



Extreme OS ONE SR v22.2.2.0 Monitoring Configuration Guide

Threshold Configuration and Resource Management

9039568-00 Rev AA
April 2026



Copyright © 2026 Extreme Networks, Inc. All rights reserved.

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks® and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, see: <https://www.extremenetworks.com/about-extreme-networks/company/legal/trademarks>

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses.

End-user license agreements and open source declarations can be found at: <https://www.extremenetworks.com/support/policies/open-source-declaration/>



Table of Contents

Abstract.....	v
Preface.....	vi
Text Conventions.....	vi
Documentation and Training.....	vii
Open Source Declarations.....	viii
Training.....	viii
Help and Support.....	viii
Subscribe to Product Announcements.....	ix
Send Feedback.....	ix
About This Document.....	10
What's New in This Document.....	10
Supported Platforms.....	10
Hardware Monitoring.....	11
Hardware Monitoring Overview	11
Temperature Sensor Monitoring.....	11
Temperature Sensor Monitoring Overview.....	11
Key Capabilities.....	12
Best Practices.....	13
Events and Alert Generation.....	13
Monitoring Temperature Sensors.....	14
Threshold and System Event Monitoring.....	15
Threshold and System Event Monitoring Overview.....	15
Key Features.....	16
Main Aspects of Resource Monitoring.....	16
Benefits of Resource Monitoring.....	16
Resources Monitored.....	16
Configurable Thresholds.....	16
Alert Actions.....	16
Limitations.....	17
Resource Threshold Monitoring.....	17
Resource Threshold Monitoring Common Interface.....	17
Resource Threshold Monitoring SNMP Trap.....	18
Resource Threshold Monitoring Configuration and Default Behavior.....	19
System Resource Monitoring.....	19
Key Features.....	20
Monitoring Statistics.....	20
CPU Monitoring.....	21
Memory Monitoring.....	21
Configuration.....	22
Hardware Component Monitoring.....	22

Fan Status Monitoring.....	22
Power Supply Monitoring.....	24
Network Resource Monitoring.....	25
Supported Resources.....	25
Resource Threshold Configuration Parameters.....	26
BGP Protocol Event Monitoring and Notification.....	26
Supported Functionalities.....	26
BGP Enterprise and Standard MIB Notifications.....	26
RASlogs.....	28
BFD Protocol Event Monitoring and Notification.....	29
Supported Notifications.....	29
gNMI Notifications.....	30
RAS Logs.....	30
CLI Commands and Statistics for BFD.....	31
Best Practices.....	31
Threshold Configuration.....	31
Monitoring Strategy.....	31
Troubleshooting.....	32
CLI Commands for Threshold Monitoring and Alerting.....	32
Configuring System Resource Monitoring.....	32
Configuring Network Resource Monitoring.....	33
gNMI Commands for Threshold Monitoring and Alerting.....	35
SNMP MIBs for Threshold Monitoring and Alerting.....	36



Abstract

The *Extreme OS ONE SR Monitoring Configuration Guide* version 22.2.2.0 provides technical procedures for configuring and managing resource monitoring. Configuration enables precise control of thresholds, polling intervals, and alert actions for system and network resources to support operational reliability for network engineers and administrators. Key features include continuous temperature sensor monitoring with automated alerting and device reboot actions, threshold-based event monitoring for CPU, memory, fans, power supplies, and network resources, and integration with gNMI/gNOI protocols for model-driven management.



Preface

Read the following topics to learn about:

- The meanings of text formats used in this document.
- Where you can find additional information and help.
- How to reach us with questions and comments.

Text Conventions

Unless otherwise noted, information in this document applies to all supported environments for the products in question. Exceptions, like command keywords associated with a specific software version, are identified in the text.

When a feature, function, or operation pertains to a specific hardware product, the product name is used. When features, functions, and operations are the same across an entire product family, such as Extreme Networks switches, the product is referred to as *the switch*.

Table 1: Notes and warnings






Icon	Notice type	Alerts you to..
	Tip	Helpful tips and notices for using the product
	Note	Useful information or instructions
	Important	Important features or instructions
	Caution	Risk of personal injury, system damage, or loss of data
	Warning	Risk of severe personal injury

Table 2: Text

Convention	Description
screen displays	This typeface indicates command syntax, or represents information as it is displayed on the screen.
The words <i>enter</i> and <i>type</i>	When you see the word <i>enter</i> in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says <i>type</i> .
Key names	Key names are written in boldface, for example Ctrl or Esc . If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press Ctrl+Alt+Del
<i>Words in italicized type</i>	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles.
NEW!	New information. In a PDF, this is searchable text.

Table 3: Command syntax

Convention	Description
bold text	Bold text indicates command names, keywords, and command options.
<i>italic text</i>	Italic text indicates variable content.
[]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ x y z }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x y	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, such as passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, <i>member [member...]</i> .
\	In command examples, the backslash indicates a “soft” line break. When a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

[Hardware and Software Compatibility](#) for Extreme Networks products

[Extreme Optics Compatibility](#)

[Other Resources](#) such as articles, white papers, and case studies

Open Source Declarations

Some software files have been licensed under certain open source licenses. Information is available on the [Open Source Declaration](#) page.

Training

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit the [Extreme Networks Training](#) page.

Help and Support

If you require assistance, contact Extreme Networks using one of the following methods:

Extreme Portal

Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.

The Hub

A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

Call GTAC

For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2800. For the support phone number in your country, visit www.extremenetworks.com/support/contact.

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number, or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any actions already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Subscribe to Product Announcements

You can subscribe to email notifications for product and software release announcements, Field Notices, and Vulnerability Notices.

1. Go to [The Hub](#).
2. In the list of categories, expand the **Product Announcements** list.
3. Select a product for which you would like to receive notifications.
4. Select **Subscribe**.
5. To select additional products, return to the **Product Announcements** list and repeat steps 3 and 4.

You can modify your product selections or unsubscribe at any time.


Send Feedback

The User Enablement team at Extreme Networks has made every effort to ensure that this document is accurate, complete, and easy to use. We strive to improve our documentation to help you in your work, so we want to hear from you. We welcome all feedback, but we especially want to know about:

- Content errors, or confusing or conflicting information.
- Improvements that would help you find relevant information.
- Broken links or usability issues.

To send feedback, email us at Product-Documentation@extremenetworks.com.

Provide as much detail as possible including the publication title, topic heading, and page number (if applicable), along with your comments and suggestions for improvement.



About This Document

[What's New in This Document](#) on page 10

[Supported Platforms](#) on page 10

What's New in This Document

The *Extreme OS ONE SR Monitoring Configuration Guide* is unchanged for release 22.2.2.0.

Supported Platforms


Extreme OS ONE Switching and Routing 22.2.1.0 and later releases support Extreme 8520, Extreme 8720, Extreme 8730, and Extreme 8820 hardware platforms.



Note

Although many software and hardware configurations are tested and supported for this release, all possible configurations and scenarios are beyond this document's scope.

For information about other releases, see the documentation for those releases.



Hardware Monitoring

[Hardware Monitoring Overview](#) on page 11

[Temperature Sensor Monitoring](#) on page 11

Hardware monitoring lets you monitor CPU and memory usage of the system, interface and optic environmental status, critical thermal zones and components, and security status, and be alerted when configured thresholds are exceeded.

Hardware Monitoring Overview

For hardware monitoring, you can create policies with default options or custom options for nondefault thresholds. When the policies are applied, you can toggle between default settings and saved custom configuration settings and apply actions and thresholds separately. For example, you can choose to use default threshold settings together with a customized subset of available actions, or you can modify some threshold settings and use the default action settings. You can also pause monitoring and actions.

Temperature Sensor Monitoring

Temperature sensor monitoring protects your Extreme OS ONE devices by continuously monitoring critical thermal zones and components. The system generates alerts automatically when temperatures exceed safe operating thresholds and can reboot devices when critical temperatures threaten hardware integrity.

Temperature Sensor Monitoring Overview

Every device has a number of temperature sensors to monitor the temperatures of critical zones and components. These are used to determine the thermal policy of a switch to regulate the temperature to keep it within optimal limits. Each sensor has predefined warning, critical, and shutdown (reboot) thresholds.

You can use this feature to:

- Monitor thermal sensors continuously to ensure optimal operating temperatures.
- Monitor warning, critical, and shutdown temperature thresholds for each sensor.

- Receive automated responses via RASlogs, SNMP traps, and alerts when temperatures cross thresholds.

Temperature sensor monitoring operates on all BMC-based platforms (Extreme 8000 series) running Extreme OS ONE Switching and Routing. The feature integrates seamlessly with existing thermal management systems regardless of your specific hardware configuration.

Key Capabilities

The temperature sensor monitoring implementation delivers four core functionalities.

Continuous Monitoring

The continuous monitoring functionality:

- Polls thermal sensors every five seconds to ensure optimal operating temperatures.
- Uses each sensor to monitor specific critical zones and components within your device hardware.

Threshold Management

The system monitors three distinct temperature thresholds for each sensor:

- Warning: The system generates alerts when temperatures exceed this level.
- Critical: The system generates critical alerts indicating immediate attention required.
- Shutdown: The system automatically reboots the switch to prevent hardware damage.

Automated Response

When temperatures cross thresholds, the system:

- Generates Reliability, Availability, and Serviceability logs (RASlogs).
- Sends SNMP traps to monitoring systems.
- Creates alerts for both threshold violations and temperature normalization (when a temperature returns to the normal range after exceeding a threshold).
- Generates alerts and logs when shutdown thresholds are reached.

RASlogs and SNMP traps are also generated at temperature normalization.

Thermal Policy Management

The temperatures reported by the thermal sensors are used to determine the fan speeds of a device. The thermal policy also reboots the device if the shutdown threshold of a specific sensor is reached.

The system manages thermal policies differently based on hardware architecture:

- BMC-based platforms (Extreme 8000 series): The BMC manages thermal policy and performs automatic reboots.

- Non-BMC platforms: Extreme OS ONE handles both thermal policy and reboot functions.

Best Practices

Monitoring Integration

- Configure your SNMP monitoring systems to receive temperature alerts.
- Set up log aggregation to capture temperature-related RASlogs.
- Establish alerting workflows for critical temperature events.

Proactive Management

- Regularly review temperature trends using sensor data.
- Investigate recurring warning threshold violations.
- Ensure adequate ventilation and cooling for your device environment.
- Monitor ambient temperature conditions in equipment locations.

Emergency Response

- Develop procedures for critical temperature alerts.
- Plan for automatic device reboots during thermal emergencies.
- Maintain backup configurations for rapid device recovery.
- Test thermal monitoring integration with your network management systems.

Events and Alert Generation

The temperature sensor monitoring implementation provides alerting for changes in temperature state. The controls on the alert frequency prevent alert flooding.

Automated Response

The system generates alerts for these temperature state changes:

- Normal to Alarm (warning threshold exceeded)
- Alarm to Critical (critical threshold exceeded)
- Normal to Critical (direct transition to critical)
- Critical to Alarm (temperature decreasing from critical)
- Alarm to Normal (temperature returning to safe levels)
- Critical to Normal (temperature returning to safe levels from critical)

Alert Frequency

The system generates one alert per sensor per event, not continuous alerts while temperatures remain elevated. This prevents alert flooding while ensuring that you receive notification when each threshold is crossed.

Monitoring Temperature Sensors

Use the **show sysinfo sensor all** command to display current temperature readings and thresholds. This command displays:


- Sensor ID and descriptive name.
- Current temperature reading (°C).
- Warning threshold (°C).
- Critical threshold (°C).
- Shutdown threshold (°C).

The following example displays temperature information:

```
device# show sysinfo sensor all
```

```
Sensor Information
```

Id	Name	Current (°C)	Warning (°C)	Critical (°C)	Shutdown (°C)
1	CPU Core	50.00	76.00	80.00	84.00
2	Sensor 1 (AFO)	39.00	54.00	60.00	66.00
3	Sensor 2 (Hot Spot)	43.00	76.00	80.00	84.00
4	Sensor 3 (AFI)	31.00	54.00	60.00	66.00



Threshold and System Event Monitoring

[Threshold and System Event Monitoring Overview](#) on page 15

[Key Features](#) on page 16

[Limitations](#) on page 17

[Resource Threshold Monitoring](#) on page 17

[System Resource Monitoring](#) on page 19

[Hardware Component Monitoring](#) on page 22

[Network Resource Monitoring](#) on page 25

[BGP Protocol Event Monitoring and Notification](#) on page 26

[BFD Protocol Event Monitoring and Notification](#) on page 29

[Best Practices](#) on page 31

[CLI Commands for Threshold Monitoring and Alerting](#) on page 32

[gNMI Commands for Threshold Monitoring and Alerting](#) on page 35

[SNMP MIBs for Threshold Monitoring and Alerting](#) on page 36

The good health status of a device and the adequate performance of various resources that it uses will ensure that the device works as intended and is available continuously to perform its tasks. Monitoring the available resources and generating alerts when user-configured thresholds are breached is the fastest way to react to changes in the health and performance of the device. Threshold and system event monitoring is a comprehensive monitoring system to help ensure proactive identification and resolution of resource constraints before they impact network operations.

Threshold and System Event Monitoring Overview

This feature provides comprehensive monitoring capabilities for system resources and hardware components. The feature monitors critical resources such as CPU, memory, fans, power supplies, and various Layer 2/Layer 3 hardware resources to ensure optimal network performance and reliability.

When resource utilization reaches configured threshold values, the system automatically generates alerts, logs events, and can trigger SNMP traps to notify you of potential issues before they impact network operations.

Key Features

Threshold and system event monitoring provides configurable mechanisms to track resource utilization and generate notifications when defined limits are exceeded.

Main Aspects of Resource Monitoring

- **Resource Utilization Tracking:** The system monitors CPU and memory usage, notifying users when high usage thresholds are exceeded.
- **Configurable Notifications:** Supports various notification actions, including RAS logs, diagnostic data collection, and GNMI state updates.
- **Internal Notifications:** Enables callbacks to other subsystems for further actions, such as proactive system shutdown in critical scenarios.
- **Resource Monitoring Interface:** Provides a common interface for monitoring different resources, allowing for flexibility and scalability.

Benefits of Resource Monitoring

- **Proactive Issue Detection:** Enables early detection of potential issues, allowing for timely intervention and minimizing downtime.
- **Improved Resource Management:** Helps administrators manage system resources more effectively, optimizing performance and capacity planning.
- **Enhanced System Reliability:** By monitoring critical resources, the system can identify and respond to issues before they impact overall system reliability.

Resources Monitored

The resource monitoring functionality monitors the following:

- **System resources:** Monitors CPU and memory utilization at the system level.
- **Hardware components:** Tracks the status of fans and power supplies .
- **Network resources:** Monitors Layer 2 and Layer 3 hardware resources including MAC tables, routing tables, and ACL entries.

Configurable Thresholds

You can configure thresholds as follows:

- Set custom high and low threshold limits for monitored resources
- Configure polling intervals for system resources
- Define the number of polling retries before triggering actions

Alert Actions

The system generates alerts using the following methods:

- **RASlogs:** Generates structured log entries when thresholds are exceeded.

- SNMP traps: Sends SNMP notifications to network management systems.

Limitations

Threshold and system event monitoring has the following limitations and scope:

- No External Alarm Framework Integration: The system generates alerts through RASlogs and SNMP traps but does not integrate with a centralized alarm framework.
- Rate-Limiting Configurations: The system lacks configurations for rate-limiting the generation of events and traps
- SNMP Walk Limitation: The system does not support SNMP walk for threshold monitor MIBs.

Resource Threshold Monitoring

The system monitors various resources, including system, Layer 2, and Layer 3 resources. Administrators can configure threshold parameters for each resource, and when a threshold is breached, multiple actions can be triggered. Available actions include:

- RASlog: Generates a RASlog entry for record-keeping and troubleshooting.
- SNMP: Sends a SNMP trap to notify administrators of the issue.

These actions enable administrators to track resource usage, identify potential issues, and perform root cause analysis (RCA).

Resource Threshold Monitoring Common Interface

The resource threshold monitoring common interface provides a standardized way for microservices to monitor their resources and trigger actions when thresholds are breached. This interface has the following key features:

- GNMI Configuration Updates: The interface subscribes to GNMI configuration updates, allowing for dynamic configuration changes.
- Polling Mechanism: The interface uses a timer to poll resource usage at configured intervals, executing actions when thresholds are exceeded.

Resource Threshold Monitoring SNMP Trap

The system utilizes Extreme-defined MIB traps to notify when preconfigured thresholds for resources are exceeded, enabling proactive monitoring and management.

Table 4: SNMP Threshold Monitoring-MIB Notifications

Trap Name and OID	Varbinds	Description
extremeThreshMonNotif .1.3.6.1.4.1.1916.1.58.0.1	extremeThreshMonResourceId extremeThreshMonNotificationType extremeThreshMonResourceLimit	This notification is generated when the resource usage reaches the configured high threshold limit or falls below the low threshold limit; and the total number of this notification sent in configured time interval has not exceeded the configured max notification count.

Table 5: SNMP Threshold Monitoring-MIB Objects

Trap Name and OID	Varbinds	Description
extremeThreshMonResourceId .1.3.6.1.4.1.1916.1.58.1	Accessible-for-notify	Specifies the unique index to identify monitored resources. Syntax - INTEGER MacAddressTable(1) VxlanTunnelTable (2) LIFTable (3) BFDSession (4) bfdIPv4Session (5), - For devices where BFD resources are separate for IPv4 and IPv6 bfdIPv6Session (6) - For devices where BFD resources are separate for IPv4 and IPv6.
extremeThreshMonNotificationType .1.3.6.1.4.1.1916.1.58.2	Accessible-for-notify	Specifies the type of notification. Syntax - INTEGER rising (1) falling (2) rising (1)- resource usage reaches the configured high threshold limit. falling (2) - resource usage falls below the configured low threshold limit.
extremeThreshMonResourceLimit .1.3.6.1.4.1.1916.1.58.3	Accessible-for-notify	Specifies the configured threshold resource usage limit for this notification.

Resource Threshold Monitoring Configuration and Default Behavior

The system allows administrators to configure resource threshold monitoring for various resources, with configuration options available in the CLI commands section.

Configuration Elements

Each resource has the following configuration elements:

- High-Limit: The upper threshold value (in percentage) that triggers an action when exceeded.
- Low-Limit: The lower threshold value (in percentage) that triggers an action when the usage falls below it after exceeding the high limit.
- Action: The action to take when a threshold is breached, with options including:
 - raslog: Generate a RAS log entry.
 - all: Perform both raslog and snmp actions.
 - snmp: Send a SNMP trap.
 - none: No action is taken.

Additionally, for CPU and memory resources, the following configuration elements are available:

- Poll-Interval: The interval (in seconds) between polls.
- Poll-Retry: The number of retries before taking action.

Configuration Path

The configuration elements are stored in the config-db at the path `/components/component\[name=chassis-0]/chassis/utilization/resources/resource\[name=*]/config`, where `*` represents the resource name. The `state-db` is populated at a similar path.

Default Behavior

When the system boots up without threshold monitoring configuration, resources are not monitored, and the default action is none. However, for CPU and memory resources, the default action is raslog, meaning that these resources are monitored for usage even without explicit configuration, and a RAS log is generated when the default high-limit value is exceeded.

Polling Mechanism

Resources are monitored using a polling mechanism, with a default interval of 10 seconds for CPU and memory resources. When a threshold is breached, the configured action is taken.

System Resource Monitoring

Extreme OS ONE tracks various statistics to monitor CPU and memory usage at the system level. The system generates notifications when configured thresholds are breached.

CPU and memory usage percentage can vary with the hardware platform and the running application services. This feature lets you configure the usage percentage of resources to configure the time between polling intervals to monitor the resources and also configure the permitted number of polling retries. In a scaled configuration, you might experience a high CPU usage percentage of all cores as well as a higher memory usage percentage.

The system uses a polling mechanism to monitor CPU and memory usage, with configurable polling intervals and retry counts. At every polling interval, the resources CPU and memory usage are read.

When a threshold is breached, the system generates RAS logs and can trigger additional actions. When this threshold upper value is crossed, the actions are executed (such as generation of a RASlog containing diagnostic information for the resource for high usage).

Key Features

The system tracks various statistics for CPU and memory usage, including:

- System-Level Monitoring: Tracks CPU and memory usage across the system.
- Configurable Thresholds: You can set threshold values for CPU and memory usage.
- Notification Actions: Supports multiple actions when thresholds are breached, including RAS log generation and callbacks to other subsystems.

Monitoring Statistics

The system tracks various statistics for CPU and memory usage, including:

- CPU Usage
 - Current per-core CPU usage
 - Average CPU usage across all cores
 - CPU load average for 1-minute, 5-minute, and 15-minute intervals
 - CPU load average percentage for 1-minute, 5-minute, and 15-minute intervals.
- Memory Usage
 - Current memory usage percentage
 - Total, free, used, and available memory
- Process-Level Statistics
 - CPU usage percentage for top processes
 - Memory usage percentage for top processes
 - Command line and process name for top CPU and memory-consuming processes.

CPU Monitoring

The system continuously monitors CPU utilization across all cores and generates alerts when usage exceeds configured thresholds.

Default Configuration

- High threshold: 80%
- Low threshold: 80%
- Polling interval: 10 seconds
- Retry count: 1 attempt
- Default action: RASlog generation

Monitored Metrics

- Per-core CPU usage percentage
- Overall CPU usage (average across all cores)
- CPU load averages (1, 5, and 15 minutes)

Sample RASlog Alert Messages

```
{"Level":"info", "Service":"monitor-svc", "LogID":"27006", "Topic":"27006", "CPU use Percent": "99.75", "Time":"2024-11-26 09:58:00.9478", "Msg":"Warning: High CPU usage detected!"}
```

```
2025-08-19 13:14:10.4834 monitor-svc[3500]: Level:info LogID:0 Topic:16 Msg:Detailed System CPU Breakdown Load Avg 15min:0.25 CPU Count:4 Load Avg Percent 1min:1 Load Avg Percent 5min:2.75 Load Avg Percent 15min:6.25 System CPU Percent Avg:4.29 Load Avg 1min:0.04 Load Avg 5min:0.11
```

```
2025-08-19 13:14:10.4837 monitor-svc[3500]: Level:info LogID:27006 Topic:16 Msg: {"system_cpu":{"avg_percent":4.29,"per_cpu_percent": [4.08,3.09,5,5],"load_avg_1min":0.04,"load_avg_5min":0.11,"load_avg_15min":0.25,"load_avg_percent_1min":1,"load_avg_percent_5min":2.75,"load_avg_percent_15min":6.25,"cpu_count":4}, "top_processes":[{"name":"k3s-server","percent":8.59}, {"name":"chassis","percent":2.94}, {"name":"tierra-svc","percent":2.27}, {"name":"containerd","percent":1.85}, {"name":"api-gw","percent":1.17}], "load_status":"low"}
```

Memory Monitoring

Memory utilization monitoring tracks system memory usage and identifies processes that consume the most memory.

Default Configuration

- High threshold: 80%
- Low threshold: 80%
- Polling interval: 10 seconds
- Retry count: 1 attempt
- Default action: RASlog generation

Monitored Metrics

- System memory usage percentage
- Total, free, used, and available memory
- Process-level memory consumption of top consumers CPU load averages (1, 5, and 15 minutes)

Sample RASlog Alert Messages

```
{ "Level": "info", "Service": "monitor-svc", "LogID": "27006", "Topic": "27006", "Memory use
Percent": "80.99", "Time": "2024-11-25 11:16:53.139 UTC +0000", "Msg": "Warning: High Memory
usage detected!"}

2025-08-19 13:13:58.4168 monitor-svc[3500]: Level:info LogID:0 Topic:16 Msg:Detailed
System CPU Breakdown Load Avg Percent 15min:6.25 CPU Count:4 Load Avg 1min:0.05 Load Avg
15min:0.25 Load Avg Percent 1min:1.25 Load Avg Percent 5min:3 System CPU Percent
Avg:11.91 Load Avg 5min:0.12

2025-08-19 13:13:58.4171 monitor-svc[3500]: Level:info LogID:27006 Topic:16 Msg:
{"system_cpu":{"avg_percent":11.91,"per_cpu_percent":
[10.2,11.11,15,11.34],"load_avg_1min":0.05,"load_avg_5min":0.12,"load_avg_15min":0.25,"loa
d_avg_percent_1min":1.25,"load_avg_percent_5min":3,"load_avg_percent_15min":6.25,"cpu_coun
t":4},"top_processes":[{"name":"k3s-server","percent":8.58},
{"name":"chassis","percent":2.94}, {"name":"tierra-svc","percent":2.27},
{"name":"containerd","percent":1.86}, {"name":"api-
gw","percent":1.17}], "load_status":"low"}
```

Configuration

You can configure threshold values and notification actions using the openconfig path `/components/component\[name=chassis-0]/chassis/utilization/resources/resource\[name=*/config/`.

For more information on configuration attributes, see the *gNMI Commands* section in the [CLI Commands for Threshold Monitoring and Alerting](#) on page 32 topic.

Hardware Component Monitoring

Monitoring of hardware components includes the device's fans and power supplies. These units are monitored for status including failures. RASlogs are generated when the status of a unit changes.

Fan Status Monitoring

The system polls fan information from the baseboard management controller (BMC) every 30 seconds and generates alerts for status changes.

Fan Events

- Fan enable: When a fan is inserted and becomes operational
- Fan fault: When a fan experiences a failure

- Fan departure: When a fan is removed from the chassis

RASlogs for Fan Events

The system generates RASlogs for fan-related failures and recoveries. Fan and PSU information is polled from the BMC every 30 seconds.

Log ID	Cause	Impact	Level	Message	Remedy
5300	A fan has been inserted into the chassis	A new fan is available	Info	fanEnable(%d): Status: %v, RPM: %v, RPM Percent: %v	N/A
5301	A fan has a fault	Faulty fan	Warning	fanFaulty(%d): Cur State: %s: Fan Failed	Check and replace the fan
5302	A fan has been removed from the chassis	The removed fan is unavailable	Info	fanDeparted(%d): Cur State: %s: Fan extracted	NA

Status Information

For the Extreme 8730 platform, the higher of the inlet and the outlet speed is used. For all other platforms, Fan_SYS_x_2 speed is used for the fan speed display.

- Fan ID and operational status
- Revolutions per minute (RPM)
- Speed percentage and level (low/medium/high)
- Airflow direction (front-to-back or back-to-front)

CLI Commands for Fan Status

The **show sysinfo fan** command displays the status of fans, including:

- Fan ID: Unique identifier for each fan.
- Status: Current status of the fan ("Up" or "Down").
- RPM: Fan speed in revolutions per minute.
- Percentage: Fan speed as a percentage of maximum speed.
- Speed Level: Fan speed level (LOW, MEDIUM, or HIGH).
- Direction: Fan airflow direction (FAN_DIR_F2B for front-to-back airflow).

gNMI Status Notification

The **show system internal sdb path components/component[name=fan-0]** command provides detailed fan information and alerting capabilities to help administrators monitor and manage fan health.

```
32d# show sysinfo fan
Fan Information
Id Status RPM Percentage SpeedLevel Direction
-----
1 Up 7400 30 LOW FAN_DIR_F2B
2 Up 7400 30 LOW FAN_DIR_F2B
3 Up 7600 30 LOW FAN_DIR_F2B
4 Up 7400 30 LOW FAN_DIR_F2B
5 Up 7400 30 LOW FAN_DIR_F2B
```

```

6 Up 7600 30 LOW FAN_DIR_F2B
7 Up 7400 30 LOW FAN_DIR_F2B

FAN_DIR_F2B - Fan Airflow Direction is FrontToBack

FanSpeedLevel - <40% [LOW], 40-70% [MEDIUM], >70% [HIGH]

```

Power Supply Monitoring

Power supply units (PSUs) are continuously monitored for status changes and operational parameters.

PSU Events

- PSU enable: When a PSU comes online and begins operation
- PSU fault: When a PSU experiences a failure condition
- PSU departure: When a PSU is removed from the system

RASlogs for PSU Events

The system generates RAS logs for PSU-related failures and recoveries.

Log ID	Cause	Impact	Level	Message	Remedy
5400 check	A PSU has been inserted into the chassis	A new PSU seated in a chassis	Info	psuPresent(%d): Status: %v	NA
5400	A PSU has come online	A new PSU is available	Info	psuEnable(): Psu Id: %d, cur_state: %s: Admitting PSU	NA
5401	A PSU has a fault	Faulty PSU	Warning	psuFaulty(%d): Cur State: %s: Psu Failed	Check and replace the PSU
5402	A PSU has been removed from the chassis	The removed PSU is unavailable	Info	psuDeparted(%d): Cur State: %s: Psu Departed	N/A

Status Information

- PSU ID and operational status
- Power supply type (AC or DC)
- Input/output current, voltage, and power measurements

CLI Commands for PSU Status

The **show sysinfo power-supply** command displays the status of PSUs, including:

- PSU ID: Unique identifier for each PSU.
- Status: Current status of the PSU ("Up" or "Present").
- Type: PSU type (for example, AC).

- Current: Input and output current in amps.
- Power: Input and output power in watts.
- Voltage: Input and output voltage in volts.

The following is an example output:

```
device# show sysinfo power-supply
PSU Information
Id  Status  Type  C[in]  C[out]  P[in]  P[out]  V[in]  V[out]
-----
1   Up      AC    1.200  18      270    210     225    11
2   Present AC     0       0       0       0       256     0
```

GNMI Status Notification

The **show system internal sdb path components/component[name=psu-0]** command provides detailed PSU information. The system provides detailed PSU information and alerting capabilities to monitor and manage PSU health.

Network Resource Monitoring

Monitoring of network resources includes Layer 2 and Layer 3 resources (such as MAC address tables and next-hop entries) as well as ACL resources (such as IPv4/IPv6 ingress and egress ACL entries) and advanced resources (such as resilient hashing and encapsulation tables).

Supported Resources

The following network resources can be monitored for utilization.

Layer 2 Resources

- MAC address table
- VXLAN tunnel table
- Logical Interface (LIF) table

Layer 3 Resources

- IPv4 and IPv6 routes
- IPv4 and IPv6 host entries
- Next-hop entries
- Equal-cost multi-path (ECMP) entries
- Bidirectional forwarding detection (BFD) sessions

ACL Resources

- IPv4/IPv6 Ingress and egress ACL entries
- MAC ACL entries
- Router ACL entries

Advanced Resources

- Resilient hashing
- Encapsulation tables
- Route/host combined tables

Resource Threshold Configuration Parameters

Each monitored resource supports the following configuration parameters:

- High limit: Percentage threshold that triggers high usage alerts
- Low limit: Percentage threshold for returning to normal status
- Action: Response when thresholds are exceeded (RASlog, SNMP, both, or none)

BGP Protocol Event Monitoring and Notification

The BGP Protocol Event Monitoring and Notification feature automates alerts about issues such as policy misconfigurations and route flaps rather than requiring manual CLI checks, improving efficiency and network stability. Key events monitored include BGP IPv4 and IPv6 connection state transitions (one of the pre-Established to Established state) and Established to one of the pre-Established state).

Supported Functionalities

You can configure BGP Protocol Event Monitoring and Notification to send the following types of data:

- gNMI notifications as part of BGP Finite State Machine (FSM) events and BGP session UP/DOWN events
- RAS trace logs as part of BGP UP/DOWN events
- SNMP traps as part of BGP FSM events for Enterprise MIB and Standard MIB

BGP Enterprise and Standard MIB Notifications

BGP reports significant events to the message bus when a BGP session changes state to Established or experiences backward transitions. These BGP traps contain session information within their payload/Varbind. The BGP Enterprise and Standard MIB define specific trap OID and Varbind lists for this purpose.

BGP standard MIB notifications are sent for the peers of IPv4 types.

Table 6: BGP Standard MIB Notifications

Trap Name and OID	Varbinds	Description
bgpEstablishedNotification 1.3.6.1.2.1.15.0.1	bgpPeerRemoteAddr bgpPeerLastError bgpPeerState	The bgpEstablishedNotification event is generated when the BGP FSM enters the established state.
bgpBackwardTransNotification 1.3.6.1.2.1.15.0.2	bgpPeerRemoteAddr, bgpPeerLastError, bgpPeerState	The bgpBackwardTransNotification event is generated when the BGP FSM moves from a higher numbered state to a lower numbered state.

The BGP enterprise MIB notifications are sent for the peers of IPv6 types.

Table 7: BGP Enterprise MIB Notifications

Trap Name and OID	Varbinds	Description
extremeBGP4V2EstablishedNotification 1.3.6.1.4.1.1916.1.51.0.1	extremeBgp4V2PeerState extremeBgp4V2PeerLocalPort extremeBgp4V2PeerRemotePort extremeBgp4V2PeerRemoteAddr	The extremeBGP4V2EstablishedNotification event is generated when the BGP FSM enters the established state.
extremeBGP4V2BackwardTransitionNotification 1.3.6.1.4.1.1916.1.51.0.2	extremeBgp4V2PeerState extremeBgp4V2PeerLocalPort extremeBgp4V2PeerRemotePort extremeBgp4V2PeerLastErrorCodeReceived, extremeBgp4V2PeerLastErrorSubCodeReceived, extremeBgp4V2PeerLastErrorReceivedText extremeBgp4V2PeerRemoteAddr	The extremeBGP4V2BackwardTransitionNotification event is generated when the BGP FSM moves from a higher numbered state to a lower numbered state.

Following is an example of a BGP SNMP trap message:

```
2025-02-17 09:49:36 <UNKNOWN> [UDP: [10.32.100.182]:53908->[10.37.32.26]:200]:
iso.3.6.1.2.1.1.3.0 = Timeticks: (25057400) 2 days, 21:36:14.00 iso.3.6.1.6.3.1.1.4.1.0 =
OID: iso.3.6.1.4.1.1916.1.51.0.2 iso.3.6.1.4.1.1916.1.51.1.2.1.13 = STRING: "Idle"
iso.3.6.1.4.1.1916.1.51.1.2.1.5 = STRING: "1000::1" iso.3.6.1.4.1.1916.1.51.1.3.1.1
= Counter32: 6 iso.3.6.1.4.1.1916.1.51.1.3.1.2 = Counter32: 6
iso.3.6.1.4.1.1916.1.51.1.2.1.6 = Counter32: 179 iso.3.6.1.4.1.1916.1.51.1.2.1.9
= Counter32: 38292 iso.3.6.1.4.1.1916.1.51.1.3.1.4 = STRING: "CEASE
OTHER_CONFIG_CHANGE"
```

gNMI Notifications

BGP neighbor configuration or state changes are published to the gNMI path `network-instances/network-instance[name=]/protocols/protocol[identifier=BGP][name=bgp]/bgp/neighbors/neighbor[neighbor-address=]/`

```

+--rw network-instances
  +--rw network-instance* [name]
    . . .
    +--rw protocols
      | +--rw protocol* [identifier name]
      |
      | +--rw bgp
      |
      | . . .
      | | +--rw neighbors
      | | | +--rw neighbor* [neighbor-address]
      | | | | +({)rw neighbor-address{-} > ../config/neighbor-address
      | | | . . .
      | | | | +--ro state oc-inet:as-number
      | | | | | +--ro session-state?
      | | | enumeration
      | | | | +--ro last-established? oc-
      | | | | | types:timeticks64
      | | | | | +--ro established-transitions? oc-
      | | | | | | yang:counter64
      | | | | | | +--ro supported-capabilities*
      | | | | | | identityref
      | | | | | | | +--ro messages
      | | | | | | | | +--ro sent
      | | | | | | | | | +--ro UPDATE? uint64
      | | | | | | | | | +--ro NOTIFICATION? uint64
      | | | | | | | | | +--ro last-notification-time? oc-
      | | | | | | | | | | types:timeticks64
      | | | | | | | | | | | +--ro last-notification-error-code? identityref
      | | | | | | | | | | | +--ro last-notification-error-subcode? identityref
      | | | | | | | | | | | +--ro received
      | | | | | | | | | | | | +--ro UPDATE? uint64
      | | | | | | | | | | | | +--ro NOTIFICATION? uint64
      | | | | | | | | | | | | +--ro last-notification-time? oc-
      | | | | | | | | | | | | | types:timeticks64
      | | | | | | | | | | | | | | +--ro last-notification-error-code? identityref
      | | | | | | | | | | | | | | +--ro last-notification-error-subcode? identityref
  
```

RASlogs

The following are Session UP/Down RASlogs:

```

2025-01-16 11:04:36.7728 bgp[15]: {"Level":"info","LogID":9008,"Topic":2,"VRF":"default-vrf","Neighbor":"192.x.x.x","Reason":"ADMIN-DOWN","Msg":"Session DOWN"}
2025-01-16 11:04:36.7729 bgp[15]: {"Level":"info","LogID":9008,"Topic":2,"VRF":"default-vrf","Neighbor":"10.x.x.x","Reason":"ADMIN-DOWN","Msg":"Session DOWN"}
2025-01-16 11:06:29.1857 bgp[15]: {"Level":"info","LogID":9008,"Topic":2,"VRF":"default-vrf","Neighbor":"10.x.x.x","Msg":"Session UP"}
2025-01-16 11:06:31.8469 bgp[15]: {"Level":"info","LogID":9008,"Topic":2,"VRF":"default-vrf","Neighbor":"10.x.x.x","Msg":"Session UP"}
  
```

BFD Protocol Event Monitoring and Notification

Bidirectional Forwarding Detection (BFD) rapidly detects communication failures between routers/network systems, enabling quick establishment of alternative paths for routing protocols. The key features include:

- Fast failure detection (in milliseconds)
- Media-independent liveness detection

BFD operation runs in unicast, point-to-point mode between two systems, using small packet sizes for efficient liveness detection

- Detection of failures in interfaces, data links, and forwarding engines

Supported Notifications

- SNMP traps for BFD UP, DOWN, and ADMIN DOWN events
- gNMI notifications for BFD state changes
- RAS trace logs for BFD session state changes and microservice events

Table 8: BFD Enterprise MIB Notifications

Trap Name and OID	Varbinds	Description
extremeBfdSessUp 1.3.6.1.4.1.1916.1.55.0.1	bfdSessDiag bfdSessInterface bfdSessSrcAddrType bfdSessSrcAddr bfdSessDstAddrType bfdSessDstAddr ifName extremeBfdVrfName	This notification is triggered when the <code>bfdSessState</code> object for an entry in the <code>bfdSessTable</code> is transitioning to the up (4) state from a different state. At this point, the <code>bfdSessDiag</code> value is set to <code>noDiagnostic</code> (0).
extremeBfdSessDown 1.3.6.1.4.1.1916.1.55.0.2	bfdSessDiag bfdSessInterface bfdSessSrcAddrType bfdSessSrcAddr bfdSessDstAddrType bfdSessDstAddr ifName extremeBfdVrfName	This notification is triggered when the <code>bfdSessState</code> object for an entry in the <code>bfdSessTable</code> is about to transition to either the down (2) or <code>adminDown</code> (1) state from another state. The <code>bfdSessDiag</code> value provides the diagnostic code indicating the reason for this state change (e.g., <code>pathDown</code> (5)).

Following is an example of a BFD SNMP trap message:

```
2025-02-09 15:13:26 <UNKNOWN> [UDP: [10.38.61.132]:54888->[10.32.90.33]:162]:
SNMPv2-MIB::sysUpTime.0 = Timeticks: (203600) 0:33:56.00 SNMPv2-MIB::snmpTrapOID.0 =
OID: EXTREME-BFD-MIB::extremeBfdSessUp EXTREME-BFD-MIB::extremeBfdVrfName.1 = STRING:
vrf1 IF-MIB::ifName.1 = STRING: ethernet 0/1:1.0 BFD-STD-MIB::bfdSessDstAddr.1
= STRING: "10.1.1.2" BFD-STD-MIB::bfdSessDstAddrType = INTEGER: ipv4(1) BFD-STD-
MIB::bfdSessSrcAddr.1 = STRING: "10.1.1.1" BFD-STD-MIB::bfdSessSrcAddrType = INTEGER:
```

```
ipv4(1) BFD-STD-MIB::bfdSessInterface.1 = INTEGER: 570425346 BFD-STD-MIB::bfdSessDiag.1 =
INTEGER: noDiagnostic(0)
```

gNMI Notifications

- Published to gNMI path `/bfd/interfaces/interface[id=*]/peers/peer[local-discriminator=*]/state`
- Include detailed session information, such as local and remote addresses, session state, and diagnostic codes

```
bfd
  interfaces
    interface* [id]
      peers
        peer* [local-discriminator]
          state
          | local-address
          | remote-address
          | subscribed-protocols
          | session-state
          | remote-session-state
          | last-failure-time
          | failure-transitions
          | local-discriminator
          | remote-discriminator
          | local-diagnostic-code
          | remote-diagnostic-code
          | remote-minimum-receive-interval
          | demand-mode-requested
          | remote-authentication-enabled
          | remote-control-plane-independent
          | applied-profile
          | local-minimum-tx-interval
          | local-minimum-rx-interval
          | local-detection-multiplier
          | remote-minimum-transmit-interval
          | remote-detection-multiplier
          | authentication-failure
          | last-up-time
```

RAS Logs

- Log BFD session state changes, including UP, DOWN, and ADMIN DOWN events
- Include session ID, DIP, and other relevant details

```
bfd[51]: Level:info LogID:28003 Topic:1 Msg:BFD session is operationally UP SessionID:1
DIP:10.1.1.2

bfd[51]: Level:info LogID:28004 Topic:1 Msg:BFD session is operationally DOWN SessionID:1
DIP:10.1.1.2

bfd[51]: Level:info LogID:28005 Topic:1 Msg:BFD session is Administratively DOWN
SessionID:1 DIP:10.1.1.2
```

CLI Commands and Statistics for BFD

For details on syntax and command parameters, see the *Extreme OS ONE SR Command Reference Guide*.

- `show bfd`: displays a summary view of BFD interface-related information
- `show bfd neighbors`: displays BFD sessions with filters by Destination IP address, Interface, VRF (Virtual Routing and Forwarding) name, and Client application type
- `show bfd profile <NAME | all>`: displays profile parameters
- `clear counters bfd`: clears BFD session counters. Clears for the specified group of sessions. Allows grouping of sessions by Destination IP address, Interface, VRF (Virtual Routing and Forwarding) name, and Client application type
- Member Ethernet: allows to configure default profile to be used for sessions that are created under the given member interface.
- `profile NAME`: creates a BFD profile and subsequent relevant commands in the sub-mode
- `interval`: configures minimum transmit interval, minimum receive interval for BFD packets at local end-point, and configures multiplier value that helps to calculate detection timeout of BFD sessions

Statistics can be checked using the `curl` command to view BFD trap statistics, including total traps sent and the last trap sent time.

```
curl 0:9005/dump-global-dbs
BFD trap data
=====
totalTrapSent          lastTrapSentTime  totalUpTrapSent
lastUpTrapSentTime  totalDownTrapSent          lastDownTrapSentTime
=====
1002          2025-02-11 09:02:59.816311          1001          2025-02-11
09:02:59.816311          1          2025-02-11 08:43:24.128967
```

Best Practices

Threshold Configuration

- Set high thresholds based on your environment's normal operating range.
- Configure low thresholds 5-10% below high thresholds to prevent alert flapping.
- Use shorter polling intervals for critical resources.
- Enable both RASlogs and SNMP traps for comprehensive monitoring.

Monitoring Strategy

- Establish baseline resource utilization before setting thresholds.
- Monitor trends over time to identify capacity planning needs.
- Coordinate with network management systems to correlate alerts.
- Review and adjust thresholds regularly based on operational experience.

Troubleshooting

- Check RAS logs for detailed resource usage information.
- Use diagnostic data collection when thresholds are exceeded.
- Monitor process-level resource consumption for root cause analysis.
- Verify hardware component status when resource alerts occur.

CLI Commands for Threshold Monitoring and Alerting

Use CLI commands to enable system resource monitoring of the CPU and memory for the device and to enable monitoring of Layer 2 and Layer 3 network resources for the device.

Use the **threshold-monitor** command to configure the high limit, low limit, and actions to perform when the thresholds for usage are exceeded. When the percentage of usage of resources (such as CPU or system memory) reaches the threshold value configured, RASlogs are generated, and diagnostic information is captured.

For CPU and memory monitoring, the default action is to generate RASlogs when the default high-limit value is exceeded even without threshold monitoring configuration. For all other monitoring, there is no default action at system boot when the default high-limit value is exceeded unless threshold monitoring is configured.

Configuring System Resource Monitoring

Use the **threshold-monitor** (CPU and memory) command to configure system resource monitoring of all CPU cores and memory. For more information about this command, see the *Extreme OS ONE SR Command Reference Guide*.

1. From privileged EXEC mode, access global configuration mode.

```
device# configure terminal
```

2. Configure the high limit, low limit, and actions to perform when the thresholds for CPU usage are exceeded. The available actions are RASlog generation (**raslog**), SNMP notifications (**snmp**), or both (**a11**). You can optionally specify a polling interval from 10 to 3600 seconds (default is 10) and a number of polling retries from 1 to 100 (default is 1).

```
device(config)# threshold-monitor cpu high-limit 90 low-limit 85 action raslog poll-  
interval 15 poll-retry 5
```

3. Configure the high limit, low limit, and actions to perform when the thresholds for memory usage are exceeded. The available actions are RASlog generation (**raslog**), SNMP notifications (**snmp**), or both (**a11**). You can optionally specify a polling interval from 10 to 3600 seconds (default is 10) and a number of polling retries from 1 to 100 (default is 1).

```
device(config)# threshold-monitor memory high-limit 80 low-limit 75 action snmp poll-  
interval 10 poll-retry 3
```

4. (Optional) Verify the system resource monitoring configuration.

```
device(config)# do show running-config threshold-monitor
```

```
threshold-monitor cpu high-limit 95 low-limit 85 actions RASLOG poll-interval 20
threshold-monitor memory high-limit 85 low-limit 75 actions ALL poll-retry 1
!
device(config)#
```

The following example shows how to configure system resource monitoring of all CPU cores and memory for the device:

```
device# configure terminal
device(config)# threshold-monitor cpu high-limit 95 low-limit 85 actions RASLOG poll-
interval 20
device(config)# threshold-monitor memory high-limit 85 low-limit 75 actions ALL poll-
retry 1
device(config)#
```

The following example shows how to display the complete threshold monitoring configuration for the device. In this example, the CPU and memory threshold monitoring configurations are included in the output:

```
device# show running-config threshold-monitor

threshold-monitor cpu high-limit 95 low-limit 85 actions RASLOG poll-interval 20
threshold-monitor memory high-limit 85 low-limit 75 actions ALL poll-retry 1
threshold-monitor mac-table high-limit 85 low-limit 80 actions ALL
threshold-monitor acl-ipv4-in high-limit 95 low-limit 85 actions RASLOG
!
device#
```

Configuring Network Resource Monitoring

Use the **threshold-monitor** command to configure monitoring of all network resources for a device. For more information about this command, see the *Extreme OS ONE SR Command Reference Guide*.

The following table lists the Layer 3 and ACL resources for which you can configure network resource monitoring, and the corresponding **threshold-monitor** keywords.

Keyword to enable monitoring	Monitored resource description	Comments
acl-ipv4-in	ACL IPv4 ingress	IPv4 and IPv6 egress ACLs share the same resource group
acl-ipv4-out	ACL IPv4 egress	
acl-ipv6-in	ACL IPv6 ingress	
acl-ipv6-out	ACL IPv6 egress	
acl-mac-in	ACL MAC ingress	-
bfd-session	BFD sessions	
ecmp	ECMP table	
host	Host table	
lif	Logical interface	
nexthop	Next-hop table	

Keyword to enable monitoring	Monitored resource description	Comments
racl-ipv4-in	RACL IPv4 ingress	
racl-ipv6-in	RACL IPv6 ingress	
resilient-hashing	Resilient hashing	
route	Route table	
vxlan-tunnel	VxLAN tunnel	

1. From privileged EXEC mode, access global configuration mode.

```
device# configure terminal
```

2. Configure the high limit, low limit, and actions to perform when the thresholds for a Layer 2 or Layer 3 resource are exceeded. Specify one of the following resource keywords: { **acl-ipv4-in** | **acl-ipv4-out** | **acl-ipv6-in** | **acl-ipv6-out** | **acl-mac-in** | **bfd-session** | **ecmp** | **host** | **lif** | **mac-table** | **nexthop** | **racl-ipv4-in** | **racl-ipv6-in** | **resilient-hashing** | **route** | **vxlan-tunnel** }. The available actions are RASlog generation (**raslog**), SNMP notifications (**snmp**), or both (**all**).

```
device(config)# threshold-monitor mac-table high-limit 85 low-limit 80 action all
```

3. (Optional) Verify the network resource monitoring configuration.

```
device(config)# do show running-config threshold-monitor
```

```
threshold-monitor mac-table high-limit 85 low-limit 80 action all
!
device(config)#
```

The following example configures threshold monitoring for ACL IPv4 ingress with RASlog generation:

```
device# configure terminal
device(config)# threshold-monitor acl-ipv4-in high-limit 95 low-limit 85 actions raslog
device(config)#
```

The following example shows how to display the complete threshold monitoring configuration for the device. In this example, the MAC table and IPv4 ACL threshold monitoring configurations are included in the output:

```
device# show running-config threshold-monitor

threshold-monitor cpu high-limit 95 low-limit 85 actions RASLOG poll-interval 20
threshold-monitor memory high-limit 85 low-limit 75 actions ALL poll-retry 1
threshold-monitor mac-table high-limit 85 low-limit 80 actions ALL
threshold-monitor acl-ipv4-in high-limit 95 low-limit 85 actions RASLOG
!
device#
```

gNMI Commands for Threshold Monitoring and Alerting

The system uses the open-config path `/components/component\[name=chassis-0]/chassis/utilization/resources` to configure resource threshold monitoring. This path includes three augmented attributes:

- **action:** Applies to all resources, specifying the action to take when a threshold is breached.
- **poll-interval** and **poll-retry:** Specific to CPU and memory resources, configuring the polling interval and retry count.

These attributes enable flexible configuration of resource threshold monitoring.

```

+--rw chassis
|   +--rw config
|   +--ro state
|   +--rw utilization
|       +--rw resources
|           +--rw resource* [name]
|               +--rw name          -> ../config/name
|               +--rw config
|                   | +--rw name?                string
|                   | +--rw used-threshold-upper?  oc-types:percentage
|                   | +--rw used-threshold-upper-clear? oc-types:percentage
|
|                   | +--rw action?              resource-action
|                   | +--rw poll-interval?      uint16
|                   | +--rw poll-retry?         uint16
|               +--ro state
|                   +--ro name?                  string
|                   +--ro used-threshold-upper?  oc-types:percentage
|                   +--ro used-threshold-upper-clear? oc-types:percentage
|
|                   +--rw action?              resource-action
|                   +--rw poll-interval?      uint16
|                   +--rw poll-retry?         uint16
|                   +--ro used?                uint64
|                   +--ro committed?           uint64
|                   +--ro free?                uint64
|                   +--ro max-limit?           uint64
|                   +--ro high-watermark?      uint64
|                   +--ro last-high-watermark? oc-types:timeticks64
|                   +--ro used-threshold-upper-exceeded? boolean

```

The following is an example command output:

```

device# show system internal sdb path /components/component[name=chassis-0]
key /components/component[name=chassis-0]
{
  "chassis": {
    "utilization": {
      "resources": {
        "resource": [
          {
            "name": "cpu",
            "state": {
              "action": "RASLOG",
              "name": "cpu",
              "poll-interval": 10,
              "poll-retry": 1,
              "used-threshold-upper": 85,
              "used-threshold-upper-clear": 80
            }
          }
        ]
      }
    }
  }
}

```

```
}
},
```

SNMP MIBs for Threshold Monitoring and Alerting

For details on SNMP MIBs, see the *Extreme Threshold Monitoring MIB* and *CPU and Memory Utilization - MIB* Trapst topics in the *Extreme OS ONE SR SNMP MIB Reference Guide*.

```
EXTREME-THRESHOLDMONITOR-MIB DEFINITIONS ::= BEGIN

IMPORTS
    MODULE-IDENTITY, OBJECT-TYPE, NOTIFICATION-TYPE, Integer32
        FROM SNMPv2-SMI
        -- RFC 2578

    MODULE-COMPLIANCE, OBJECT-GROUP, NOTIFICATION-GROUP
        FROM SNMPv2-CONF
        -- RFC 2580

    extremeAgent
        FROM EXTREME-BASE-MIB;

extremeThresholdMonitorMIB MODULE-IDENTITY
    LAST-UPDATED "202403180000Z" -- 18 March 2024 00:00:00 GMT
    ORGANIZATION "Extreme Networks, Inc."
    CONTACT-INFO
        "Postal: Extreme Networks, Inc.
         2121 RDU Center Drive,
         Morrisville, NC 27560.
         E-mail: support@extremenetworks.com
         WWW: http://www.extremenetworks.com"

    DESCRIPTION
        "This MIB is used to monitor L2/L3/TCAM hardware resource
         utilization on the managed device."

    REVISION "202403180000Z" -- 18 March 2024 00:00:00 GMT"
    DESCRIPTION
        "Updated extremeHWResourceOverallUsage with bits & extremeHWResourceID
         with TCAM MAC/IPV4/IPV6 ingress & egress integer resource IDs."

    REVISION "202309200000Z" -- 20 September 2023 00:00:00 GMT
    DESCRIPTION
        "Deprecated extremeResourceThreshMonNotif,
         extremeThreshMonResourceId, extremeThreshMonNotifType
         and extremeThreshMonResourceLimit.

         Added extremeHWResourceUsageAlert,
         extremeHWResourceOverallUsage and extremeHWResourceTable.

         extremeResourceThreshMonNotif notification which is for status
         change of an individual resource is replaced by
         extremeHWResourceUsageAlert notification which will give the
         comprehensive status of all resources in a bitmap
         extremeHWResourceOverallUsage."

    REVISION "202205110000Z" -- 11 May 2022 00:00:00 GMT
    DESCRIPTION
        "Initial version"
    ::= { extremeAgent 58 }

    extremeThresholdMonNotifObjects OBJECT IDENTIFIER ::= { extremeThresholdMonitorMIB 0 }
-- Deprecated objects
```

```

-- extremeThreshMonResourceId      OBJECT-TYPE      ::= { extremeThresholdMonitorMIB 1 }
-- extremeThreshMonNotifType       OBJECT-TYPE      ::= { extremeThresholdMonitorMIB 2 }
-- extremeThreshMonResourceLimit   OBJECT-TYPE      ::= { extremeThresholdMonitorMIB 3 }
extremeThreshMonObjects            OBJECT IDENTIFIER ::= { extremeThresholdMonitorMIB 4 }

extremeHWResourceOverallUsage      OBJECT-TYPE
    SYNTAX          BITS {
        macAddressTable (0),
        vxlanTunnelTable (1),
        lifTable (2),
        bfdSession (3),
        bfdIPv4Session (4),
        bfdIPv6Session (5),
        ipv4Route (6),
        ipv6Route (7),
        routeTable (8),
        ipv4Host (9),
        ipv6Host (10),
        hostTable (11),
        nextHop (12),
        nextHopTable (13),
        ecmp (14),
        ecmpTable (15),
        routeHostTable (16),
        encapTable (17),
        resilientHashing (18),
        tcamMacIngress (19),
        tcamMacEgress (20),
        tcamIPv4Ingress (21),
        tcamIPv4Egress (22),
        tcamIPv6Ingress (23),
        tcamIPv6Egress (24)
    }
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "L2/L3 Resource usage status of the monitored resources whether
        resource usage reaches low or high limit based on the threshold
        limit configuration. Each bit represents the individual resource
        usage status. If the resource usage reaches the configured high
        threshold limit then the corresponding bit is set to 1. If the
        resource usage falls below the configured low threshold limit then
        corresponding bit is set to 0. If a resource is not supported in
        this system, the bit value should be 0.

        The bit 'macAddressTable (0)' indicates the usage status of MAC
        table utilization.
        The bit 'vxlanTunnelTable (1)' indicates the usage status of VXLAN
        tunnel scale.
        The bit 'lifTable (2)' indicates the usage status of LIF scale.
        The bit 'bfdSession (3)' indicates the usage status of BFD session
        scale.
        The bit 'bfdIPv4Session (4)' indicates the usage status of IPv4
        BFD session scale.
        The bit 'bfdIPv6Session (5)' indicates the usage status of IPv6
        BFD session scale.
        The bit 'ipv4Route (6)' indicates the usage status of IPv4 routes
        supported by current route profile.
        The bit 'ipv6Route (7)' indicates the usage status of IPv6 routes
        supported by current route profile.
        The bit 'routeTable (8)' indicates the usage status of route table
        utilization.
        The bit 'ipv4Host (9)' indicates the usage status of IPv4 host
        supported by current route profile."

```

```

The bit 'ipv6Host (10)' indicates the usage status of IPv6 host
supported by current route profile.
The bit 'HostTable (11)' indicates the usage status of host table
utilization.
The bit 'nextHop (12)' indicates the usage status of Next Hops
supported by the current route profile.
The bit 'nextHopTable (13)' indicates the usage status of Next Hop
table utilization.
The bit 'ecmp (14)' indicates the usage status of ECMP Next Hops
supported by the current route profile.
The bit 'ecmpTable (15)' indicates the usage status of ECMP table
utilization.
The bit 'routeHostTable (16)' indicates the usage status of
hardware space shared between routes (IPv4 and IPv6) and neighbors
(ARP and ND).
The bit 'encapTable (17)' indicates the usage status of ENCAP
hardware space.
The bit 'resilientHashing (18)' indicates the usage status of
Resilient Hashing Next Hops supported by the current route profile.
The bit 'tcamMacIngress (19)', indicates the usage status of TCAM
L2 Ingress table utilization.
The bit 'tcamMacEgress (20)', indicates the usage status of TCAM
L2 Egress table utilization.
The bit 'tcamIPv4Ingress (21)', indicates the usage status of TCAM
IPv4 Ingress table utilization.
The bit 'tcamIPv4Egress (22)', indicates the usage status of TCAM
IPv4 Egress table utilization.
The bit 'tcamIPv6Ingress (23)', indicates the usage status of TCAM
IPv6 Ingress table utilization.
The bit 'tcamIPv6Egress (24)', indicates the usage status of TCAM
IPv6 Egress table utilization."
 ::= { extremeThreshMonObjects 1 }

extremeHWResourceUsageTable OBJECT-TYPE
    SYNTAX          SEQUENCE OF ExtremeHWResourceUsageTableEntry
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "A table of L2/L3 hardware resources monitored for utilization."
    ::= { extremeThreshMonObjects 2 }

extremeHWResourceUsageTableEntry OBJECT-TYPE
    SYNTAX          ExtremeHWResourceUsageTableEntry
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "The conceptual row of extremeHWResourceUsageTable."
    INDEX {
        extremeHWResourceID
    }
    ::= { extremeHWResourceUsageTable 1 }

ExtremeHWResourceUsageTableEntry ::= SEQUENCE {
    extremeHWResourceID          INTEGER,
    extremeHWResourceUsageHighLimit  INTEGER,
    extremeHWResourceUsageLowLimit  INTEGER,
    extremeHWResourceUsage          INTEGER
}

extremeHWResourceID OBJECT-TYPE
    SYNTAX          INTEGER {
        macAddressTable (0),
        vxlanTunnelTable (1),
        lifTable (2),

```

```

        bfdSession          (3),
        bfdIPv4Session      (4),
        bfdIPv6Session      (5),
        ipv4Route           (6),
        ipv6Route           (7),
        routeTable          (8),
        ipv4Host            (9),
        ipv6Host            (10),
        hostTable           (11),
        nextHop             (12),
        nextHopTable        (13),
        ecmp                (14),
        ecmpTable           (15),
        routeHostTable      (16),
        encapTable          (17),
        resilientHashing    (18),
        tcamMacIngress       (19),
        tcamMacEgress        (20),
        tcamIPv4Ingress      (21),
        tcamIPv4Egress       (22),
        tcamIPv6Ingress      (23),
        tcamIPv6Egress       (24)
    }
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "Resource ID of the monitored L2/L3 hardware resource."
 ::= { extremeHWResourceUsageTableEntry 1 }

extremeHWResourceUsageHighLimit    OBJECT-TYPE
    SYNTAX      Integer32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "High threshold limit of hardware resource usage in percentage."
 ::= { extremeHWResourceUsageTableEntry 2 }

extremeHWResourceUsageLowLimit     OBJECT-TYPE
    SYNTAX      Integer32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "Low threshold limit of hardware resource usage in percentage."
 ::= { extremeHWResourceUsageTableEntry 3 }

extremeHWResourceUsage             OBJECT-TYPE
    SYNTAX  INTEGER {
        normal (0),
        high   (1)
    }
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "Hardware resource usage status. It is 'normal' if the usage
        falls below low threshold limit and 'high' if it goes to or
        above high threshold limit. If a resource is not supported for
        monitoring, then the status should be '0'."
 ::= { extremeHWResourceUsageTableEntry 4 }

-- Deprecated object
-- extremeResourceThreshMonNotif  NOTIFICATION-TYPE ::=
{ extremeThresholdMonNotifObjects 1}
    extremeHWResourceUsageAlert    NOTIFICATION-TYPE
        OBJECTS {

```

```

        extremeHWResourceOverallUsage
    }
    STATUS current
    DESCRIPTION
        "This notification is generated when any of the HW resource'usage
        goes to/above the configured high threshold limit or falls below
        the low threshold limit; and the total number of this notification
        sent in the configured time interval shall not exceed the
        configured max notification count."
    ::= { extremeThresholdMonNotifObjects 2}

--
-- Compliance Statements
--

extremeThreshMonMIBConformance OBJECT IDENTIFIER ::= { extremeThresholdMonitorMIB 2 }

extremeThreshMonObjectsGroup OBJECT-GROUP
    OBJECTS {
        extremeHWResourceOverallUsage,
        extremeHWResourceUsageHighLimit,
        extremeHWResourceUsageLowLimit,
        extremeHWResourceUsage
    }
    STATUS current
    DESCRIPTION
        "A collection of management objects for hardware resource threshold
        monitoring."
    ::= { extremeThreshMonMIBConformance 1 }

extremeThreshMonNotifGroup NOTIFICATION-GROUP
    NOTIFICATIONS {
        extremeHWResourceUsageAlert
    }
    STATUS current
    DESCRIPTION
        "A collection of hardware resource threshold monitoring
        notifications."
    ::= { extremeThreshMonMIBConformance 2 }

extremeThreshMonMIBCompliances MODULE-COMPLIANCE
    STATUS current
    DESCRIPTION
        "The compliance statement for SNMP entities implementing
        EXTREME-THRESHOLDMONITOR-MIB."

    MODULE -- this module
        MANDATORY-GROUPS {
            extremeThreshMonObjectsGroup,
            extremeThreshMonNotifGroup
        }
    ::= { extremeThreshMonMIBConformance 3 }

END

```