVOLTAGE_3000	ALL_16GLWL_SFP	VOLTAGE	<=3000	NONE	M, A, C
defALL_16GLWL_SFPVOLTAGE_3600	ALL_16GLWL_SFP	VOLTAGE	>=3600	NONE	M, A, C
defALL_16GSWL_SFPCURRENT_12	ALL_16GSWL_SFP	CURRENT	>=12	NONE	M, A, C
defALL_16GSWL_SFPCURRENT_2	ALL_16GSWL_SFP	CURRENT	<=2	NONE	M, A, C
defALL_16GSWL_SFPRXP_1259	ALL_16GSWL_SFP	RXP	>=1259	NONE	M, A, C
defALL_16GSWL_SFPRXP_32	ALL_16GSWL_SFP	RXP	<=32	NONE	M, A, C

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_16GSWL_SFPSFP_TEMP_85	ALL_16GSWL_SFP	SFP_TEMP	>=85	NONE	M, A, C
defALL_16GSWL_SFPSFP_TEMP_n5	ALL_16GSWL_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_16GSWL_SFPTXP_1259	ALL_16GSWL_SFP	TXP	>=1259	NONE	M, A, C
defALL_16GSWL_SFPTXP_25	ALL_16GSWL_SFP	TXP	<=25	NONE	—
defALL_16GSWL_SFPTXP_251	ALL_16GSWL_SFP	TXP	<=251	NONE	M, A, C
defALL_16GSWL_SFPVOLTAGE_3000	ALL_16GSWL_SFP	VOLTAGE	<=3000	NONE	M, A, C
defALL_16GSWL_SFPVOLTAGE_3600	ALL_16GSWL_SFP	VOLTAGE	>=3600	NONE	M, A, C
defALL_25Km_16GLWL_SFPCURRENT_1	ALL_25Km_16GLWL_SFP	CURRENT	<=1	NONE	M, A, C
defALL_25Km_16GLWL_SFPCURRENT_90	ALL_25Km_16GLWL_SFP	CURRENT	>=90	NONE	M, A, C
defALL_25Km_16GLWL_SFPRXP_2238	ALL_25Km_16GLWL_SFP	RXP	>=2238	NONE	M, A, C
defALL_25Km_16GLWL_SFPRXP_28	ALL_25Km_16GLWL_SFP	RXP	<=28	NONE	M, A, C
defALL_25Km_16GLWL_SFPSFP_TEMP_75	ALL_25Km_16GLWL_SFP	SFP_TEMP	>=75	NONE	M, A, C
defALL_25Km_16GLWL_SFPSFP_TEMP_n5	ALL_25Km_16GLWL_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_25Km_16GLWL_SFPTXP_158	ALL_25Km_16GLWL_SFP	TXP	<=158	NONE	M, A, C
defALL_25Km_16GLWL_SFPTXP_4466	ALL_25Km_16GLWL_SFP	TXP	>=4466	NONE	M, A, C
defALL_25Km_16GLWL_SFPVOLTAGE_2850	ALL_25Km_16GLWL_SFP	VOLTAGE	<=2850	NONE	—
defALL_25Km_16GLWL_SFPVOLTAGE_3000	ALL_25Km_16GLWL_SFP	VOLTAGE	<=3000	NONE	M, A, C
defALL_25Km_16GLWL_SFPVOLTAGE_3600	ALL_25Km_16GLWL_SFP	VOLTAGE	>=3600	NONE	M, A, C
defALL_25Km_16GLWL_SFPVOLTAGE_3750	ALL_25Km_16GLWL_SFP	VOLTAGE	>=3750	NONE	—
defALL_25Km_32GELWL_JBA_SFPCURRENT_6	ALL_25Km_32GELWL_JBA_SFP	CURRENT	<=6	NONE	M, A, C
defALL_25Km_32GELWL_JBA_SFPCURRENT_90	ALL_25Km_32GELWL_JBA_SFP	CURRENT	>=90	NONE	M, A, C
defALL_25Km_32GELWL_JBA_SFPRXP_46	ALL_25Km_32GELWL_JBA_SFP	RXP	<=46	NONE	M, A, C
defALL_25Km_32GELWL_JBA_SFPRXP_6310	ALL_25Km_32GELWL_JBA_SFP	RXP	>=6310	NONE	M, A, C
defALL_25Km_32GELWL_JBA_SFPSFP_TEMP_85	ALL_25Km_32GELWL_JBA_SFP	SFP_TEMP	>=85	NONE	M, A, C

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_25Km_32GELWL_JBA_SFPSFP_TEMP_n5	ALL_25Km_32GELWL_JBA_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_25Km_32GELWL_JBA_SFPTXP_6310	ALL_25Km_32GELWL_JBA_SFP	TXP	>=6310	NONE	M, A, C
defALL_25Km_32GELWL_JBA_SFPTXP_794	ALL_25Km_32GELWL_JBA_SFP	TXP	<=794	NONE	M, A, C
defALL_25Km_32GELWL_JBA_SFPVOLTAGE_2970	ALL_25Km_32GELWL_JBA_SFP	VOLTAGE	<=2970	NONE	M, A, C
defALL_25Km_32GELWL_JBA_SFPVOLTAGE_3630	ALL_25Km_32GELWL_JBA_SFP	VOLTAGE	>=3630	NONE	M, A, C
defALL_25Km_32GELWL_SFPCURRENT_10	ALL_25Km_32GELWL_SFP	CURRENT	<=10	NONE	M, A, C
defALL_25Km_32GELWL_SFPCURRENT_130	ALL_25Km_32GELWL_SFP	CURRENT	>=130	NONE	M, A, C
defALL_25Km_32GELWL_SFPCURRENT_70	ALL_25Km_32GELWL_SFP	CURRENT	>=70	NONE	—
defALL_25Km_32GELWL_SFPRXP_1995	ALL_25Km_32GELWL_SFP	RXP	>=1995	NONE	M, A, C
defALL_25Km_32GELWL_SFPRXP_31	ALL_25Km_32GELWL_SFP	RXP	<=31	NONE	M, A, C
defALL_25Km_32GELWL_SFPSFP_TEMP_75	ALL_25Km_32GELWL_SFP	SFP_TEMP	>=75	NONE	M, A, C
defALL_25Km_32GELWL_SFPSFP_TEMP_n5	ALL_25Km_32GELWL_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_25Km_32GELWL_SFPTXP_6309	ALL_25Km_32GELWL_SFP	TXP	>=6309	NONE	M, A, C
defALL_25Km_32GELWL_SFPTXP_794	ALL_25Km_32GELWL_SFP	TXP	<=794	NONE	M, A, C
defALL_25Km_32GELWL_SFPVOLTAGE_3000	ALL_25Km_32GELWL_SFP	VOLTAGE	<=3000	NONE	M, A, C
defALL_25Km_32GELWL_SFPVOLTAGE_3600	ALL_25Km_32GELWL_SFP	VOLTAGE	>=3600	NONE	M, A, C
defALL_25Km_64GELWL_SFPCURRENT_6	ALL_25Km_64GELWL_SFP	CURRENT	<=6	NONE	M, A, C
defALL_25Km_64GELWL_SFPCURRENT_90	ALL_25Km_64GELWL_SFP	CURRENT	>=90	NONE	M, A, C
defALL_25Km_64GELWL_SFPRXP_46	ALL_25Km_64GELWL_SFP	RXP	<=46	NONE	M, A, C
defALL_25Km_64GELWL_SFPRXP_6310	ALL_25Km_64GELWL_SFP	RXP	>=6310	NONE	M, A, C
defALL_25Km_64GELWL_SFPSFP_TEMP_85	ALL_25Km_64GELWL_SFP	SFP_TEMP	>=85	NONE	M, A, C

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_25Km_64GELWL_SFPSFP_TEMP_n5	ALL_25Km_64GELWL_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_25Km_64GELWL_SFPTXP_6310	ALL_25Km_64GELWL_SFP	TXP	>=6310	NONE	M, A, C
defALL_25Km_64GELWL_SFPTXP_794	ALL_25Km_64GELWL_SFP	TXP	<=794	NONE	M, A, C
defALL_25Km_64GELWL_SFPVOLTAGE_2970	ALL_25Km_64GELWL_SFP	VOLTAGE	<=2970	NONE	M, A, C
defALL_25Km_64GELWL_SFPVOLTAGE_3630	ALL_25Km_64GELWL_SFP	VOLTAGE	>=3630	NONE	M, A, C
defALL_2K_QSFPCURRENT_3	ALL_2K_QSFP	CURRENT	<=3	NONE	M, A, C
defALL_2K_QSFPCURRENT_39	ALL_2K_QSFP	CURRENT	>=39	NONE	M, A, C
defALL_2K_QSFPRXP_2000	ALL_2K_QSFP	RXP	>=2000	NONE	M, A, C
defALL_2K_QSFPRXP_25	ALL_2K_QSFP	RXP	<=25	NONE	M, A, C
defALL_2K_QSFPSFP_TEMP_85	ALL_2K_QSFP	SFP_TEMP	>=85	NONE	M, A, C
defALL_2K_QSFPSFP_TEMP_n15	ALL_2K_QSFP	SFP_TEMP	<=	NONE	M, A, C
defALL_2K_QSFPVOLTAGE_2900	ALL_2K_QSFP	VOLTAGE	<=2900	NONE	M, A, C
defALL_2K_QSFPVOLTAGE_3600	ALL_2K_QSFP	VOLTAGE	>=3600	NONE	M, A, C
defALL_2Km_32GLWL_QSFPCURRENT_1	ALL_2Km_32GLWL_QSFP	CURRENT	<=1	NONE	M, A, C
defALL_2Km_32GLWL_QSFPCURRENT_75	ALL_2Km_32GLWL_QSFP	CURRENT	>=75	NONE	M, A, C
defALL_2Km_32GLWL_QSFPRXP_3548	ALL_2Km_32GLWL_QSFP	RXP	>=3548	NONE	M, A, C
defALL_2Km_32GLWL_QSFPRXP_77	ALL_2Km_32GLWL_QSFP	RXP	<=77	NONE	M, A, C
defALL_2Km_32GLWL_QSFPSFP_TEMP_75	ALL_2Km_32GLWL_QSFP	SFP_TEMP	>=75	NONE	M, A, C
defALL_2Km_32GLWL_QSFPSFP_TEMP_n5	ALL_2Km_32GLWL_QSFP	SFP_TEMP	<=	NONE	M, A, C
defALL_2Km_32GLWL_QSFPTXP_223	ALL_2Km_32GLWL_QSFP	TXP	<=223	NONE	M, A, C
defALL_2Km_32GLWL_QSFPTXP_4466	ALL_2Km_32GLWL_QSFP	TXP	>=4466	NONE	M, A, C
defALL_2Km_32GLWL_QSFPVOLTAGE_3010	ALL_2Km_32GLWL_QSFP	VOLTAGE	<=3010	NONE	M, A, C
defALL_2Km_32GLWL_QSFPVOLTAGE_3604	ALL_2Km_32GLWL_QSFP	VOLTAGE	>=3604	NONE	M, A, C
defALL_2Km_GEN7LWL_QSFPCURRENT_10	ALL_2Km_GEN7LWL_QSFP	CURRENT	<=10	NONE	M, A, C
defALL_2Km_GEN7LWL_QSFPCURRENT_80	ALL_2Km_GEN7LWL_QSFP	CURRENT	>=80	NONE	M, A, C

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_2Km_GEN7LWL_QSFPRXP_63	ALL_2Km_GEN7LWL_QSFP	RXP	<=63	NONE	M, A, C
defALL_2Km_GEN7LWL_QSFPRXP_6309	ALL_2Km_GEN7LWL_QSFP	RXP	>=6309	NONE	M, A, C
defALL_2Km_GEN7LWL_QSFPSFP_TEMP_80	ALL_2Km_GEN7LWL_QSFP	SFP_TEMP	>=80	NONE	M, A, C
defALL_2Km_GEN7LWL_QSFPSFP_TEMP_n5	ALL_2Km_GEN7LWL_QSFP	SFP_TEMP	<=	NONE	M, A, C
defALL_2Km_GEN7LWL_QSFPTXP_177	ALL_2Km_GEN7LWL_QSFP	TXP	<=177	NONE	M, A, C
defALL_2Km_GEN7LWL_QSFPTXP_5623	ALL_2Km_GEN7LWL_QSFP	TXP	>=5623	NONE	M, A, C
defALL_2Km_GEN7LWL_QSFPVOLTAGE_2970	ALL_2Km_GEN7LWL_QSFP	VOLTAGE	<=2970	NONE	M, A, C
defALL_2Km_GEN7LWL_QSFPVOLTAGE_3630	ALL_2Km_GEN7LWL_QSFP	VOLTAGE	>=3630	NONE	M, A, C
defALL_32GLWL_JDB_SFPCURRENT_1	ALL_32GLWL_JDB_SFP	CURRENT	<=1	NONE	M, A, C
defALL_32GLWL_JDB_SFPCURRENT_110	ALL_32GLWL_JDB_SFP	CURRENT	>=110	NONE	M, A, C
defALL_32GLWL_JDB_SFPRXP_2512	ALL_32GLWL_JDB_SFP	RXP	>=2512	NONE	M, A, C
defALL_32GLWL_JDB_SFPRXP_43	ALL_32GLWL_JDB_SFP	RXP	<=43	NONE	M, A, C
defALL_32GLWL_JDB_SFPSFP_TEMP_80	ALL_32GLWL_JDB_SFP	SFP_TEMP	>=80	NONE	M, A, C
defALL_32GLWL_JDB_SFPSFP_TEMP_n10	ALL_32GLWL_JDB_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_32GLWL_JDB_SFPTXP_199	ALL_32GLWL_JDB_SFP	TXP	<=199	NONE	M, A, C
defALL_32GLWL_JDB_SFPTXP_3162	ALL_32GLWL_JDB_SFP	TXP	>=3162	NONE	M, A, C
defALL_32GLWL_JDB_SFPVOLTAGE_3000	ALL_32GLWL_JDB_SFP	VOLTAGE	<=3000	NONE	M, A, C
defALL_32GLWL_JDB_SFPVOLTAGE_3600	ALL_32GLWL_JDB_SFP	VOLTAGE	>=3600	NONE	M, A, C
defALL_32GLWL_SFPCURRENT_1	ALL_32GLWL_SFP	CURRENT	<=1	NONE	M, A, C
defALL_32GLWL_SFPCURRENT_60	ALL_32GLWL_SFP	CURRENT	>=60	NONE	—
defALL_32GLWL_SFPCURRENT_90	ALL_32GLWL_SFP	CURRENT	>=90	NONE	M, A, C
defALL_32GLWL_SFPRXP_1995	ALL_32GLWL_SFP	RXP	>=1995	NONE	M, A, C
defALL_32GLWL_SFPRXP_39	ALL_32GLWL_SFP	RXP	<=39	NONE	M, A, C
defALL_32GLWL_SFPSFP_TEMP_75	ALL_32GLWL_SFP	SFP_TEMP	>=75	NONE	M, A, C
defALL_32GLWL_SFPSFP_TEMP_n5	ALL_32GLWL_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_32GLWL_SFPTXP_125	ALL_32GLWL_SFP	TXP	<=125	NONE	M, A, C

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_32GLWL_SFPTXP_1584	ALL_32GLWL_SFP	TXP	>=1584	NONE	M, A, C
defALL_32GLWL_SFPVOLTAGE_3000	ALL_32GLWL_SFP	VOLTAGE	<=3000	NONE	M, A, C
defALL_32GLWL_SFPVOLTAGE_3600	ALL_32GLWL_SFP	VOLTAGE	>=3600	NONE	M, A, C
defALL_32GSWL_QSFPCURRENT_1	ALL_32GSWL_QSFP	CURRENT	<=1	NONE	M, A, C
defALL_32GSWL_QSFPCURRENT_10	ALL_32GSWL_QSFP	CURRENT	>=10	NONE	—
defALL_32GSWL_QSFPCURRENT_43	ALL_32GSWL_QSFP	CURRENT	>=43	NONE	M, A, C
defALL_32GSWL_QSFPRXP_2187	ALL_32GSWL_QSFP	RXP	>=2187	NONE	—
defALL_32GSWL_QSFPRXP_91	ALL_32GSWL_QSFP	RXP	<=91	NONE	—
defALL_32GSWL_QSFPSFP_TEMP_75	ALL_32GSWL_QSFP	SFP_TEMP	>=75	NONE	M, A, C
defALL_32GSWL_QSFPSFP_TEMP_n5	ALL_32GSWL_QSFP	SFP_TEMP	<=	NONE	M, A, C
defALL_32GSWL_QSFPTXP_3467	ALL_32GSWL_QSFP	TXP	>=3467	NONE	M, A, C
defALL_32GSWL_QSFPTXP_72	ALL_32GSWL_QSFP	TXP	<=72	NONE	M, A, C
defALL_32GSWL_QSFPVOLTAGE_2970	ALL_32GSWL_QSFP	VOLTAGE	<=2970	NONE	M, A, C
defALL_32GSWL_QSFPVOLTAGE_3630	ALL_32GSWL_QSFP	VOLTAGE	>=3630	NONE	M, A, C
defALL_32GSWL_SFPCURRENT_1	ALL_32GSWL_SFP	CURRENT	<=1	NONE	M, A, C
defALL_32GSWL_SFPCURRENT_13	ALL_32GSWL_SFP	CURRENT	>=13	NONE	M, A, C
defALL_32GSWL_SFPRXP_2187	ALL_32GSWL_SFP	RXP	>=2187	NONE	M, A, C
defALL_32GSWL_SFPRXP_63	ALL_32GSWL_SFP	RXP	<=63	NONE	M, A, C
defALL_32GSWL_SFPSFP_TEMP_85	ALL_32GSWL_SFP	SFP_TEMP	>=85	NONE	M, A, C
defALL_32GSWL_SFPSFP_TEMP_n5	ALL_32GSWL_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_32GSWL_SFPTXP_251	ALL_32GSWL_SFP	TXP	<=251	NONE	M, A, C
defALL_32GSWL_SFPTXP_3162	ALL_32GSWL_SFP	TXP	>=3162	NONE	M, A, C
defALL_32GSWL_SFPVOLTAGE_2970	ALL_32GSWL_SFP	VOLTAGE	<=2970	NONE	M, A, C
defALL_32GSWL_SFPVOLTAGE_3630	ALL_32GSWL_SFP	VOLTAGE	>=3630	NONE	M, A, C
defALL_40G_QSFPCURRENT_1	ALL_FCOE_40G_QSFP	CURRENT	<=1	NONE	M, A, C
defALL_40G_QSFPCURRENT_10	ALL_FCOE_40G_QSFP	CURRENT	>=10	NONE	—
defALL_40G_QSFPCURRENT_12	ALL_FCOE_40G_QSFP	CURRENT	>=12	NONE	M, A, C
defALL_40G_QSFPRXP_2188	ALL_FCOE_40G_QSFP	RXP	>=2188	NONE	M, A, C
defALL_40G_QSFPRXP_44	ALL_FCOE_40G_QSFP	RXP	<=44	NONE	M, A, C

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_40G_QSFPSFP_TEMP_75	ALL_FCOE_40G_QSFP	SFP_TEMP	>=75	NONE	M, A, C
defALL_40G_QSFPSFP_TEMP_n5	ALL_FCOE_40G_QSFP	SFP_TEMP	<=	NONE	M, A, C
defALL_40G_QSFPVOLTAGE_2970	ALL_FCOE_40G_QSFP	VOLTAGE	<=2970	NONE	M, A, C
defALL_40G_QSFPVOLTAGE_3630	ALL_FCOE_40G_QSFP	VOLTAGE	>=3630	NONE	M, A, C
defALL_40Km_16GELWL_SFPCURRENT_10	ALL_40Km_16GELWL_SFP	CURRENT	<=10	NONE	M, A, C
defALL_40Km_16GELWL_SFPCURRENT_130	ALL_40Km_16GELWL_SFP	CURRENT	>=130	NONE	M, A, C
defALL_40Km_16GELWL_SFPRXP_30	ALL_40Km_16GELWL_SFP	RXP	<=30	NONE	M, A, C
defALL_40Km_16GELWL_SFPRXP_794	ALL_40Km_16GELWL_SFP	RXP	>=794	NONE	M, A, C
defALL_40Km_16GELWL_SFPSFP_TEMP_75	ALL_40Km_16GELWL_SFP	SFP_TEMP	>=75	NONE	M, A, C
defALL_40Km_16GELWL_SFPSFP_TEMP_n5	ALL_40Km_16GELWL_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_40Km_16GELWL_SFPTXP_3162	ALL_40Km_16GELWL_SFP	TXP	>=3162	NONE	M, A, C
defALL_40Km_16GELWL_SFPTXP_794	ALL_40Km_16GELWL_SFP	TXP	<=794	NONE	M, A, C
defALL_40Km_16GELWL_SFPVOLTAGE_3000	ALL_40Km_16GELWL_SFP	VOLTAGE	<=3000	NONE	M, A, C
defALL_40Km_16GELWL_SFPVOLTAGE_3600	ALL_40Km_16GELWL_SFP	VOLTAGE	>=3600	NONE	M, A, C
defALL_40Km_32GELWL_SFPCURRENT_120	ALL_40Km_32GELWL_SFP	CURRENT	>=120	NONE	M, A, C
defALL_40Km_32GELWL_SFPCURRENT_30	ALL_40Km_32GELWL_SFP	CURRENT	<=30	NONE	M, A, C
defALL_40Km_32GELWL_SFPRXP_1995	ALL_40Km_32GELWL_SFP	RXP	>=1995	NONE	M, A, C
defALL_40Km_32GELWL_SFPRXP_40	ALL_40Km_32GELWL_SFP	RXP	<=40	NONE	M, A, C
defALL_40Km_32GELWL_SFPSFP_TEMP_80	ALL_40Km_32GELWL_SFP	SFP_TEMP	>=80	NONE	M, A, C
defALL_40Km_32GELWL_SFPSFP_TEMP_n10	ALL_40Km_32GELWL_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_40Km_32GELWL_SFPTXP_2512	ALL_40Km_32GELWL_SFP	TXP	>=2512	NONE	M, A, C
defALL_40Km_32GELWL_SFPTXP_320	ALL_40Km_32GELWL_SFP	TXP	<=320	NONE	M, A, C
defALL_40Km_32GELWL_SFPVOLTAGE_3000	ALL_40Km_32GELWL_SFP	VOLTAGE	<=3000	NONE	M, A, C

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_40Km_32GELWL_SFPVOLTAGE_3600	ALL_40Km_32GELWL_SFP	VOLTAGE	>=3600	NONE	M, A, C
defALL_64GLWL_SFPCURRENT_6	ALL_64GLWL_SFP	CURRENT	<=6	NONE	M, A, C
defALL_64GLWL_SFPCURRENT_90	ALL_64GLWL_SFP	CURRENT	>=90	NONE	M, A, C
defALL_64GLWL_SFPRXP_178	ALL_64GLWL_SFP	RXP	<=178	NONE	M, A, C
defALL_64GLWL_SFPRXP_5248	ALL_64GLWL_SFP	RXP	>=5248	NONE	M, A, C
defALL_64GLWL_SFPSFP_TEMP_85	ALL_64GLWL_SFP	SFP_TEMP	>=85	NONE	M, A, C
defALL_64GLWL_SFPSFP_TEMP_n5	ALL_64GLWL_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_64GLWL_SFPTXP_178	ALL_64GLWL_SFP	TXP	<=178	NONE	M, A, C
defALL_64GLWL_SFPTXP_5248	ALL_64GLWL_SFP	TXP	>=5248	NONE	M, A, C
defALL_64GLWL_SFPVOLTAGE_2970	ALL_64GLWL_SFP	VOLTAGE	<=2970	NONE	M, A, C
defALL_64GLWL_SFPVOLTAGE_3630	ALL_64GLWL_SFP	VOLTAGE	>=3630	NONE	M, A, C
defALL_64GSWL_SFPCURRENT_10	ALL_64GSWL_SFP	CURRENT	>=10	NONE	M, A, C
defALL_64GSWL_SFPCURRENT_2	ALL_64GSWL_SFP	CURRENT	<=2	NONE	M, A, C
defALL_64GSWL_SFPRXP_45	ALL_64GSWL_SFP	RXP	<=45	NONE	M, A, C
defALL_64GSWL_SFPRXP_5012	ALL_64GSWL_SFP	RXP	>=5012	NONE	M, A, C
defALL_64GSWL_SFPSFP_TEMP_80	ALL_64GSWL_SFP	SFP_TEMP	>=80	NONE	M, A, C
defALL_64GSWL_SFPSFP_TEMP_n5	ALL_64GSWL_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_64GSWL_SFPTXP_5012	ALL_64GSWL_SFP	TXP	>=5012	NONE	M, A, C
defALL_64GSWL_SFPTXP_83	ALL_64GSWL_SFP	TXP	<=83	NONE	M, A, C
defALL_64GSWL_SFPVOLTAGE_2970	ALL_64GSWL_SFP	VOLTAGE	<=2970	NONE	M, A, C
defALL_64GSWL_SFPVOLTAGE_3630	ALL_64GSWL_SFP	VOLTAGE	>=3630	NONE	M, A, C
defALL_64GSWL_SFP_DD_CURRENT_10	ALL_64GSWL_SFP_DD	CURRENT	>=10	NONE	M, A, C
defALL_64GSWL_SFP_DD_CURRENT_2	ALL_64GSWL_SFP_DD	CURRENT	<=2	NONE	M, A, C
defALL_64GSWL_SFP_DD_RXP_45	ALL_64GSWL_SFP_DD	RXP	<=45	NONE	M, A, C
defALL_64GSWL_SFP_DD_RXP_5012	ALL_64GSWL_SFP_DD	RXP	>=5012	NONE	M, A, C
defALL_64GSWL_SFP_DD_SFP_TEMP_80	ALL_64GSWL_SFP_DD	SFP_TEMP	>=80	NONE	M, A, C
defALL_64GSWL_SFP_DD_SFP_TEMP_n5	ALL_64GSWL_SFP_DD	SFP_TEMP	<=	NONE	M, A, C
defALL_64GSWL_SFP_DD_TXP_5012	ALL_64GSWL_SFP_DD	TXP	>=5012	NONE	M, A, C
defALL_64GSWL_SFP_DD_TXP_83	ALL_64GSWL_SFP_DD	TXP	<=83	NONE	M, A, C

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_64GSWL_SFP_DD_VOLTAGE_2970	ALL_64GSWL_SFP_DD	VOLTAGE	<=2970	NONE	M, A, C
defALL_64GSWL_SFP_DD_VOLTAGE_3630	ALL_64GSWL_SFP_DD	VOLTAGE	>=3630	NONE	M, A, C
defALL_80Km_8GELWL_SFPCURRENT_10	ALL_80Km_8GELWL_SFP	CURRENT	<=10	NONE	M, A, C
defALL_80Km_8GELWL_SFPCURRENT_130	ALL_80Km_8GELWL_SFP	CURRENT	>=130	NONE	M, A, C
defALL_80Km_8GELWL_SFPRXP_250	ALL_80Km_8GELWL_SFP	RXP	>=250	NONE	M, A, C
defALL_80Km_8GELWL_SFPRXP_3	ALL_80Km_8GELWL_SFP	RXP	<=3	NONE	M, A, C
defALL_80Km_8GELWL_SFPSFP_TEMP_75	ALL_80Km_8GELWL_SFP	SFP_TEMP	>=75	NONE	M, A, C
defALL_80Km_8GELWL_SFPSFP_TEMP_n5	ALL_80Km_8GELWL_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_80Km_8GELWL_SFPTXP_2510	ALL_80Km_8GELWL_SFP	TXP	>=2510	NONE	M, A, C
defALL_80Km_8GELWL_SFPTXP_630	ALL_80Km_8GELWL_SFP	TXP	<=630	NONE	M, A, C
defALL_80Km_8GELWL_SFPVOLTAGE_3000	ALL_80Km_8GELWL_SFP	VOLTAGE	<=3000	NONE	M, A, C
defALL_80Km_8GELWL_SFPVOLTAGE_3600	ALL_80Km_8GELWL_SFP	VOLTAGE	>=3600	NONE	M, A, C
defALL_BE_PORTSBAD_OS_5M_10	ALL_BE_PORTS	BAD_OS	>10	5MIN	M, A, C
defALL_BE_PORTSBAD_OS_D_100	ALL_BE_PORTS	BAD_OS	>100	DAY	M, A, C
defALL_BE_PORTSCRC_5M_10	ALL_BE_PORTS	CRC	>10	5MIN	M, A, C
defALL_BE_PORTSCRC_D_100	ALL_BE_PORTS	CRC	>100	DAY	M, A, C
defALL_BE_PORTSFRM_LONG_5M_10	ALL_BE_PORTS	FRM_LONG	>10	5MIN	M, A, C
defALL_BE_PORTSFRM_LONG_D_100	ALL_BE_PORTS	FRM_LONG	>100	DAY	M, A, C
defALL_BE_PORTSFRM_TRUNC_5M_10	ALL_BE_PORTS	FRM_TRUNC	>10	5MIN	M, A, C
defALL_BE_PORTSFRM_TRUNC_D_100	ALL_BE_PORTS	FRM_TRUNC	>100	DAY	M, A, C
defALL_BE_PORTSITW_5M_10	ALL_BE_PORTS	ITW	>10	5MIN	M, A, C
defALL_BE_PORTSITW_D_100	ALL_BE_PORTS	ITW	>100	DAY	M, A, C
defALL_BE_PORTSLR_5M_10	ALL_BE_PORTS	LR	>10	5MIN	M, A, C
defALL_BE_PORTSLR_D_100	ALL_BE_PORTS	LR	>100	DAY	M, A, C

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_BE_PORTS_LATENCY_CLEAR	ALL_BE_PORTS	BE_LATENCY_IMPACT	==IO_LATENCY_CLEAR	NONE	M, A, C
defALL_BE_PORTS_LATENCY_IMPACT	ALL_BE_PORTS	BE_LATENCY_IMPACT	==IO_PERF_IMPACT	NONE	M, A, C
defALL_CIRCUITSCIR_PKTLOSS_PER_05	ALL_CIRCUITS	CIR_PKTLOSS	>0	MIN	A
defALL_CIRCUITSCIR_PKTLOSS_PER_1	ALL_CIRCUITS	CIR_PKTLOSS	>0	MIN	M
defALL_CIRCUITSCIR_PKTLOSS_PER_5	ALL_CIRCUITS	CIR_PKTLOSS	>0	MIN	C
defALL_CIRCUITSCIR_STATE_0	ALL_CIRCUITS	CIR_STATE	>0	MIN	A
defALL_CIRCUITSCIR_STATE_3	ALL_CIRCUITS	CIR_STATE	>3	MIN	M
defALL_CIRCUITSCIR_STATE_5	ALL_CIRCUITS	CIR_STATE	>5	MIN	C
defALL_CIRCUITSCIR_UTIL_60	ALL_CIRCUITS	CIR_UTIL	>60	MIN	—
defALL_CIRCUITSCIR_UTIL_75	ALL_CIRCUITS	CIR_UTIL	>75	MIN	—
defALL_CIRCUITSCIR_UTIL_90	ALL_CIRCUITS	CIR_UTIL	>90	MIN	—
defALL_CIRCUITS_IP_JITTER_PER_05	ALL_CIRCUITS	IP_JITTER	>5	NONE	A
defALL_CIRCUITS_IP_JITTER_PER_15	ALL_CIRCUITS	IP_JITTER	>15	NONE	M
defALL_CIRCUITS_IP_JITTER_PER_20	ALL_CIRCUITS	IP_JITTER	>20	NONE	C
defALL_CIRCUITS_IP_PKTLOSS_P_05	ALL_CIRCUITS	CIR_IP_PKTLOSS	>0	MIN	A
defALL_CIRCUITS_IP_PKTLOSS_P_1	ALL_CIRCUITS	CIR_IP_PKTLOSS	>0	MIN	M
defALL_CIRCUITS_IP_PKTLOSS_P_5	ALL_CIRCUITS	CIR_IP_PKTLOSS	>0	MIN	C
defALL_CIRCUITS_IP_RTT_250	ALL_CIRCUITS	IP_RTT	>250	NONE	M, A, C
defALL_CIRCUITS_IP_UTIL_P_60	ALL_CIRCUITS	CIR_IP_UTIL	>60	MIN	—
defALL_CIRCUITS_IP_UTIL_P_75	ALL_CIRCUITS	CIR_IP_UTIL	>75	MIN	—
defALL_CIRCUITS_IP_UTIL_P_90	ALL_CIRCUITS	CIR_IP_UTIL	>90	MIN	—
defALL_CIRCUITS_JITTER_PER_05	ALL_CIRCUITS	JITTER	>5	NONE	A
defALL_CIRCUITS_JITTER_PER_15	ALL_CIRCUITS	JITTER	>15	NONE	M
defALL_CIRCUITS_JITTER_PER_20	ALL_CIRCUITS	JITTER	>20	NONE	C
defALL_CIRCUITS_RTT_250	ALL_CIRCUITS	RTT	>250	NONE	M, A, C
defALL_CIRCUIT_F_QOS_PKTLOSS_PER_05	ALL_CIRCUIT_F_QOS	CIR_QOS_PKTLOSS	>0	MIN	A
defALL_CIRCUIT_F_QOS_PKTLOSS_PER_1	ALL_CIRCUIT_F_QOS	CIR_QOS_PKTLOSS	>0	MIN	M

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_CIRCUIT_F_QOS_PKTLOSS_PER_5	ALL_CIRCUIT_F_QOS	CIR_QOS_PKTLOSS	>0	MIN	C
defALL_CIRCUIT_F_QOS_STATE_CHG_4	ALL_CIRCUIT_F_QOS	STATE_CHG	>4	MIN	—
defALL_CIRCUIT_F_QOS_UTIL_PER_50	ALL_CIRCUIT_F_QOS	CIR_QOS_UTIL	>50	HOUR	—
defALL_CIRCUIT_F_QOS_UTIL_PER_75	ALL_CIRCUIT_F_QOS	CIR_QOS_UTIL	>75	HOUR	—
defALL_CIRCUIT_F_QOS_UTIL_PER_90	ALL_CIRCUIT_F_QOS	CIR_QOS_UTIL	>90	HOUR	—
defALL_CIRCUIT_HIGH_QOS_PKTLOSS_PER_05	ALL_CIRCUIT_HIGH_QOS	CIR_QOS_PKTLOSS	>0	MIN	A
defALL_CIRCUIT_HIGH_QOS_PKTLOSS_PER_1	ALL_CIRCUIT_HIGH_QOS	CIR_QOS_PKTLOSS	>0	MIN	M
defALL_CIRCUIT_HIGH_QOS_PKTLOSS_PER_5	ALL_CIRCUIT_HIGH_QOS	CIR_QOS_PKTLOSS	>0	MIN	C
defALL_CIRCUIT_HIGH_QOS_STATE_CHG_4	ALL_CIRCUIT_HIGH_QOS	STATE_CHG	>4	MIN	—
defALL_CIRCUIT_HIGH_QOS_UTIL_PER_50	ALL_CIRCUIT_HIGH_QOS	CIR_QOS_UTIL	>50	HOUR	—
defALL_CIRCUIT_HIGH_QOS_UTIL_PER_75	ALL_CIRCUIT_HIGH_QOS	CIR_QOS_UTIL	>75	HOUR	—
defALL_CIRCUIT_HIGH_QOS_UTIL_PER_90	ALL_CIRCUIT_HIGH_QOS	CIR_QOS_UTIL	>90	HOUR	—
defALL_CIRCUIT_IP_HIGH_QOS_PKTLOSS_P_05	ALL_CIRCUIT_IP_HIGH_QOS	CIR_QOS_PKTLOSS	>0	MIN	A
defALL_CIRCUIT_IP_HIGH_QOS_PKTLOSS_P_1	ALL_CIRCUIT_IP_HIGH_QOS	CIR_QOS_PKTLOSS	>0	MIN	M
defALL_CIRCUIT_IP_HIGH_QOS_PKTLOSS_P_5	ALL_CIRCUIT_IP_HIGH_QOS	CIR_QOS_PKTLOSS	>0	MIN	C
defALL_CIRCUIT_IP_HIGH_QOS_STATE_CHG_4	ALL_CIRCUIT_IP_HIGH_QOS	STATE_CHG	>4	MIN	—
defALL_CIRCUIT_IP_HIGH_QOS_UTIL_P_50	ALL_CIRCUIT_IP_HIGH_QOS	CIR_QOS_UTIL	>50	HOUR	—
defALL_CIRCUIT_IP_HIGH_QOS_UTIL_P_75	ALL_CIRCUIT_IP_HIGH_QOS	CIR_QOS_UTIL	>75	HOUR	—
defALL_CIRCUIT_IP_HIGH_QOS_UTIL_P_90	ALL_CIRCUIT_IP_HIGH_QOS	CIR_QOS_UTIL	>90	HOUR	—
defALL_CIRCUIT_IP_LOW_QOS_PKTLOSS_P_05	ALL_CIRCUIT_IP_LOW_QOS	CIR_QOS_PKTLOSS	>0	MIN	A
defALL_CIRCUIT_IP_LOW_QOS_PKTLOSS_P_1	ALL_CIRCUIT_IP_LOW_QOS	CIR_QOS_PKTLOSS	>0	MIN	M

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_CIRCUIT_IP_LOW_QOS_PKTLOSS_P_5	ALL_CIRCUIT_IP_LOW_QOS	CIR_QOS_PKTLOSS	>0	MIN	C
defALL_CIRCUIT_IP_LOW_QOS_STATE_CHG_4	ALL_CIRCUIT_IP_LOW_QOS	STATE_CHG	>4	MIN	—
defALL_CIRCUIT_IP_LOW_QOS_UTIL_P_50	ALL_CIRCUIT_IP_LOW_QOS	CIR_QOS_UTIL	>50	HOUR	—
defALL_CIRCUIT_IP_LOW_QOS_UTIL_P_75	ALL_CIRCUIT_IP_LOW_QOS	CIR_QOS_UTIL	>75	HOUR	—
defALL_CIRCUIT_IP_LOW_QOS_UTIL_P_90	ALL_CIRCUIT_IP_LOW_QOS	CIR_QOS_UTIL	>90	HOUR	—
defALL_CIRCUIT_IP_MED_QOS_PKTLOSS_P_05	ALL_CIRCUIT_IP_MED_QOS	CIR_QOS_PKTLOSS	>0	MIN	A
defALL_CIRCUIT_IP_MED_QOS_PKTLOSS_P_1	ALL_CIRCUIT_IP_MED_QOS	CIR_QOS_PKTLOSS	>0	MIN	M
defALL_CIRCUIT_IP_MED_QOS_PKTLOSS_P_5	ALL_CIRCUIT_IP_MED_QOS	CIR_QOS_PKTLOSS	>0	MIN	C
defALL_CIRCUIT_IP_MED_QOS_STATE_CHG_4	ALL_CIRCUIT_IP_MED_QOS	STATE_CHG	>4	MIN	—
defALL_CIRCUIT_IP_MED_QOS_UTIL_P_50	ALL_CIRCUIT_IP_MED_QOS	CIR_QOS_UTIL	>50	HOUR	—
defALL_CIRCUIT_IP_MED_QOS_UTIL_P_75	ALL_CIRCUIT_IP_MED_QOS	CIR_QOS_UTIL	>75	HOUR	—
defALL_CIRCUIT_IP_MED_QOS_UTIL_P_90	ALL_CIRCUIT_IP_MED_QOS	CIR_QOS_UTIL	>90	HOUR	—
defALL_CIRCUIT_LOW_QOS_PKTLOSS_PER_05	ALL_CIRCUIT_LOW_QOS	CIR_QOS_PKTLOSS	>0	MIN	A
defALL_CIRCUIT_LOW_QOS_PKTLOSS_PER_1	ALL_CIRCUIT_LOW_QOS	CIR_QOS_PKTLOSS	>0	MIN	M
defALL_CIRCUIT_LOW_QOS_PKTLOSS_PER_5	ALL_CIRCUIT_LOW_QOS	CIR_QOS_PKTLOSS	>0	MIN	C
defALL_CIRCUIT_LOW_QOS_STATE_CHG_4	ALL_CIRCUIT_LOW_QOS	STATE_CHG	>4	MIN	—
defALL_CIRCUIT_LOW_QOS_UTIL_PER_50	ALL_CIRCUIT_LOW_QOS	CIR_QOS_UTIL	>50	HOUR	—
defALL_CIRCUIT_LOW_QOS_UTIL_PER_75	ALL_CIRCUIT_LOW_QOS	CIR_QOS_UTIL	>75	HOUR	—
defALL_CIRCUIT_LOW_QOS_UTIL_PER_90	ALL_CIRCUIT_LOW_QOS	CIR_QOS_UTIL	>90	HOUR	—
defALL_CIRCUIT_MED_QOS_PKTLOSS_PER_05	ALL_CIRCUIT_MED_QOS	CIR_QOS_PKTLOSS	>0	MIN	A
defALL_CIRCUIT_MED_QOS_PKTLOSS_PER_1	ALL_CIRCUIT_MED_QOS	CIR_QOS_PKTLOSS	>0	MIN	M

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_CIRCUIT_MED_QOS_PKTLOSS_PER_5	ALL_CIRCUIT_MED_QOS	CIR_QOS_PKTLOSS	>0	MIN	C
defALL_CIRCUIT_MED_QOS_STATE_CHG_4	ALL_CIRCUIT_MED_QOS	STATE_CHG	>4	MIN	—
defALL_CIRCUIT_MED_QOS_UTIL_PER_50	ALL_CIRCUIT_MED_QOS	CIR_QOS_UTIL	>50	HOUR	—
defALL_CIRCUIT_MED_QOS_UTIL_PER_75	ALL_CIRCUIT_MED_QOS	CIR_QOS_UTIL	>75	HOUR	—
defALL_CIRCUIT_MED_QOS_UTIL_PER_90	ALL_CIRCUIT_MED_QOS	CIR_QOS_UTIL	>90	HOUR	—
defALL_DPIP_EXTN_FLOW_A	ALL_DP	IP_EXTN_FLOW	>65	none	A
defALL_DPIP_EXTN_FLOW_C	ALL_DP	IP_EXTN_FLOW	>90	none	C
defALL_DPIP_EXTN_FLOW_M	ALL_DP	IP_EXTN_FLOW	>80	none	M
defALL_DPIP_EXTN_FLOW_MAX	ALL_DP	IP_EXTN_FLOW	>99	none	C, M, A
defALL_D_PORTSCRC_1	ALL_D_PORTS	CRC	>1	MIN	A
defALL_D_PORTSCRC_2	ALL_D_PORTS	CRC	>2	MIN	M
defALL_D_PORTSCRC_3	ALL_D_PORTS	CRC	>3	MIN	C
defALL_D_PORTSCRC_D1000	ALL_D_PORTS	CRC	>1000	DAY	M
defALL_D_PORTSCRC_D1500	ALL_D_PORTS	CRC	>1500	DAY	C
defALL_D_PORTSCRC_D500	ALL_D_PORTS	CRC	>500	DAY	A
defALL_D_PORTSCRC_H30	ALL_D_PORTS	CRC	>30	HOUR	A
defALL_D_PORTSCRC_H60	ALL_D_PORTS	CRC	>60	HOUR	M
defALL_D_PORTSCRC_H90	ALL_D_PORTS	CRC	>90	HOUR	C
defALL_D_PORTSITW_1	ALL_D_PORTS	ITW	>1	MIN	A
defALL_D_PORTSITW_2	ALL_D_PORTS	ITW	>2	MIN	M
defALL_D_PORTSITW_3	ALL_D_PORTS	ITW	>3	MIN	C
defALL_D_PORTSITW_D1000	ALL_D_PORTS	ITW	>1000	DAY	M
defALL_D_PORTSITW_D1500	ALL_D_PORTS	ITW	>1500	DAY	C
defALL_D_PORTSITW_D500	ALL_D_PORTS	ITW	>500	DAY	A
defALL_D_PORTSITW_H30	ALL_D_PORTS	ITW	>30	HOUR	A
defALL_D_PORTSITW_H60	ALL_D_PORTS	ITW	>60	HOUR	M
defALL_D_PORTSITW_H90	ALL_D_PORTS	ITW	>90	HOUR	C
defALL_D_PORTSLOSS_SYNC_1	ALL_D_PORTS	LOSS_SYNC	>1	MIN	A
defALL_D_PORTSLOSS_SYNC_2	ALL_D_PORTS	LOSS_SYNC	>2	MIN	M
defALL_D_PORTSLOSS_SYNC_3	ALL_D_PORTS	LOSS_SYNC	>3	MIN	C
defALL_D_PORTSLOSS_SYNC_D1000	ALL_D_PORTS	LOSS_SYNC	>1000	DAY	M
defALL_D_PORTSLOSS_SYNC_D1500	ALL_D_PORTS	LOSS_SYNC	>1500	DAY	C

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_D_PORTSLOSS_SYNC_D500	ALL_D_PORTS	LOSS_SYNC	>500	DAY	A
defALL_D_PORTSLOSS_SYNC_H30	ALL_D_PORTS	LOSS_SYNC	>30	HOUR	A
defALL_D_PORTSLOSS_SYNC_H60	ALL_D_PORTS	LOSS_SYNC	>60	HOUR	M
defALL_D_PORTSLOSS_SYNC_H90	ALL_D_PORTS	LOSS_SYNC	>90	HOUR	C
defALL_ETH_PORTSSFP_STATE_FAULTY	ALL_ETH_PORTS	SFP_STATE	==FAULTY	NONE	M, A, C
defALL_ETH_PORTSSFP_STATE_IN	ALL_ETH_PORTS	SFP_STATE	==IN	NONE	M, A, C
defALL_ETH_PORTSSFP_STATE_OUT	ALL_ETH_PORTS	SFP_STATE	==OUT	NONE	M, A, C
defALL_EXT_GE_PORTSCRC_0	ALL_EXT_GE_PORTS	GE_CRC	>0	MIN	A
defALL_EXT_GE_PORTSCRC_1	ALL_EXT_GE_PORTS	GE_CRC	>1	MIN	M, C
defALL_EXT_GE_PORTSLOS_0	ALL_EXT_GE_PORTS	GE_LOS_OF_SIG	>0	MIN	A
defALL_EXT_GE_PORTSLOS_1	ALL_EXT_GE_PORTS	GE_LOS_OF_SIG	>1	MIN	M, C
defALL_E_PORTSC3TXTO_10	ALL_E_PORTS	C3TXTO	>10	MIN	M
defALL_E_PORTSC3TXTO_20	ALL_E_PORTS	C3TXTO	>20	MIN	C
defALL_E_PORTSC3TXTO_5	ALL_E_PORTS	C3TXTO	>5	MIN	A
defALL_E_PORTSCRC_0	ALL_E_PORTS	CRC	>0	MIN	A
defALL_E_PORTSCRC_10	ALL_E_PORTS	CRC	>10	MIN	M
defALL_E_PORTSCRC_2	ALL_E_PORTS	CRC	>2	MIN	A
defALL_E_PORTSCRC_20	ALL_E_PORTS	CRC	>20	MIN	M
defALL_E_PORTSCRC_21	ALL_E_PORTS	CRC	>21	MIN	C
defALL_E_PORTSCRC_40	ALL_E_PORTS	CRC	>40	MIN	C
defALL_E_PORTSENCR_BLK	ALL_E_PORTS	ENCR_BLK	>0	MIN	M, A, C
defALL_E_PORTSENCR_DISC	ALL_E_PORTS	ENCR_DISC	>30	NONE	M, A, C
defALL_E_PORTSENCR_SHORT_FRM	ALL_E_PORTS	ENCR_SHRT_FRM	>0	MIN	M, A, C
defALL_E_PORTSITW_15	ALL_E_PORTS	ITW	>15	MIN	A
defALL_E_PORTSITW_20	ALL_E_PORTS	ITW	>20	MIN	A
defALL_E_PORTSITW_21	ALL_E_PORTS	ITW	>21	MIN	M
defALL_E_PORTSITW_40	ALL_E_PORTS	ITW	>40	MIN	M
defALL_E_PORTSITW_41	ALL_E_PORTS	ITW	>41	MIN	C
defALL_E_PORTSITW_80	ALL_E_PORTS	ITW	>80	MIN	C
defALL_E_PORTSLOSS_SYNC_0	ALL_E_PORTS	LOSS_SYNC	>0	MIN	A
defALL_E_PORTSLOSS_SYNC_3	ALL_E_PORTS	LOSS_SYNC	>3	MIN	M
defALL_E_PORTSLOSS_SYNC_5	ALL_E_PORTS	LOSS_SYNC	>5	MIN	C
defALL_E_PORTSLR_10	ALL_E_PORTS	LR	>10	MIN	M
defALL_E_PORTSLR_11	ALL_E_PORTS	LR	>11	MIN	C

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_E_PORTSLR_2	ALL_E_PORTS	LR	>2	MIN	A
defALL_E_PORTSLR_20	ALL_E_PORTS	LR	>20	MIN	C
defALL_E_PORTSLR_4	ALL_E_PORTS	LR	>4	MIN	A
defALL_E_PORTSLR_5	ALL_E_PORTS	LR	>5	MIN	M
defALL_E_PORTSPE_0	ALL_E_PORTS	PE	>0	MIN	A
defALL_E_PORTSPE_10	ALL_E_PORTS	PE	>10	MIN	C
defALL_E_PORTSPE_2	ALL_E_PORTS	PE	>2	MIN	A
defALL_E_PORTSPE_3	ALL_E_PORTS	PE	>3	MIN	M
defALL_E_PORTSPE_5	ALL_E_PORTS	PE	>5	MIN	C
defALL_E_PORTSPE_7	ALL_E_PORTS	PE	>7	MIN	M
defALL_E_PORTSROUTING_ERR_1	ALL_E_PORTS	ROUTING_ERR	>1	SEC	A, M, C, B
defALL_E_PORTSROUTING_ERR_ROR_2	ALL_E_PORTS	defALL_E_PORTSROUTING_ERR_1	>2	MIN	A, M, C, B
defALL_E_PORTSRX_60	ALL_E_PORTS	RX	>60	HOUR	A
defALL_E_PORTSRX_75	ALL_E_PORTS	RX	>75	HOUR	M
defALL_E_PORTSRX_90	ALL_E_PORTS	RX	>90	HOUR	C, B
defALL_E_PORTSRX_95	ALL_E_PORTS	RX	>95	MIN	C, B, M, A
defALL_E_PORTSSTATE_CHG_10	ALL_E_PORTS	STATE_CHG	>10	MIN	M
defALL_E_PORTSSTATE_CHG_11	ALL_E_PORTS	STATE_CHG	>11	MIN	C
defALL_E_PORTSSTATE_CHG_2	ALL_E_PORTS	STATE_CHG	>2	MIN	A
defALL_E_PORTSSTATE_CHG_20	ALL_E_PORTS	STATE_CHG	>20	MIN	C
defALL_E_PORTSSTATE_CHG_4	ALL_E_PORTS	STATE_CHG	>4	MIN	A
defALL_E_PORTSSTATE_CHG_5	ALL_E_PORTS	STATE_CHG	>5	MIN	M
defALL_E_PORTSTX_60	ALL_E_PORTS	TX	>60	HOUR	A
defALL_E_PORTSTX_75	ALL_E_PORTS	TX	>75	HOUR	M
defALL_E_PORTSTX_90	ALL_E_PORTS	TX	>90	HOUR	C, B
defALL_E_PORTSTX_95	ALL_E_PORTS	TX	>95	MIN	C, M, A, B
defALL_E_PORTSUTIL_60	ALL_E_PORTS	UTIL	>60	HOUR	A
defALL_E_PORTSUTIL_75	ALL_E_PORTS	UTIL	>75	HOUR	M
defALL_E_PORTSUTIL_90	ALL_E_PORTS	UTIL	>90	HOUR	C, B
defALL_E_PORTSUTIL_95	ALL_E_PORTS	UTIL	>95	MIN	C, M, A, B
defALL_FANFAN_STATE_FAULTY	ALL_FAN	FAN_STATE	==FAULTY	NONE	M, A, C, B
defALL_FANFAN_STATE_ON	ALL_FAN	FAN_STATE	==ON	NONE	M, A, C, B
defALL_FANFAN_STATE_OUT	ALL_FAN	FAN_STATE	==OUT	NONE	M, A, C, B
defALL_FAN_AIR_FLOW_MISMATCH	CHASSIS	FAN_AIRFLOW_MISMATCH	==TRUE	NONE	M, A, C, B

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_F_PORTSDEV_LOGIN_DIST_BALANCED	ALL_F_PORTS	DEV_LOGIN_DIST	==BALANCED	NONE	C, M, A
defALL_F_PORTSDEV_LOGIN_DIST_BALANCE_FAILED	ALL_F_PORTS	DEV_LOGIN_DIST	==BALANCE_FAILED	NONE	C, M, A
defALL_F_PORTSDEV_LOGIN_DIST_IMBALANCED	ALL_F_PORTS	DEV_LOGIN_DIST	==IMBALANCED	NONE	C, M, A
defALL_F_PORTSDEV_NPIV_LOGINS_PER_60	ALL_F_PORTS	DEV_NPIV_LOGINS	>60	NONE	A
defALL_F_PORTSDEV_NPIV_LOGINS_PER_75	ALL_F_PORTS	DEV_NPIV_LOGINS	>75	NONE	M
defALL_F_PORTSDEV_NPIV_LOGINS_PER_90	ALL_F_PORTS	DEV_NPIV_LOGINS	>90	NONE	C
defALL_F_PORTSROUTING_ERR_1	ALL_F_PORTS	ROUTING_ERR	>1	SEC	A, M, C, B
defALL_F_PORTSROUTING_ERR_ROR_2	ALL_F_PORTS	defALL_F_PORTSROUTING_ERR_1	>2	MIN	A, M, C, B
defALL_F_PORTSRX_90	ALL_F_PORTS	RX	>90	HOUR	B
defALL_F_PORTSRX_95	ALL_F_PORTS	RX	>95	MIN	B
defALL_F_PORTSTX_90	ALL_F_PORTS	TX	>90	HOUR	B
defALL_F_PORTSTX_95	ALL_F_PORTS	TX	>95	MIN	B
defALL_F_PORTSUTIL_90	ALL_F_PORTS	UTIL	>90	HOUR	B
defALL_F_PORTSUTIL_95	ALL_F_PORTS	UTIL	>95	MIN	B
defALL_GEN7SWL_QSFPCURRENT_5	ALL_GEN7SWL_QSFP	CURRENT	<=5	NONE	M, A, C
defALL_GEN7SWL_QSFPCURRENT_9	ALL_GEN7SWL_QSFP	CURRENT	>=9	NONE	M, A, C
defALL_GEN7SWL_QSFPRXP_3467	ALL_GEN7SWL_QSFP	RXP	>=3467	NONE	M, A, C
defALL_GEN7SWL_QSFPRXP_46	ALL_GEN7SWL_QSFP	RXP	<=46	NONE	M, A, C
defALL_GEN7SWL_QSFPSFP_TEMP_80	ALL_GEN7SWL_QSFP	SFP_TEMP	>=80	NONE	M, A, C
defALL_GEN7SWL_QSFPSFP_TEMP_n10	ALL_GEN7SWL_QSFP	SFP_TEMP	<=	NONE	M, A, C
defALL_GEN7SWL_QSFPTXP_3467	ALL_GEN7SWL_QSFP	TXP	>=3467	NONE	M, A, C
defALL_GEN7SWL_QSFPTXP_72	ALL_GEN7SWL_QSFP	TXP	<=72	NONE	M, A, C
defALL_GEN7SWL_QSFPVOLTAGE_3100	ALL_GEN7SWL_QSFP	VOLTAGE	<=3100	NONE	M, A, C
defALL_GEN7SWL_QSFPVOLTAGE_3500	ALL_GEN7SWL_QSFP	VOLTAGE	>=3500	NONE	M, A, C
defALL_HOST_PORTSC3TXTO_10	ALL_HOST_PORTS	C3TXTO	>10	MIN	M
defALL_HOST_PORTSC3TXTO_11	ALL_HOST_PORTS	C3TXTO	>11	MIN	C
defALL_HOST_PORTSC3TXTO_2	ALL_HOST_PORTS	C3TXTO	>2	MIN	A

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_HOST_PORTSC3TXTO_20	ALL_HOST_PORTS	C3TXTO	>20	MIN	C
defALL_HOST_PORTSC3TXTO_3	ALL_HOST_PORTS	C3TXTO	>3	MIN	M
defALL_HOST_PORTSC3TXTO_4	ALL_HOST_PORTS	C3TXTO	>4	MIN	A
defALL_HOST_PORTSCRC_0	ALL_HOST_PORTS	CRC	>0	MIN	A
defALL_HOST_PORTSCRC_10	ALL_HOST_PORTS	CRC	>10	MIN	M
defALL_HOST_PORTSCRC_2	ALL_HOST_PORTS	CRC	>2	MIN	A
defALL_HOST_PORTSCRC_20	ALL_HOST_PORTS	CRC	>20	MIN	M
defALL_HOST_PORTSCRC_21	ALL_HOST_PORTS	CRC	>21	MIN	C
defALL_HOST_PORTSCRC_40	ALL_HOST_PORTS	CRC	>40	MIN	C
defALL_HOST_PORTSITW_15	ALL_HOST_PORTS	ITW	>15	MIN	A
defALL_HOST_PORTSITW_20	ALL_HOST_PORTS	ITW	>20	MIN	A
defALL_HOST_PORTSITW_21	ALL_HOST_PORTS	ITW	>21	MIN	M
defALL_HOST_PORTSITW_40	ALL_HOST_PORTS	ITW	>40	MIN	M
defALL_HOST_PORTSITW_41	ALL_HOST_PORTS	ITW	>41	MIN	C
defALL_HOST_PORTSITW_80	ALL_HOST_PORTS	ITW	>80	MIN	C
defALL_HOST_PORTSLOSS_SYNC_0	ALL_HOST_PORTS	LOSS_SYNC	>0	MIN	A
defALL_HOST_PORTSLOSS_SYNC_3	ALL_HOST_PORTS	LOSS_SYNC	>3	MIN	M
defALL_HOST_PORTSLOSS_SYNC_5	ALL_HOST_PORTS	LOSS_SYNC	>5	MIN	C
defALL_HOST_PORTSLR_10	ALL_HOST_PORTS	LR	>10	MIN	M
defALL_HOST_PORTSLR_11	ALL_HOST_PORTS	LR	>11	MIN	C
defALL_HOST_PORTSLR_2	ALL_HOST_PORTS	LR	>2	MIN	A
defALL_HOST_PORTSLR_20	ALL_HOST_PORTS	LR	>20	MIN	C
defALL_HOST_PORTSLR_4	ALL_HOST_PORTS	LR	>4	MIN	A
defALL_HOST_PORTSLR_5	ALL_HOST_PORTS	LR	>5	MIN	M
defALL_HOST_PORTSPE_0	ALL_HOST_PORTS	PE	>0	MIN	A
defALL_HOST_PORTSPE_10	ALL_HOST_PORTS	PE	>10	MIN	C
defALL_HOST_PORTSPE_2	ALL_HOST_PORTS	PE	>2	MIN	A
defALL_HOST_PORTSPE_3	ALL_HOST_PORTS	PE	>3	MIN	M
defALL_HOST_PORTSPE_5	ALL_HOST_PORTS	PE	>5	MIN	C
defALL_HOST_PORTSPE_7	ALL_HOST_PORTS	PE	>7	MIN	M
defALL_HOST_PORTSRX_60	ALL_HOST_PORTS	RX	>60	HOUR	A
defALL_HOST_PORTSRX_75	ALL_HOST_PORTS	RX	>75	HOUR	M
defALL_HOST_PORTSRX_90	ALL_HOST_PORTS	RX	>90	HOUR	C
defALL_HOST_PORTSRX_95	ALL_HOST_PORTS	RX	>95	MIN	C, M, A

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_HOST_PORTSSTATE_CHG_10	ALL_HOST_PORTS	STATE_CHG	>10	MIN	M
defALL_HOST_PORTSSTATE_CHG_11	ALL_HOST_PORTS	STATE_CHG	>11	MIN	C
defALL_HOST_PORTSSTATE_CHG_2	ALL_HOST_PORTS	STATE_CHG	>2	MIN	A
defALL_HOST_PORTSSTATE_CHG_20	ALL_HOST_PORTS	STATE_CHG	>20	MIN	C
defALL_HOST_PORTSSTATE_CHG_4	ALL_HOST_PORTS	STATE_CHG	>4	MIN	A
defALL_HOST_PORTSSTATE_CHG_5	ALL_HOST_PORTS	STATE_CHG	>5	MIN	M
defALL_HOST_PORTSTX_60	ALL_HOST_PORTS	TX	>60	HOUR	A
defALL_HOST_PORTSTX_75	ALL_HOST_PORTS	TX	>75	HOUR	M
defALL_HOST_PORTSTX_90	ALL_HOST_PORTS	TX	>90	HOUR	C
defALL_HOST_PORTSTX_95	ALL_HOST_PORTS	TX	>95	MIN	C, M, A
defALL_HOST_PORTSUTIL_60	ALL_HOST_PORTS	UTIL	>60	HOUR	A
defALL_HOST_PORTSUTIL_75	ALL_HOST_PORTS	UTIL	>75	HOUR	M
defALL_HOST_PORTSUTIL_90	ALL_HOST_PORTS	UTIL	>90	HOUR	C
defALL_HOST_PORTSUTIL_95	ALL_HOST_PORTS	UTIL	>95	MIN	C, M, A
defALL_LOCAL_PIDSIT_FLOW_16	ALL_LOCAL_PIDS	IT_FLOW	>16	NONE	M
defALL_LOCAL_PIDSIT_FLOW_32	ALL_LOCAL_PIDS	IT_FLOW	>32	NONE	C
defALL_LOCAL_PIDSIT_FLOW_8	ALL_LOCAL_PIDS	IT_FLOW	>8	NONE	A
defALL_OTHER_F_PORTSC3TXTO_10	ALL_OTHER_F_PORTS	C3TXTO	>10	MIN	M
defALL_OTHER_F_PORTSC3TXTO_11	ALL_OTHER_F_PORTS	C3TXTO	>11	MIN	C
defALL_OTHER_F_PORTSC3TXTO_2	ALL_OTHER_F_PORTS	C3TXTO	>2	MIN	A
defALL_OTHER_F_PORTSC3TXTO_20	ALL_OTHER_F_PORTS	C3TXTO	>20	MIN	C
defALL_OTHER_F_PORTSC3TXTO_3	ALL_OTHER_F_PORTS	C3TXTO	>3	MIN	M
defALL_OTHER_F_PORTSC3TXTO_4	ALL_OTHER_F_PORTS	C3TXTO	>4	MIN	A
defALL_OTHER_F_PORTSCRC_0	ALL_OTHER_F_PORTS	CRC	>0	MIN	A
defALL_OTHER_F_PORTSCRC_10	ALL_OTHER_F_PORTS	CRC	>10	MIN	M
defALL_OTHER_F_PORTSCRC_2	ALL_OTHER_F_PORTS	CRC	>2	MIN	A
defALL_OTHER_F_PORTSCRC_20	ALL_OTHER_F_PORTS	CRC	>20	MIN	M
defALL_OTHER_F_PORTSCRC_21	ALL_OTHER_F_PORTS	CRC	>21	MIN	C
defALL_OTHER_F_PORTSCRC_40	ALL_OTHER_F_PORTS	CRC	>40	MIN	C
defALL_OTHER_F_PORTSITW_15	ALL_OTHER_F_PORTS	ITW	>15	MIN	A

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_OTHER_F_PORTSITW_20	ALL_OTHER_F_PORTS	ITW	>20	MIN	A
defALL_OTHER_F_PORTSITW_21	ALL_OTHER_F_PORTS	ITW	>21	MIN	M
defALL_OTHER_F_PORTSITW_40	ALL_OTHER_F_PORTS	ITW	>40	MIN	M
defALL_OTHER_F_PORTSITW_41	ALL_OTHER_F_PORTS	ITW	>41	MIN	C
defALL_OTHER_F_PORTSITW_80	ALL_OTHER_F_PORTS	ITW	>80	MIN	C
defALL_OTHER_F_PORTSLOSS_SYNC_0	ALL_OTHER_F_PORTS	LOSS_SYNC	>0	MIN	A
defALL_OTHER_F_PORTSLOSS_SYNC_3	ALL_OTHER_F_PORTS	LOSS_SYNC	>3	MIN	M
defALL_OTHER_F_PORTSLOSS_SYNC_5	ALL_OTHER_F_PORTS	LOSS_SYNC	>5	MIN	C
defALL_OTHER_F_PORTSLR_10	ALL_OTHER_F_PORTS	LR	>10	MIN	M
defALL_OTHER_F_PORTSLR_11	ALL_OTHER_F_PORTS	LR	>11	MIN	C
defALL_OTHER_F_PORTSLR_2	ALL_OTHER_F_PORTS	LR	>2	MIN	A
defALL_OTHER_F_PORTSLR_20	ALL_OTHER_F_PORTS	LR	>20	MIN	C
defALL_OTHER_F_PORTSLR_4	ALL_OTHER_F_PORTS	LR	>4	MIN	A
defALL_OTHER_F_PORTSLR_5	ALL_OTHER_F_PORTS	LR	>5	MIN	M
defALL_OTHER_F_PORTSPE_0	ALL_OTHER_F_PORTS	PE	>0	MIN	A
defALL_OTHER_F_PORTSPE_10	ALL_OTHER_F_PORTS	PE	>10	MIN	C
defALL_OTHER_F_PORTSPE_2	ALL_OTHER_F_PORTS	PE	>2	MIN	A
defALL_OTHER_F_PORTSPE_3	ALL_OTHER_F_PORTS	PE	>3	MIN	M
defALL_OTHER_F_PORTSPE_5	ALL_OTHER_F_PORTS	PE	>5	MIN	C
defALL_OTHER_F_PORTSPE_7	ALL_OTHER_F_PORTS	PE	>7	MIN	M
defALL_OTHER_F_PORTSRX_60	ALL_OTHER_F_PORTS	RX	>60	HOUR	A
defALL_OTHER_F_PORTSRX_75	ALL_OTHER_F_PORTS	RX	>75	HOUR	M
defALL_OTHER_F_PORTSRX_90	ALL_OTHER_F_PORTS	RX	>90	HOUR	C
defALL_OTHER_F_PORTSRX_95	ALL_OTHER_F_PORTS	RX	>95	MIN	C, M, A
defALL_OTHER_F_PORTSSTATE_CHG_10	ALL_OTHER_F_PORTS	STATE_CHG	>10	MIN	M
defALL_OTHER_F_PORTSSTATE_CHG_11	ALL_OTHER_F_PORTS	STATE_CHG	>11	MIN	C
defALL_OTHER_F_PORTSSTATE_CHG_2	ALL_OTHER_F_PORTS	STATE_CHG	>2	MIN	A
defALL_OTHER_F_PORTSSTATE_CHG_20	ALL_OTHER_F_PORTS	STATE_CHG	>20	MIN	C
defALL_OTHER_F_PORTSSTATE_CHG_4	ALL_OTHER_F_PORTS	STATE_CHG	>4	MIN	A
defALL_OTHER_F_PORTSSTATE_CHG_5	ALL_OTHER_F_PORTS	STATE_CHG	>5	MIN	M

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_OTHER_F_PORTSTX_60	ALL_OTHER_F_PORTS	TX	>60	HOUR	A
defALL_OTHER_F_PORTSTX_75	ALL_OTHER_F_PORTS	TX	>75	HOUR	M
defALL_OTHER_F_PORTSTX_90	ALL_OTHER_F_PORTS	TX	>90	HOUR	C
defALL_OTHER_F_PORTSTX_95	ALL_OTHER_F_PORTS	TX	>95	MIN	C, M, A
defALL_OTHER_F_PORTSUTIL_60	ALL_OTHER_F_PORTS	UTIL	>60	HOUR	A
defALL_OTHER_F_PORTSUTIL_75	ALL_OTHER_F_PORTS	UTIL	>75	HOUR	M
defALL_OTHER_F_PORTSUTIL_90	ALL_OTHER_F_PORTS	UTIL	>90	HOUR	C
defALL_OTHER_F_PORTSUTIL_95	ALL_OTHER_F_PORTS	UTIL	>95	MIN	C, M, A
defALL_OTHER_SFPCURRENT_50	ALL_OTHER_SFP	CURRENT	>=50	NONE	M, A, C
defALL_OTHER_SFPRXP_5000	ALL_OTHER_SFP	RXP	>=5000	NONE	M, A, C
defALL_OTHER_SFPSFP_TEMP_85	ALL_OTHER_SFP	SFP_TEMP	>=85	NONE	M, A, C
defALL_OTHER_SFPSFP_TEMP_n13	ALL_OTHER_SFP	SFP_TEMP	<=	NONE	M, A, C
defALL_OTHER_SFPTXP_5000	ALL_OTHER_SFP	TXP	>=5000	NONE	M, A, C
defALL_OTHER_SFPVOLTAGE_2960	ALL_OTHER_SFP	VOLTAGE	<=2960	NONE	M, A, C
defALL_OTHER_SFPVOLTAGE_3630	ALL_OTHER_SFP	VOLTAGE	>=3630	NONE	M, A, C
defALL_PORTSLF_0	ALL_PORTS	LF	>0	MIN	A
defALL_PORTSLF_3	ALL_PORTS	LF	>3	MIN	M
defALL_PORTSLF_5	ALL_PORTS	LF	>5	MIN	C
defALL_PORTSLOSS_SIGNAL_0	ALL_PORTS	LOSS_SIGNAL	>0	MIN	A
defALL_PORTSLOSS_SIGNAL_3	ALL_PORTS	LOSS_SIGNAL	>3	MIN	M
defALL_PORTSLOSS_SIGNAL_5	ALL_PORTS	LOSS_SIGNAL	>5	MIN	C
defALL_PORTSSFP_STATE_FAULTY	ALL_PORTS	SFP_STATE	==FAULTY	NONE	M, A, C, B
defALL_PORTSSFP_STATE_IN	ALL_PORTS	SFP_STATE	==IN	NONE	M, A, C, B
defALL_PORTSSFP_STATE_OUT	ALL_PORTS	SFP_STATE	==OUT	NONE	M, A, C, B
defALL_PORTS_IO_FRAME_LOSS	ALL_PORTS	DEV_LATENCY_IMPACT	==IO_FRAME_LOSS	NONE	A, B
defALL_PORTS_IO_FRAME_LOSS_UNQUAR	ALL_PORTS	DEV_LATENCY_IMPACT	==IO_FRAME_LOSS	NONE	M, C
defALL_PORTS_IO_LATENCY_CLEAR	ALL_PORTS	DEV_LATENCY_IMPACT	==IO_LATENCY_CLEAR	NONE	M, A, C, B
defALL_PORTS_IO_PERF_IMPACT	ALL_PORTS	DEV_LATENCY_IMPACT	==IO_PERF_IMPACT	NONE	A, B
defALL_PORTS_IO_PERF_IMPACT_UNQUAR	ALL_PORTS	DEV_LATENCY_IMPACT	==IO_PERF_IMPACT	NONE	M, C
defALL_PORTS_OVERSUBSCRIBED	ALL_PORTS	PORT_BANDWIDTH	==OVERSUBSCRIBED	NONE	M, A, C, B

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_PORTS_OVERSUBSCRIPTION_CLEAR	ALL_PORTS	PORT_BANDWIDTH	==OVERSUBSCRIPTION_CLEAR	NONE	M, A, C, B
defALL_PSPS_STATE_FAULTY	ALL_PS	PS_STATE	==FAULTY	NONE	M, A, C, B
defALL_PSPS_STATE_ON	ALL_PS	PS_STATE	==ON	NONE	M, A, C, B
defALL_PSPS_STATE_OUT	ALL_PS	PS_STATE	==OUT	NONE	M, A, C, B
defALL_SLOTSBLADE_STATE_FAULTY	ALL_SLOTS	BLADE_STATE	==FAULTY	NONE	M, A, C, B
defALL_SLOTSBLADE_STATE_OFF	ALL_SLOTS	BLADE_STATE	==OFF	NONE	M, A, C, B
defALL_SLOTSBLADE_STATE_ON	ALL_SLOTS	BLADE_STATE	==ON	NONE	M, A, C, B
defALL_SLOTSBLADE_STATE_OUT	ALL_SLOTS	BLADE_STATE	==OUT	NONE	M, A, C, B
defALL_TARGET_PORTSC3TXTO_0	ALL_TARGET_PORTS	C3TXTO	>0	MIN	A
defALL_TARGET_PORTSC3TXTO_10	ALL_TARGET_PORTS	C3TXTO	>10	MIN	C
defALL_TARGET_PORTSC3TXTO_2	ALL_TARGET_PORTS	C3TXTO	>2	MIN	A
defALL_TARGET_PORTSC3TXTO_3	ALL_TARGET_PORTS	C3TXTO	>3	MIN	M
defALL_TARGET_PORTSC3TXTO_5	ALL_TARGET_PORTS	C3TXTO	>5	MIN	M
defALL_TARGET_PORTSC3TXTO_6	ALL_TARGET_PORTS	C3TXTO	>6	MIN	C
defALL_TARGET_PORTSCRC_0	ALL_TARGET_PORTS	CRC	>0	MIN	A
defALL_TARGET_PORTSCRC_10	ALL_TARGET_PORTS	CRC	>10	MIN	M
defALL_TARGET_PORTSCRC_11	ALL_TARGET_PORTS	CRC	>11	MIN	C
defALL_TARGET_PORTSCRC_2	ALL_TARGET_PORTS	CRC	>2	MIN	A
defALL_TARGET_PORTSCRC_20	ALL_TARGET_PORTS	CRC	>20	MIN	C
defALL_TARGET_PORTSCRC_5	ALL_TARGET_PORTS	CRC	>5	MIN	M
defALL_TARGET_PORTSITW_10	ALL_TARGET_PORTS	ITW	>10	MIN	A
defALL_TARGET_PORTSITW_11	ALL_TARGET_PORTS	ITW	>11	MIN	M
defALL_TARGET_PORTSITW_20	ALL_TARGET_PORTS	ITW	>20	MIN	M
defALL_TARGET_PORTSITW_21	ALL_TARGET_PORTS	ITW	>21	MIN	C
defALL_TARGET_PORTSITW_40	ALL_TARGET_PORTS	ITW	>40	MIN	C
defALL_TARGET_PORTSITW_5	ALL_TARGET_PORTS	ITW	>5	MIN	A
defALL_TARGET_PORTSLOSS_SYNC_0	ALL_TARGET_PORTS	LOSS_SYNC	>0	MIN	A
defALL_TARGET_PORTSLOSS_SYNC_3	ALL_TARGET_PORTS	LOSS_SYNC	>3	MIN	M
defALL_TARGET_PORTSLOSS_SYNC_5	ALL_TARGET_PORTS	LOSS_SYNC	>5	MIN	C
defALL_TARGET_PORTSLR_0	ALL_TARGET_PORTS	LR	>0	MIN	A
defALL_TARGET_PORTSLR_10	ALL_TARGET_PORTS	LR	>10	MIN	C
defALL_TARGET_PORTSLR_2	ALL_TARGET_PORTS	LR	>2	MIN	A

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_TARGET_PORTSLR_3	ALL_TARGET_PORTS	LR	>3	MIN	M
defALL_TARGET_PORTSLR_5	ALL_TARGET_PORTS	LR	>5	MIN	M
defALL_TARGET_PORTSLR_6	ALL_TARGET_PORTS	LR	>6	MIN	C
defALL_TARGET_PORTSPE_0	ALL_TARGET_PORTS	PE	>0	MIN	A
defALL_TARGET_PORTSPE_2	ALL_TARGET_PORTS	PE	>2	MIN	A
defALL_TARGET_PORTSPE_3	ALL_TARGET_PORTS	PE	>3	MIN	M
defALL_TARGET_PORTSPE_4	ALL_TARGET_PORTS	PE	>4	MIN	M
defALL_TARGET_PORTSPE_5	ALL_TARGET_PORTS	PE	>5	MIN	C
defALL_TARGET_PORTSPE_6	ALL_TARGET_PORTS	PE	>6	MIN	C
defALL_TARGET_PORTSRX_60	ALL_TARGET_PORTS	RX	>60	HOUR	A
defALL_TARGET_PORTSRX_75	ALL_TARGET_PORTS	RX	>75	HOUR	M
defALL_TARGET_PORTSRX_90	ALL_TARGET_PORTS	RX	>90	HOUR	C
defALL_TARGET_PORTSRX_95	ALL_TARGET_PORTS	RX	>95	MIN	C, M, A
defALL_TARGET_PORTSSTATE_CHG_0	ALL_TARGET_PORTS	STATE_CHG	>0	MIN	A
defALL_TARGET_PORTSSTATE_CHG_15	ALL_TARGET_PORTS	STATE_CHG	>15	MIN	C
defALL_TARGET_PORTSSTATE_CHG_2	ALL_TARGET_PORTS	STATE_CHG	>2	MIN	A
defALL_TARGET_PORTSSTATE_CHG_3	ALL_TARGET_PORTS	STATE_CHG	>3	MIN	M
defALL_TARGET_PORTSSTATE_CHG_7	ALL_TARGET_PORTS	STATE_CHG	>7	MIN	M
defALL_TARGET_PORTSSTATE_CHG_8	ALL_TARGET_PORTS	STATE_CHG	>8	MIN	C
defALL_TARGET_PORTSTX_60	ALL_TARGET_PORTS	TX	>60	HOUR	A
defALL_TARGET_PORTSTX_75	ALL_TARGET_PORTS	TX	>75	HOUR	M
defALL_TARGET_PORTSTX_90	ALL_TARGET_PORTS	TX	>90	HOUR	C
defALL_TARGET_PORTSTX_95	ALL_TARGET_PORTS	TX	>95	MIN	C, M, A
defALL_TARGET_PORTSUTIL_60	ALL_TARGET_PORTS	UTIL	>60	HOUR	A
defALL_TARGET_PORTSUTIL_75	ALL_TARGET_PORTS	UTIL	>75	HOUR	M
defALL_TARGET_PORTSUTIL_90	ALL_TARGET_PORTS	UTIL	>90	HOUR	C
defALL_TARGET_PORTSUTIL_95	ALL_TARGET_PORTS	UTIL	>95	MIN	C, M, A
defALL_TSTEMP_IN_RANGE	ALL_TS	TEMP	==IN_RANGE	NONE	M, A, C, B
defALL_TSTEMP_OUT_OF_RANGE	ALL_TS	TEMP	==OUT_OF_RANGE	NONE	M, A, C, B
defALL_TUNNELSSTATE_CHG_0	ALL_TUNNELS	TUNNEL_STATE	>0	MIN	A
defALL_TUNNELSSTATE_CHG_1	ALL_TUNNELS	TUNNEL_STATE	>1	MIN	M

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_TUNNELSSTATE_CHG_3	ALL_TUNNELS	TUNNEL_STATE	>3	MIN	C
defALL_TUNNELSUTIL_PER_50	ALL_TUNNELS	TUNNEL_UTIL	>50	HOUR	—
defALL_TUNNELSUTIL_PER_75	ALL_TUNNELS	TUNNEL_UTIL	>75	HOUR	—
defALL_TUNNELSUTIL_PER_90	ALL_TUNNELS	TUNNEL_UTIL	>90	HOUR	—
defALL_TUNNELS_IP_UTIL_P_50	ALL_TUNNELS	TUNNEL_IP_UTIL	>50	HOUR	—
defALL_TUNNELS_IP_UTIL_P_75	ALL_TUNNELS	TUNNEL_IP_UTIL	>75	HOUR	—
defALL_TUNNELS_IP_UTIL_P_90	ALL_TUNNELS	TUNNEL_IP_UTIL	>90	HOUR	—
defALL_TUNNEL_F_QOS_PKTLOSS_PER_05	ALL_TUNNEL_F_QOS	PKTLOSS	>0	MIN	A
defALL_TUNNEL_F_QOS_PKTLOSS_PER_1	ALL_TUNNEL_F_QOS	PKTLOSS	>0	MIN	M
defALL_TUNNEL_F_QOS_PKTLOSS_PER_5	ALL_TUNNEL_F_QOS	PKTLOSS	>0	MIN	C
defALL_TUNNEL_F_QOS_UTIL_PER_50	ALL_TUNNEL_F_QOS	QOS_UTIL	>50	HOUR	—
defALL_TUNNEL_F_QOS_UTIL_PER_75	ALL_TUNNEL_F_QOS	QOS_UTIL	>75	HOUR	—
defALL_TUNNEL_F_QOS_UTIL_PER_90	ALL_TUNNEL_F_QOS	QOS_UTIL	>90	HOUR	—
defALL_TUNNEL_HIGH_QOS_PKTLOSS_PER_05	ALL_TUNNEL_HIGH_QOS	PKTLOSS	>0	MIN	A
defALL_TUNNEL_HIGH_QOS_PKTLOSS_PER_1	ALL_TUNNEL_HIGH_QOS	PKTLOSS	>0	MIN	M
defALL_TUNNEL_HIGH_QOS_PKTLOSS_PER_5	ALL_TUNNEL_HIGH_QOS	PKTLOSS	>0	MIN	C
defALL_TUNNEL_HIGH_QOS_UTIL_PER_50	ALL_TUNNEL_HIGH_QOS	QOS_UTIL	>50	HOUR	—
defALL_TUNNEL_HIGH_QOS_UTIL_PER_75	ALL_TUNNEL_HIGH_QOS	QOS_UTIL	>75	HOUR	—
defALL_TUNNEL_HIGH_QOS_UTIL_PER_90	ALL_TUNNEL_HIGH_QOS	QOS_UTIL	>90	HOUR	—
defALL_TUNNEL_IP_HIGH_QOS_PKTLOSS_P_05	ALL_TUNNEL_IP_HIGH_QOS	PKTLOSS	>0	MIN	A
defALL_TUNNEL_IP_HIGH_QOS_PKTLOSS_P_1	ALL_TUNNEL_IP_HIGH_QOS	PKTLOSS	>0	MIN	M
defALL_TUNNEL_IP_HIGH_QOS_PKTLOSS_P_5	ALL_TUNNEL_IP_HIGH_QOS	PKTLOSS	>0	MIN	C
defALL_TUNNEL_IP_HIGH_QOS_UTIL_P_50	ALL_TUNNEL_IP_HIGH_QOS	QOS_UTIL	>50	HOUR	—
defALL_TUNNEL_IP_HIGH_QOS_UTIL_P_75	ALL_TUNNEL_IP_HIGH_QOS	QOS_UTIL	>75	HOUR	—

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_TUNNEL_IP_HIGH_QOS_UTIL_P_90	ALL_TUNNEL_IP_HIGH_QOS	QOS_UTIL	>90	HOUR	—
defALL_TUNNEL_IP_LOW_QOS_PKTLOSS_P_05	ALL_TUNNEL_IP_LOW_QOS	PKTLOSS	>0	MIN	A
defALL_TUNNEL_IP_LOW_QOS_PKTLOSS_P_1	ALL_TUNNEL_IP_LOW_QOS	PKTLOSS	>0	MIN	M
defALL_TUNNEL_IP_LOW_QOS_PKTLOSS_P_5	ALL_TUNNEL_IP_LOW_QOS	PKTLOSS	>0	MIN	C
defALL_TUNNEL_IP_LOW_QOS_UTIL_P_50	ALL_TUNNEL_IP_LOW_QOS	QOS_UTIL	>50	HOUR	—
defALL_TUNNEL_IP_LOW_QOS_UTIL_P_75	ALL_TUNNEL_IP_LOW_QOS	QOS_UTIL	>75	HOUR	—
defALL_TUNNEL_IP_LOW_QOS_UTIL_P_90	ALL_TUNNEL_IP_LOW_QOS	QOS_UTIL	>90	HOUR	—
defALL_TUNNEL_IP_MED_QOS_PKTLOSS_P_05	ALL_TUNNEL_IP_MED_QOS	PKTLOSS	>0	MIN	A
defALL_TUNNEL_IP_MED_QOS_PKTLOSS_P_1	ALL_TUNNEL_IP_MED_QOS	PKTLOSS	>0	MIN	M
defALL_TUNNEL_IP_MED_QOS_PKTLOSS_P_5	ALL_TUNNEL_IP_MED_QOS	PKTLOSS	>0	MIN	C
defALL_TUNNEL_IP_MED_QOS_UTIL_P_50	ALL_TUNNEL_IP_MED_QOS	QOS_UTIL	>50	HOUR	—
defALL_TUNNEL_IP_MED_QOS_UTIL_P_75	ALL_TUNNEL_IP_MED_QOS	QOS_UTIL	>75	HOUR	—
defALL_TUNNEL_IP_MED_QOS_UTIL_P_90	ALL_TUNNEL_IP_MED_QOS	QOS_UTIL	>90	HOUR	—
defALL_TUNNEL_LOW_QOS_PKTLOSS_PER_05	ALL_TUNNEL_LOW_QOS	PKTLOSS	>0	MIN	A
defALL_TUNNEL_LOW_QOS_PKTLOSS_PER_1	ALL_TUNNEL_LOW_QOS	PKTLOSS	>0	MIN	M
defALL_TUNNEL_LOW_QOS_PKTLOSS_PER_5	ALL_TUNNEL_LOW_QOS	PKTLOSS	>0	MIN	C
defALL_TUNNEL_LOW_QOS_UTIL_PER_50	ALL_TUNNEL_LOW_QOS	QOS_UTIL	>50	HOUR	—
defALL_TUNNEL_LOW_QOS_UTIL_PER_75	ALL_TUNNEL_LOW_QOS	QOS_UTIL	>75	HOUR	—
defALL_TUNNEL_LOW_QOS_UTIL_PER_90	ALL_TUNNEL_LOW_QOS	QOS_UTIL	>90	HOUR	—
defALL_TUNNEL_MED_QOS_PKTLOSS_PER_05	ALL_TUNNEL_MED_QOS	PKTLOSS	>0	MIN	A
defALL_TUNNEL_MED_QOS_PKTLOSS_PER_1	ALL_TUNNEL_MED_QOS	PKTLOSS	>0	MIN	M

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defALL_TUNNEL_MED_QOS_PKTLOSS_PER_5	ALL_TUNNEL_MED_QOS	PKTLOSS	>0	MIN	C
defALL_TUNNEL_MED_QOS_UTIL_PER_50	ALL_TUNNEL_MED_QOS	QOS_UTIL	>50	HOUR	—
defALL_TUNNEL_MED_QOS_UTIL_PER_75	ALL_TUNNEL_MED_QOS	QOS_UTIL	>75	HOUR	—
defALL_TUNNEL_MED_QOS_UTIL_PER_90	ALL_TUNNEL_MED_QOS	QOS_UTIL	>90	HOUR	—
defALL_WWNWWN_FAULTY	ALL_WWN	WWN	==FAULTY	NONE	M, A, C, B
defALL_WWNWWN_ON	ALL_WWN	WWN	==ON	NONE	M, A, C, B
defALL_WWNWWN_OUT	ALL_WWN	WWN	==OUT	NONE	M, A, C, B
defCHASSISBAD_FAN_CRIT	CHASSIS	BAD_FAN	>=2	NONE	M, A, C, B
defCHASSISBAD_FAN_MARG	CHASSIS	BAD_FAN	>=1	NONE	M, A, C, B
defCHASSISBAD_PWR_CRIT	CHASSIS	BAD_PWR	>=3	NONE	M, A, C, B
defCHASSISBAD_PWR_MARG	CHASSIS	BAD_PWR	>=2	NONE	M, A, C, B
defCHASSISBAD_TEMP_CRIT	CHASSIS	BAD_TEMP	>=2	NONE	M, A, C, B
defCHASSISBAD_TEMP_CRIT	CHASSIS	BAD_TEMP	>=2	NONE	M, A, C, B
defCHASSISBAD_TEMP_MARG	CHASSIS	BAD_TEMP	>=1	NONE	M, A, C, B
defCHASSISBAD_TEMP_MARG	CHASSIS	BAD_TEMP	>=1	NONE	M, A, C, B
defCHASSISCERTS_EXPIRED	CHASSIS	EXPIRED_CERTS	>0	NONE	M, A, C, B
defCHASSISCERT_VALIDITY_15	ALL_CERTS	DAYS_TO_EXPIRE	<15	NONE	C
defCHASSISCERT_VALIDITY_20	ALL_CERTS	DAYS_TO_EXPIRE	<20	NONE	M
defCHASSISCERT_VALIDITY_30	ALL_CERTS	DAYS_TO_EXPIRE	<30	NONE	A
defCHASSISCPU_80	CHASSIS	CPU	>=80	NONE	M, A, C, B
defCHASSISDOWN_CORE_1	CHASSIS	DOWN_CORE	>=1	NONE	M, A, C, B
defCHASSISDOWN_CORE_2	CHASSIS	DOWN_CORE	>=2	NONE	M, A, C, B
defCHASSISETH_MGMT_PORT_STATE_DOWN	CHASSIS	ETH_MGMT_PORT_STATE	==DOWN	NONE	M, A, C, B
defCHASSISETH_MGMT_PORT_STATE_UP	CHASSIS	ETH_MGMT_PORT_STATE	==UP	NONE	M, A, C, B
defCHASSISFAULTY_BLADE_1	CHASSIS	FAULTY_BLADE	>=1	NONE	M, A, C, B
defCHASSISFLASH_USAGE_90	CHASSIS	FLASH_USAGE	>=90	NONE	M, A, C, B
defCHASSISHA_SYNC_0	CHASSIS	HA_SYNC	==0	NONE	M, A, C, B
defCHASSISMEMORY_USAGE_75	CHASSIS	MEMORY_USAGE	>=75	NONE	—
defCHASSISMEMORY_USAGE_STATE_CRIT	CHASSIS	MEMORY_USAGE_STATE	==CRITICAL	NONE	S

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defCHASSISMEMORY_USAGE_STATE_IN_RANGE	CHASSIS	MEMORY_USAGE_STATE	==IN_RANGE	NONE	S
defCHASSISMEMORY_USAGE_STATE_WARN	CHASSIS	MEMORY_USAGE_STATE	==WARNING	NONE	S
defCHASSISSYSTEM_TEMP_CRIT	CHASSIS	SYSTEM_TEMP	==CRIT_OUT_OF_RANGE	NONE	M, A, C, B
defCHASSISSYSTEM_TEMP_MARG	CHASSIS	SYSTEM_TEMP	==MARG_OUT_OF_RANGE	NONE	M, A, C, B
defCHASSISTRUFOS_CERT_DAYS_TO_EXPIRE_60	CHASSIS	TRUFOS_CERT_DAYS_TO_EXPIRE	<60	NONE	S
defCHASSISTRUFOS_CERT_EXPIRED	CHASSIS	TRUFOS_CERT_EXPIRED	==TRUE	NONE	S
defCHASSISTRUFOS_CERT_EXPIRED_MARG	CHASSIS	TRUFOS_CERT_EXPIRED	==TRUE	NONE	—
defCHASSISTRUFOS_CERT_INSTALLED	CHASSIS	TRUFOS_CERT_INSTALLED	==FALSE	NONE	S
defCHASSISWWN_DOWN_1	CHASSIS	WWN_DOWN	>=1	NONE	M, A, C, B
defCHASSISZONE_CFGSZ_PER_70	CHASSIS	ZONE_CFGSZ_PER	>70	NONE	A
defCHASSISZONE_CFGSZ_PER_80	CHASSIS	ZONE_CFGSZ_PER	>80	NONE	M
defCHASSISZONE_CFGSZ_PER_90	CHASSIS	ZONE_CFGSZ_PER	>90	NONE	C
defIO_ABORT	sys_flow_monitor	IO_ABORT	>100	10SEC	A, C, M
defIO_TIMEOUT	sys_flow_monitor	IO_TIMEOUT	>100	10SEC	A, C, M
defITL_RES_USAGE_P_60	CHASSIS	ITL_RES_USAGE	>60	NONE	A
defITL_RES_USAGE_P_75	CHASSIS	ITL_RES_USAGE	>75	NONE	M
defITL_RES_USAGE_P_90	CHASSIS	ITL_RES_USAGE	>90	NONE	C
defIT_RES_USAGE_P_60	CHASSIS	IT_RES_USAGE	>60	NONE	A
defIT_RES_USAGE_P_75	CHASSIS	IT_RES_USAGE	>75	NONE	M
defIT_RES_USAGE_P_90	CHASSIS	IT_RES_USAGE	>90	NONE	C
defMAX_RD_PENDING_IO_100	sys_flow_monitor_nvme	MAX_RD_PENDING_IO	>100	10SEC	A
defMAX_RD_PENDING_IO_125	sys_flow_monitor_nvme	MAX_RD_PENDING_IO	>125	10SEC	M
defMAX_RD_PENDING_IO_150	sys_flow_monitor_nvme	MAX_RD_PENDING_IO	>150	10SEC	C
defMAX_RD_PENDING_IO_200	sys_flow_monitor_scsi	MAX_RD_PENDING_IO	>200	10SEC	A

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defMAX_RD_PENDING_IO_250	sys_flow_monitor_scsi	MAX_RD_PENDING_IO	>250	10SEC	M
defMAX_RD_PENDING_IO_300	sys_flow_monitor_scsi	MAX_RD_PENDING_IO	>300	10SEC	C
defMAX_WR_PENDING_IO_100	sys_flow_monitor_nvme	MAX_WR_PENDING_IO	>100	10SEC	A
defMAX_WR_PENDING_IO_125	sys_flow_monitor_nvme	MAX_WR_PENDING_IO	>125	10SEC	M
defMAX_WR_PENDING_IO_150	sys_flow_monitor_nvme	MAX_WR_PENDING_IO	>150	10SEC	C
defMAX_WR_PENDING_IO_200	sys_flow_monitor_scsi	MAX_WR_PENDING_IO	>200	10SEC	A
defMAX_WR_PENDING_IO_250	sys_flow_monitor_scsi	MAX_WR_PENDING_IO	>250	10SEC	M
defMAX_WR_PENDING_IO_300	sys_flow_monitor_scsi	MAX_WR_PENDING_IO	>300	10SEC	C
defNON_E_F_PORTSCRC_0	NON_E_F_PORTS	CRC	>0	MIN	A
defNON_E_F_PORTSCRC_10	NON_E_F_PORTS	CRC	>10	MIN	M
defNON_E_F_PORTSCRC_2	NON_E_F_PORTS	CRC	>2	MIN	A
defNON_E_F_PORTSCRC_20	NON_E_F_PORTS	CRC	>20	MIN	M
defNON_E_F_PORTSCRC_21	NON_E_F_PORTS	CRC	>21	MIN	C
defNON_E_F_PORTSCRC_40	NON_E_F_PORTS	CRC	>40	MIN	C
defNON_E_F_PORTSITW_15	NON_E_F_PORTS	ITW	>15	MIN	A
defNON_E_F_PORTSITW_20	NON_E_F_PORTS	ITW	>20	MIN	A
defNON_E_F_PORTSITW_21	NON_E_F_PORTS	ITW	>21	MIN	M
defNON_E_F_PORTSITW_40	NON_E_F_PORTS	ITW	>40	MIN	M
defNON_E_F_PORTSITW_41	NON_E_F_PORTS	ITW	>41	MIN	C
defNON_E_F_PORTSITW_80	NON_E_F_PORTS	ITW	>80	MIN	C
defNON_E_F_PORTSLOSS_SYNC_0	NON_E_F_PORTS	LOSS_SYNC	>0	MIN	A
defNON_E_F_PORTSLOSS_SYNC_3	NON_E_F_PORTS	LOSS_SYNC	>3	MIN	M
defNON_E_F_PORTSLOSS_SYNC_5	NON_E_F_PORTS	LOSS_SYNC	>5	MIN	C
defNON_E_F_PORTSLR_10	NON_E_F_PORTS	LR	>10	MIN	M
defNON_E_F_PORTSLR_11	NON_E_F_PORTS	LR	>11	MIN	C
defNON_E_F_PORTSLR_2	NON_E_F_PORTS	LR	>2	MIN	A
defNON_E_F_PORTSLR_20	NON_E_F_PORTS	LR	>20	MIN	C
defNON_E_F_PORTSLR_4	NON_E_F_PORTS	LR	>4	MIN	A
defNON_E_F_PORTSLR_5	NON_E_F_PORTS	LR	>5	MIN	M
defNON_E_F_PORTSPE_0	NON_E_F_PORTS	PE	>0	MIN	A
defNON_E_F_PORTSPE_10	NON_E_F_PORTS	PE	>10	MIN	C

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defNON_E_F_PORTSPE_2	NON_E_F_PORTS	PE	>2	MIN	A
defNON_E_F_PORTSPE_3	NON_E_F_PORTS	PE	>3	MIN	M
defNON_E_F_PORTSPE_5	NON_E_F_PORTS	PE	>5	MIN	C
defNON_E_F_PORTSPE_7	NON_E_F_PORTS	PE	>7	MIN	M
defNON_E_F_PORTSRX_60	NON_E_F_PORTS	RX	>60	HOUR	A
defNON_E_F_PORTSRX_75	NON_E_F_PORTS	RX	>75	HOUR	M
defNON_E_F_PORTSRX_90	NON_E_F_PORTS	RX	>90	HOUR	C
defNON_E_F_PORTSRX_95	NON_E_F_PORTS	RX	>95	MIN	C, M, A
defNON_E_F_PORTSSTATE_CHG_10	NON_E_F_PORTS	STATE_CHG	>10	MIN	M
defNON_E_F_PORTSSTATE_CHG_11	NON_E_F_PORTS	STATE_CHG	>11	MIN	C
defNON_E_F_PORTSSTATE_CHG_2	NON_E_F_PORTS	STATE_CHG	>2	MIN	A
defNON_E_F_PORTSSTATE_CHG_20	NON_E_F_PORTS	STATE_CHG	>20	MIN	C
defNON_E_F_PORTSSTATE_CHG_4	NON_E_F_PORTS	STATE_CHG	>4	MIN	A
defNON_E_F_PORTSSTATE_CHG_5	NON_E_F_PORTS	STATE_CHG	>5	MIN	M
defNON_E_F_PORTSTX_60	NON_E_F_PORTS	TX	>60	HOUR	A
defNON_E_F_PORTSTX_75	NON_E_F_PORTS	TX	>75	HOUR	M
defNON_E_F_PORTSTX_90	NON_E_F_PORTS	TX	>90	HOUR	C
defNON_E_F_PORTSTX_95	NON_E_F_PORTS	TX	>95	MIN	C, M, A
defNON_E_F_PORTSUTIL_60	NON_E_F_PORTS	UTIL	>60	HOUR	A
defNON_E_F_PORTSUTIL_75	NON_E_F_PORTS	UTIL	>75	HOUR	M
defNON_E_F_PORTSUTIL_90	NON_E_F_PORTS	UTIL	>90	HOUR	C
defNON_E_F_PORTSUTIL_95	NON_E_F_PORTS	UTIL	>95	MIN	C, M, A
defPORT_MAX_RD_PENDING_IO_200	ALL_TARGET_PORTS	PORT_MAX_RD_PENDING_IO	>200	10SEC	A
defPORT_MAX_RD_PENDING_IO_250	ALL_TARGET_PORTS	PORT_MAX_RD_PENDING_IO	>250	10SEC	M
defPORT_MAX_RD_PENDING_IO_300	ALL_TARGET_PORTS	PORT_MAX_RD_PENDING_IO	>300	10SEC	C
defPORT_MAX_WR_PENDING_IO_200	ALL_TARGET_PORTS	PORT_MAX_WR_PENDING_IO	>200	10SEC	A
defPORT_MAX_WR_PENDING_IO_250	ALL_TARGET_PORTS	PORT_MAX_WR_PENDING_IO	>250	10SEC	M
defPORT_MAX_WR_PENDING_IO_300	ALL_TARGET_PORTS	PORT_MAX_WR_PENDING_IO	>300	10SEC	C
defRD_1stDATA_TIME_11000	sys_flow_monitor_scsi	RD_1stDATA_TIME	>11000	10SEC	M

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defRD_1stDATA_TIME_1750	sys_flow_monitor_nvme	RD_1stDATA_TIME	>1750	10SEC	A
defRD_1stDATA_TIME_3500	sys_flow_monitor_scsi	RD_1stDATA_TIME	>3500	10SEC	A
defRD_1stDATA_TIME_45000	sys_flow_monitor_nvme	RD_1stDATA_TIME	>45000	10SEC	C
defRD_1stDATA_TIME_5500	sys_flow_monitor_nvme	RD_1stDATA_TIME	>5500	10SEC	M
defRD_1stDATA_TIME_90000	sys_flow_monitor_scsi	RD_1stDATA_TIME	>90000	10SEC	C
defRD_1stDATA_TIME_IO_11000	sys_flow_monitor_scsi	RD_1stDATA_TIME	>11000	IO	M
defRD_1stDATA_TIME_IO_1750	sys_flow_monitor_nvme	RD_1stDATA_TIME	>1750	IO	A
defRD_1stDATA_TIME_IO_3500	sys_flow_monitor_scsi	RD_1stDATA_TIME	>3500	IO	A
defRD_1stDATA_TIME_IO_45000	sys_flow_monitor_nvme	RD_1stDATA_TIME	>45000	IO	C
defRD_1stDATA_TIME_IO_5500	sys_flow_monitor_nvme	RD_1stDATA_TIME	>5500	IO	M
defRD_1stDATA_TIME_IO_90000	sys_flow_monitor_scsi	RD_1stDATA_TIME	>90000	IO	C
defRD_1stDATA_TIME_VIOL_10	ALL_F_PORTS	RD_1stDATA_TIME_VIOL	>10	10SEC	M
defRD_1stDATA_TIME_VIOL_20	ALL_F_PORTS	RD_1stDATA_TIME_VIOL	>20	10SEC	C
defRD_1stDATA_TIME_VIOL_5	ALL_F_PORTS	RD_1stDATA_TIME_VIOL	>5	10SEC	A
defRD_PENDING_IO_100	sys_flow_monitor_nvme	RD_PENDING_IO	>100	10SEC	A
defRD_PENDING_IO_125	sys_flow_monitor_nvme	RD_PENDING_IO	>125	10SEC	M
defRD_PENDING_IO_150	sys_flow_monitor_nvme	RD_PENDING_IO	>150	10SEC	C
defRD_PENDING_IO_200	sys_flow_monitor_scsi	RD_PENDING_IO	>200	10SEC	A
defRD_PENDING_IO_250	sys_flow_monitor_scsi	RD_PENDING_IO	>250	10SEC	M
defRD_PENDING_IO_300	sys_flow_monitor_scsi	RD_PENDING_IO	>300	10SEC	C
defRD_STATUS_TIME_100000	sys_flow_monitor_scsi	RD_STATUS_TIME	>100000	10SEC	C
defRD_STATUS_TIME_12000	sys_flow_monitor_scsi	RD_STATUS_TIME	>12000	10SEC	M
defRD_STATUS_TIME_2000	sys_flow_monitor_nvme	RD_STATUS_TIME	>2000	10SEC	A
defRD_STATUS_TIME_4000	sys_flow_monitor_scsi	RD_STATUS_TIME	>4000	10SEC	A

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defRD_STATUS_TIME_50000	sys_flow_monitor_nvme	RD_STATUS_TIME	>50000	10SEC	C
defRD_STATUS_TIME_6000	sys_flow_monitor_nvme	RD_STATUS_TIME	>6000	10SEC	M
defRD_STATUS_TIME_IO_100000	sys_flow_monitor_scsi	RD_STATUS_TIME	>100000	IO	C
defRD_STATUS_TIME_IO_12000	sys_flow_monitor_scsi	RD_STATUS_TIME	>12000	IO	M
defRD_STATUS_TIME_IO_2000	sys_flow_monitor_nvme	RD_STATUS_TIME	>2000	IO	A
defRD_STATUS_TIME_IO_4000	sys_flow_monitor_scsi	RD_STATUS_TIME	>4000	IO	A
defRD_STATUS_TIME_IO_50000	sys_flow_monitor_nvme	RD_STATUS_TIME	>50000	IO	C
defRD_STATUS_TIME_IO_6000	sys_flow_monitor_nvme	RD_STATUS_TIME	>6000	IO	M
defRD_STATUS_TIME_VIOL_10	ALL_F_PORTS	RD_STATUS_TIME_VIOL	>10	10SEC	M
defRD_STATUS_TIME_VIOL_20	ALL_F_PORTS	RD_STATUS_TIME_VIOL	>20	10SEC	C
defRD_STATUS_TIME_VIOL_5	ALL_F_PORTS	RD_STATUS_TIME_VIOL	>5	10SEC	A
defSWITCHASC_UPLOAD_FAILURE_1	CHASSIS	ASC_UPLOAD_FAILURE	>=1	WEEK	A
defSWITCHASC_UPLOAD_FAILURE_2	CHASSIS	ASC_UPLOAD_FAILURE	>=2	WEEK	M
defSWITCHASC_UPLOAD_FAILURE_3	CHASSIS	ASC_UPLOAD_FAILURE	>=3	WEEK	B, C
defSWITCHBB_FCR_CNT_MAX	SWITCH	BB_FCR_CNT	>16	NONE	M, A, C
defSWITCHDID_CHG_1	SWITCH	DID_CHG	>1	MIN	M, A, C
defSWITCHEPORT_DOWN_1	SWITCH	EPORT_DOWN	>1	MIN	A
defSWITCHEPORT_DOWN_2	SWITCH	EPORT_DOWN	>2	MIN	M
defSWITCHEPORT_DOWN_4	SWITCH	EPORT_DOWN	>4	MIN	C
defSWITCHERR_PORTS_P_10	SWITCH	ERR_PORTS	>=10	NONE	M
defSWITCHERR_PORTS_P_11	SWITCH	ERR_PORTS	>=11	NONE	C
defSWITCHERR_PORTS_P_25	SWITCH	ERR_PORTS	>=25	NONE	C
defSWITCHERR_PORTS_P_5	SWITCH	ERR_PORTS	>=5	NONE	A
defSWITCHERR_PORTS_P_6	SWITCH	ERR_PORTS	>=6	NONE	M
defSWITCHFAB_CFG_1	SWITCH	FAB_CFG	>1	MIN	A
defSWITCHFAB_CFG_2	SWITCH	FAB_CFG	>2	MIN	M

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defSWITCHFAB_CFG_4	SWITCH	FAB_CFG	>4	MIN	C
defSWITCHFAB_SEG_1	SWITCH	FAB_SEG	>1	MIN	A
defSWITCHFAB_SEG_2	SWITCH	FAB_SEG	>2	MIN	M
defSWITCHFAB_SEG_4	SWITCH	FAB_SEG	>4	MIN	C
defSWITCHFAULTY_PORTS_10	SWITCH	FAULTY_PORTS	>=10	NONE	M
defSWITCHFAULTY_PORTS_11	SWITCH	FAULTY_PORTS	>=11	NONE	C
defSWITCHFAULTY_PORTS_25	SWITCH	FAULTY_PORTS	>=25	NONE	C
defSWITCHFAULTY_PORTS_5	SWITCH	FAULTY_PORTS	>=5	NONE	A
defSWITCHFAULTY_PORTS_6	SWITCH	FAULTY_PORTS	>=6	NONE	M
defSWITCHFLOGI_4	SWITCH	FLOGI	>4	MIN	A
defSWITCHFLOGI_6	SWITCH	FLOGI	>6	MIN	M
defSWITCHFLOGI_8	SWITCH	FLOGI	>8	MIN	C
defSWITCHL2_DEVCNT_PER_60	SWITCH	L2_DEVCNT_PER	>60	NONE	A
defSWITCHL2_DEVCNT_PER_75	SWITCH	L2_DEVCNT_PER	>75	NONE	M
defSWITCHL2_DEVCNT_PER_90	SWITCH	L2_DEVCNT_PER	>90	NONE	C
defSWITCHLSAN_DEVCNT_PER_60	SWITCH	LSAN_DEVCNT_PER	>60	NONE	A
defSWITCHLSAN_DEVCNT_PER_75	SWITCH	LSAN_DEVCNT_PER	>75	NONE	M
defSWITCHLSAN_DEVCNT_PER_90	SWITCH	LSAN_DEVCNT_PER	>90	NONE	C
defSWITCHMARG_PORTS_10	SWITCH	MARG_PORTS	>=10	NONE	M
defSWITCHMARG_PORTS_11	SWITCH	MARG_PORTS	>=11	NONE	C
defSWITCHMARG_PORTS_25	SWITCH	MARG_PORTS	>=25	NONE	C
defSWITCHMARG_PORTS_5	SWITCH	MARG_PORTS	>=5	NONE	A
defSWITCHMARG_PORTS_6	SWITCH	MARG_PORTS	>=6	NONE	M
defSWITCHMARG_SFPS_PER_10	SWITCH	MARG_SFPS	>=10	NONE	M
defSWITCHMARG_SFPS_PER_20	SWITCH	MARG_SFPS	>=20	NONE	M
defSWITCHMARG_SFPS_PER_21	SWITCH	MARG_SFPS	>=21	NONE	C
defSWITCHMARG_SFPS_PER_30	SWITCH	MARG_SFPS	>=30	NONE	C
defSWITCHMARG_SFPS_PER_5	SWITCH	MARG_SFPS	>=5	NONE	A
defSWITCHSEC_AUTH_FAIL_0	SWITCH	SEC_AUTH_FAIL	>0	MIN	A
defSWITCHSEC_AUTH_FAIL_2	SWITCH	SEC_AUTH_FAIL	>2	MIN	M
defSWITCHSEC_AUTH_FAIL_4	SWITCH	SEC_AUTH_FAIL	>4	MIN	C
defSWITCHSEC_CMD_0	SWITCH	SEC_CMD	>0	MIN	A
defSWITCHSEC_CMD_2	SWITCH	SEC_CMD	>2	MIN	M
defSWITCHSEC_CMD_4	SWITCH	SEC_CMD	>4	MIN	C

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defSWITCHSEC_DCC_0	SWITCH	SEC_DCC	>0	MIN	A
defSWITCHSEC_DCC_2	SWITCH	SEC_DCC	>2	MIN	M
defSWITCHSEC_DCC_4	SWITCH	SEC_DCC	>4	MIN	C
defSWITCHSEC_HTTP_0	SWITCH	SEC_HTTP	>0	MIN	A
defSWITCHSEC_HTTP_2	SWITCH	SEC_HTTP	>2	MIN	M
defSWITCHSEC_HTTP_4	SWITCH	SEC_HTTP	>4	MIN	C
defSWITCHSEC_LV_0	SWITCH	SEC_LV	>0	MIN	A
defSWITCHSEC_LV_2	SWITCH	SEC_LV	>2	MIN	M
defSWITCHSEC_LV_4	SWITCH	SEC_LV	>4	MIN	C
defSWITCHSEC_SCC_0	SWITCH	SEC_SCC	>0	MIN	A
defSWITCHSEC_SCC_2	SWITCH	SEC_SCC	>2	MIN	M
defSWITCHSEC_SCC_4	SWITCH	SEC_SCC	>4	MIN	C
defSWITCHSEC_TELNET_0	SWITCH	SEC_TELNET	>0	MIN	A
defSWITCHSEC_TELNET_2	SWITCH	SEC_TELNET	>2	MIN	M
defSWITCHSEC_TELNET_4	SWITCH	SEC_TELNET	>4	MIN	C
defSWITCHSEC_TS_D10	SWITCH	SEC_TS	>10	DAY	C
defSWITCHSEC_TS_D2	SWITCH	SEC_TS	>2	DAY	A
defSWITCHSEC_TS_D4	SWITCH	SEC_TS	>4	DAY	M
defSWITCHSEC_TS_H1	SWITCH	SEC_TS	>1	HOUR	A
defSWITCHSEC_TS_H2	SWITCH	SEC_TS	>2	HOUR	M
defSWITCHSEC_TS_H4	SWITCH	SEC_TS	>4	HOUR	C
defSWITCHZONE_CFGSZ_PER_70	SWITCH	ZONE_CFGSZ_PER	>70	NONE	A
defSWITCHZONE_CFGSZ_PER_80	SWITCH	ZONE_CFGSZ_PER	>80	NONE	M
defSWITCHZONE_CFGSZ_PER_90	SWITCH	ZONE_CFGSZ_PER	>90	NONE	C
defSWITCHZONE_CHG_10	SWITCH	ZONE_CHG	>10	DAY	C
defSWITCHZONE_CHG_2	SWITCH	ZONE_CHG	>2	DAY	A
defSWITCHZONE_CHG_5	SWITCH	ZONE_CHG	>5	DAY	M
defWR_1stXFER_RDY_11000	sys_flow_monitor_scsi	WR_1stXFER_RDY	>11000	10SEC	M
defWR_1stXFER_RDY_17500	sys_flow_monitor_nvme	WR_1stXFER_RDY	>1750	10SEC	A
defWR_1stXFER_RDY_3500	sys_flow_monitor_scsi	WR_1stXFER_RDY	>3500	10SEC	A
defWR_1stXFER_RDY_45000	sys_flow_monitor_nvme	WR_1stXFER_RDY	>45000	10SEC	C

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defWR_1stXFER_RDY_5500	sys_flow_monitor_nvme	WR_1stXFER_RDY	>5500	10SEC	M
defWR_1stXFER_RDY_90000	sys_flow_monitor_scsi	WR_1stXFER_RDY	>90000	10SEC	C
defWR_1stXFER_RDY_IO_11000	sys_flow_monitor_scsi	WR_1stXFER_RDY	>11000	IO	M
defWR_1stXFER_RDY_IO_17500	sys_flow_monitor_nvme	WR_1stXFER_RDY	>1750	IO	A
defWR_1stXFER_RDY_IO_3500	sys_flow_monitor_scsi	WR_1stXFER_RDY	>3500	IO	A
defWR_1stXFER_RDY_IO_45000	sys_flow_monitor_nvme	WR_1stXFER_RDY	>45000	IO	C
defWR_1stXFER_RDY_IO_5500	sys_flow_monitor_nvme	WR_1stXFER_RDY	>5500	IO	M
defWR_1stXFER_RDY_IO_90000	sys_flow_monitor_scsi	WR_1stXFER_RDY	>90000	IO	C
defWR_1stXFER_RDY_VIOL_10	ALL_F_PORTS	WR_1stXFER_RDY_VIOL	>10	10SEC	M
defWR_1stXFER_RDY_VIOL_20	ALL_F_PORTS	WR_1stXFER_RDY_VIOL	>20	10SEC	C
defWR_1stXFER_RDY_VIOL_5	ALL_F_PORTS	WR_1stXFER_RDY_VIOL	>5	10SEC	A
defWR_PENDING_IO_100	sys_flow_monitor_nvme	WR_PENDING_IO	>100	10SEC	A
defWR_PENDING_IO_125	sys_flow_monitor_nvme	WR_PENDING_IO	>125	10SEC	M
defWR_PENDING_IO_150	sys_flow_monitor_nvme	WR_PENDING_IO	>150	10SEC	C
defWR_PENDING_IO_200	sys_flow_monitor_scsi	WR_PENDING_IO	>200	10SEC	A
defWR_PENDING_IO_250	sys_flow_monitor_scsi	WR_PENDING_IO	>250	10SEC	M
defWR_PENDING_IO_300	sys_flow_monitor_scsi	WR_PENDING_IO	>300	10SEC	C
defWR_STATUS_TIME_100000	sys_flow_monitor_scsi	WR_STATUS_TIME	>100000	10SEC	C
defWR_STATUS_TIME_12000	sys_flow_monitor_scsi	WR_STATUS_TIME	>12000	10SEC	M
defWR_STATUS_TIME_2000	sys_flow_monitor_nvme	WR_STATUS_TIME	>2000	10SEC	A
defWR_STATUS_TIME_4000	sys_flow_monitor_scsi	WR_STATUS_TIME	>4000	10SEC	A
defWR_STATUS_TIME_50000	sys_flow_monitor_nvme	WR_STATUS_TIME	>50000	10SEC	C
defWR_STATUS_TIME_6000	sys_flow_monitor_nvme	WR_STATUS_TIME	>6000	10SEC	M
defWR_STATUS_TIME_IO_100000	sys_flow_monitor_scsi	WR_STATUS_TIME	>100000	IO	C

Rule	Group	Monitoring System	Threshold	TimeBase	Policy
defWR_STATUS_TIME_IO_12000	sys_flow_monitor_scsi	WR_STATUS_TIME	>12000	IO	M
defWR_STATUS_TIME_IO_2000	sys_flow_monitor_nvme	WR_STATUS_TIME	>2000	IO	A
defWR_STATUS_TIME_IO_4000	sys_flow_monitor_scsi	WR_STATUS_TIME	>4000	IO	A
defWR_STATUS_TIME_IO_50000	sys_flow_monitor_nvme	WR_STATUS_TIME	>50000	IO	C
defWR_STATUS_TIME_IO_6000	sys_flow_monitor_nvme	WR_STATUS_TIME	>6000	IO	M
defWR_STATUS_TIME_VIOL_10	ALL_F_PORTS	WR_STATUS_TIME_VIOL	>10	10SEC	M
defWR_STATUS_TIME_VIOL_20	ALL_F_PORTS	WR_STATUS_TIME_VIOL	>20	10SEC	C
defWR_STATUS_TIME_VIOL_5	ALL_F_PORTS	WR_STATUS_TIME_VIOL	>5	10SEC	A
def_IO_ERROR	sys_flow_monitor	IO_ERROR	>100	10SEC	A, C, M
def_RESERVE	sys_flow_monitor	RESERVE	>100	10SEC	A, C, M

Firmware Upgrade and Downgrade Considerations for MAPS

This section describes the specific considerations for MAPS when you upgrade or downgrade a firmware version.

Firmware Upgrade Considerations for MAPS

If you upgrade from Fabric OS v9.1.x to Fabric OS v9.2.x, the Rule Quiet Time and Global Quiet Time configurations are supported in default notification mode with a warning message. When you enable the adaptive notification feature, the Rule Quiet Time and Global Quiet Time configurations are no longer supported. They are only available in default notification mode.

Firmware Downgrade Considerations for MAPS

Consider the following items when you downgrade from the Fabric OS v9.2.x release:

- The custom rules configured for the ZONE_CFGSZ_PER monitoring system with the CHASSIS group are not supported in the pre-Fabric OS v9.2.x releases.
- The MEMORY_USAGE_STATE monitoring system is not supported in the pre-Fabric OS v9.2.x releases.
- The custom rules configured for the STATE_CHG monitoring system with the CIRCUIT QOS group are not supported in the pre-Fabric OS v9.2.x releases.
- Upon downgrading to the pre-Fabric OS v9.2.x versions with the adaptive notification enabled, a warning message appears. The Default and custom Quiet Time configurations that were created in default notification mode will be supported after downgrading.

MAPS Scale Numbers

This section describes the MAPS scale numbers on the MAPS basic elements.

Table 44: MAPS Scale Numbers

MAPS Elements	Maximum Number Allowed
Maximum number of rules allowed per policy (Default, Custom, and RoR)	550
Maximum number of policies allowed	20
Maximum number of custom rules allowed	500
Maximum number of RoR rules allowed in a policy	50
Maximum number of custom groups allowed	64
Maximum number of alerts for flows in 10 seconds	1000
Maximum length of a rule name	72 characters
Maximum length of a group name	32 characters
Maximum length of an RoR rule name	72 characters
Maximum length of a policy name	31 characters
Maximum number of quarantine ports per chassis	32
Maximum number of ports that can be quarantined per slow-drain port	32

MAPS supports FENCE/DECOM actions for the default rules and the associated RoR rules that are associated with the DEV_LATENCY_IMPACT monitoring system. When you configure the default rules on the DEV_LATENCY_IMPACT monitoring system with the disruptive actions, the mutual exclusion is validated as described in the following table. For example, if both SDDQ and FENCE actions are configured, an error message is displayed and the rule is not configured.

Table 45: Mutually Exclusive MAPS Actions

Action	Mutually Exclusive Actions
RASLOG	NONE
SNMP	NONE
EMAIL	NONE
SDDQ	FENCE, DECOM
FENCE	SDDQ, TOGGLE
DECOM	SDDQ, TOGGLE
UNQUAR	FENCE, DECOM
TOGGLE	FENCE, DECOM
FMS	NONE
FPIN	NONE
RE_BALANCE	NONE

New and Deprecated CLI Commands, Monitoring Systems, and Default Monitoring Rules

This section provides information on the updates that are made in the Fabric OS v9.2.0 release.

New MAPS CLI Commands

No new MAPS CLI commands are introduced in the Fabric OS v9.2.0 release.

Modified MAPS CLI Commands

The following MAPS CLI commands are enhanced in the Fabric OS v9.2.0 release:

- The `mapsconfig` command is enhanced to enable the adaptive notifications feature or to enable default MAPS notifications as follows:
  `mapsconfig --notification {default|adaptive}`
- The `mapsconfig --show` command output is enhanced to display the maps notification behavior.
- The `fpiprofile` command is enhanced to show the oversubscription state and tx threshold.
- The `mapssam --show` command is enhanced to show the free kernel memory with the existing physical memory.
- The `mapspolicy --show -summary` command is enhanced to show the newly added default monitoring policy dflt_always_active_policy.

New MAPS Monitoring Systems

The following MAPS monitoring systems are introduced in the Fabric OS v9.2.0 release:

- MEMORY_USAGE_STATE

Modified MAPS Monitoring Systems

The following MAPS monitoring systems are modified in the Fabric OS v9.2.0 release:

- The ZONE_CFGSZ_PER monitoring system is enhanced to monitor the chassis-wide zone configuration size by using CHASSIS as a group.
- The STATE_CHG monitoring system is enhanced to monitor the Circuit QOS states.

New Default MAPS Policies

The new default monitoring policy dflt_always_active_policy is introduced in the Fabric OS 9.2.0 release:

New MAPS Monitoring Certificates Support

The following new MAPS monitoring certificates are supported in the Fabric OS v9.2.0 release:

- EXTN SW
- EXTN CA

New Default Monitoring Rules

The following MAPS default monitoring rules are introduced in the Fabric OS v9.2.0 release:

- defALL_CIRCUIT_F_QOS_STATE_CHG_4
- defALL_CIRCUIT_HIGH_QOS_STATE_CHG_4
- defALL_CIRCUIT_LOW_QOS_STATE_CHG_4
- defALL_CIRCUIT_MED_QOS_STATE_CHG_4
- defALL_CIRCUIT_IP_HIGH_QOS_STATE_CHG_4
- defALL_CIRCUIT_IP_LOW_QOS_STATE_CHG_4
- defALL_CIRCUIT_IP_MED_QOS_STATE_CHG_4
- defCHASSISMEMORY_USAGE_STATE_CRIT
- defCHASSISMEMORY_USAGE_STATE_WARN
- defCHASSISMEMORY_USAGE_STATE_IN_RANGE
- defCHASSISZONE_CFGSZ_PER_70
- defCHASSISZONE_CFGSZ_PER_80
- defCHASSISZONE_CFGSZ_PER_90

For more information, see Default Monitoring Rules.

Modified Default Monitoring Rules

No default monitoring rules are modified in the Fabric OS v9.2.0 release.

Revision History

The revision history provides a list of the important changes made in each version of the document.

FOS-92x-MAPS-UG100; April 28, 2023

Initial document version.

Documentation Legal Notice

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

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

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


BROADCOM®