

Oracle® Enterprise Manager

System Monitoring Plug-in Metric Reference Manual for Network
Management

Release 6 (6.0)

E12020-01

January 2008

E12020-01

Copyright © 2008, Oracle. All rights reserved.

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. This document is not warranted to be error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software--Restricted Rights (June 1987). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs.

Oracle, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

Contents

Preface	iii
Audience	iii
Documentation Accessibility	iii
Related Documents	iv
Conventions	iv
How to Use This Manual	ix
Structure of the Metric Reference Manual	ix
Background Information on Metrics, Thresholds, and Alerts	x
1 Check Point Firewall Metrics	
1.1 Configuration Management Metrics	1-1
1.1.1 Firewall Summary Metrics	1-1
1.1.2 System Kernel Memory Metrics	1-2
1.1.3 Hash Kernel Memory (HMEM) Metrics	1-2
1.1.4 Network Interfaces Metrics	1-3
1.2 10-Megabit Network Cards Statistics Metrics	1-4
1.3 100-Megabit Network Cards Statistics Metrics	1-4
1.4 Chains Metrics	1-4
1.5 Connections Metrics	1-5
1.6 Cookies Metrics	1-5
1.7 CPU Metrics	1-5
1.8 CPU and Memory Utilization by Processes Metrics	1-6
1.9 Disk Storage Statistics Metrics	1-6
1.10 Firewall Memory Metrics	1-6
1.11 Firewall Memory Utilization Metrics	1-7
1.12 Fragments Metrics	1-7
1.13 Gigabit Network Cards Statistics Metrics	1-8
1.14 Hash Kernel Memory Metrics	1-8
1.15 Host Performance Memory Metrics	1-9
1.16 Host Performance Memory Metrics (Non-Linux)	1-10
1.17 Inspection Statistics Metrics	1-11
1.18 Load Metrics	1-11
1.19 Network Interface Packets Metrics	1-11
1.20 Network Interfaces Metrics	1-13

1.21	Response Metrics	1-13
1.22	Session Details Metrics	1-14
1.23	System Information Metrics	1-15
1.24	VPN Configuration Metrics.....	1-15
1.25	VPN Statistics	1-16

2 Juniper Netscreen Firewall Metrics

2.1	Address Resolution Protocol (ARP) Configuration Metrics.....	2-1
2.2	Address Resolution Protocol (ARP) Mappings Metrics.....	2-1
2.3	Division of Attacks Metrics	2-2
2.4	Dropped Packets Division on the Firewall Metrics	2-2
2.5	Firewall CPU Utilization Metrics	2-3
2.6	Firewall Memory Utilization Metrics.....	2-4
2.7	Interface Traffic Metrics	2-4
2.8	Netscreen Firewall Traffic Information Per Policy Metrics	2-4
2.9	Network Interfaces Configuration Metrics	2-5
2.10	Policy Settings Metrics	2-5
2.11	Response Metrics	2-5
2.12	Session Information Metrics.....	2-6
2.13	URL Filter Configuration Metrics.....	2-6
2.14	NSRP Virtual Security Device Group Configuration Metrics.....	2-6
2.15	NSRP Virtual Security Device Interface Configuration Metrics	2-7
2.16	NSRP Virtual Security Device Member Configuration Metrics.....	2-8
2.17	NSRP Virtual Security Device Status Metrics.....	2-8

3 F5 BIG-IP Local Traffic Manager Metrics

3.1	Configuration Management Metrics.....	3-1
3.1.1	Switch Configuration Metrics	3-1
3.1.2	Virtual Server Configuration Metrics	3-2
3.2	IP Interfaces Metrics	3-2
3.3	Nodes Metrics.....	3-3
3.4	Physical Interfaces Metrics	3-4
3.5	Profile Authentication Metrics.....	3-4
3.6	Profile FTP Metrics	3-5
3.7	Profile Persistence Metrics.....	3-5
3.8	Profile TCP Metrics.....	3-6
3.9	Profile UDP Metrics.....	3-6
3.10	Response Metrics	3-7
3.11	Server Pool Members Metrics	3-7
3.12	Server Pools Metrics	3-8
3.13	Switch Metrics	3-9
3.14	User Management Metrics.....	3-10
3.15	Virtual Server Statistics Metrics	3-10
3.16	iRule Metrics	3-11

4 Load Balancer Switch

4.1	IP Address Map.....	4-1
4.1.1	Host Name.....	4-1
4.2	IP Interfaces.....	4-1
4.3	Nodes.....	4-1
4.4	Physical Interfaces.....	4-2
4.4.1	Physical Interface Bits In Rate (Kb per second).....	4-2
4.4.2	Physical Interface Bits Out Rate (Kb per second).....	4-3
4.4.3	Physical Interface Status	4-4
4.5	Response.....	4-4
4.5.1	Switch Status	4-4
4.6	Server Group Members.....	4-5
4.6.1	Real Server Address Status	4-5
4.6.2	Real Server Status (F5)	4-6
4.7	Server Pools	4-7
4.7.1	Server Group Bits In Rate (Kb per second)	4-7
4.7.2	Server Group Bits Out Rate (Kb per second).....	4-8
4.7.3	Server Group Connections per second (F5)	4-8
4.7.4	Server Group Current Connections (F5)	4-9
4.7.5	Server Group Highest Connections	4-10
4.7.6	Server Group Total Connections	4-10
4.8	Switch	4-10
4.8.1	BitsIn Rate (Kbps)	4-10
4.8.2	BitsOut Rate (Kbps).....	4-11
4.8.3	CPU Utilization (%).....	4-11
4.8.4	Current Connections	4-12
4.8.5	Highest Connections	4-12
4.8.6	MemoryUtilization %.....	4-13
4.8.7	MemUsedbyKrnل bytes	4-13
4.8.8	Switch Conns/sec	4-13
4.8.9	Switch FailoverState	4-14
4.8.10	Total Connections	4-14
4.9	Virtual Servers	4-14
4.9.1	Host Name.....	4-14
4.9.2	Virtual Server Bits In Rate (Kb per second)	4-15
4.9.3	Virtual Server Bits Out Rate (Kb per second).....	4-15
4.9.4	Virtual Server connections (%)	4-16
4.9.5	Virtual Server Connections per second	4-17
4.9.6	Virtual Server Current Connections	4-18
4.9.7	Virtual Server Highest Connections	4-18
4.9.8	Virtual Server Status.....	4-18
4.9.9	Virtual Server Total Connections	4-19

Preface

This manual is a compilation of the following metrics provided in Oracle Enterprise Manager for the following:

- System-monitoring plug-in metrics for network management
- Load Balancer Switch metrics

Note: The Load Balancer Switch-related metrics are not part of the System Monitoring Plug-in release. They are part of the Enterprise Manager Grid Control release.

Audience

This document is intended for Oracle Enterprise Manager users interested in the following metrics:

- System-monitoring plug-in metrics for network management
- Load Balancer Switch metrics

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible, with good usability, to the disabled community. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at

<http://www.oracle.com/accessibility/>

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

TTY Access to Oracle Support Services

Oracle provides dedicated Text Telephone (TTY) access to Oracle Support Services within the United States of America 24 hours a day, seven days a week. For TTY support, call 800.446.2398.

Related Documents

For more information, see the following documents in the Oracle Enterprise Manager 10g Release 2 documentation set:

- *Oracle Enterprise Manager System Monitoring Plug-in Installation Guide for Check Point Firewall*
- *Oracle Enterprise Manager System Monitoring Plug-in Installation Guide for Juniper Networks NetScreen Firewall*
- *Oracle Enterprise Manager System Monitoring Plug-in Installation Guide for F5 BIG-IP Local Traffic Manager*
- *Oracle Enterprise Manager Concepts*
- *Oracle Enterprise Manager Grid Control Quick Installation Guide*
- *Oracle Enterprise Manager Grid Control Quick Installation Guide*
- *Oracle Enterprise Manager Grid Control Installation and Basic Configuration*
- *Oracle Enterprise Manager Configuration for Oracle Collaboration Suite*
- *Oracle Enterprise Manager Advanced Configuration*
- *Oracle Enterprise Manager Policy Reference Manual*
- *Oracle Enterprise Manager Extensibility*
- *Oracle Enterprise Manager Command Line Interface*
- *Oracle Enterprise Manager SNMP Support Reference Guide*
- *Oracle Enterprise Manager Licensing Information*

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

How to Use This Manual

The *System Monitoring Plug-in Metric Reference Manual for Network Management* lists all the plug-ins metrics for network management that Enterprise Manager monitors. This manual shows all the metric help available online, eliminating the need to have the Grid Control Console up and running.

This preface describes:

- [Structure of the Metric Reference Manual](#)
- [Background Information on Metrics, Thresholds, and Alerts](#)

Structure of the Metric Reference Manual

This manual contains chapters for Check Point Firewall, Juniper Networks Netscreen Firewall, and F5 BIG-IP Local Traffic Manager. The metrics in these chapters appear in alphabetical order according to category.

Metric Information

The information for each metric comprises a description and user action if available:

- Description
Provides an explanation following the metric name. This text defines the metric and, when available, provides additional information pertinent to the metric.
- User Action
Suggests how to solve the problem causing the alert.

Definitions of Columns in Metric Summary Tables

The Metric Summary table is part of the overall metric information. The following table provides descriptions of columns in the Enterprise Manager Metric Summary table.

Column Header	Column Definition
Target Version	Version of the target, for example, 9.0.2.x and 10.1.0.x. The x at the end of a version (for example, 9.0.2.x) represents the subsequent patchsets associated with that release.

Column Header	Column Definition
Server Evaluation Frequency	The rate at which the metric is evaluated to determine whether it has crossed its threshold. For server-generated alerts, the evaluation frequency is determined by Oracle Database internals. For example, if the evaluation frequency is 10 minutes, when the Average File Write Time degrades to the point an alert should trigger, it could be almost 10 minutes before Enterprise Manager receives an indication of the alert. This column is present in the Metric Collection Summary table only for Oracle Database 10g metrics.
Collection Schedule	The rate at which the Management Agent collects data. The collection frequency for a metric comes from the Enterprise Manager default collection file for that target type.
Upload Interval	The rate at which the Management Agent moves data to the Management Repository. For example, upload every n th collection. The upload frequency for a metric comes from the Enterprise Manager default collection file for that target type. This column is present in the Metric Collection Summary table only when the Upload Frequency is different from the Collection Frequency.
Comparison Operator	The comparison method Enterprise Manager uses to evaluate the metric value against the threshold values.
Default Warning Threshold	Value that indicates whether a warning alert should be initiated. If the evaluation of the warning threshold value returns a result of TRUE for the specified number of consecutive occurrences defined for the metric, an alert triggers at the warning severity level.
Default Critical Threshold	Value that indicates whether a critical alert should be initiated. If the evaluation of the critical threshold value returns a result of TRUE for the specified number of consecutive occurrences defined for the metric, an alert triggers at the critical severity level.
Consecutive Number of Occurrences Preceding Notification	Consecutive number of times a metric's value reaches either the warning threshold or critical threshold before a notification is sent.
Alert Text	Message indicating why the alert was generated. Words that display between percent signs (%) denote variables. For example, <code>Disk Utilization for %keyValue% is %value%%</code> could translate to <code>Disk Utilization for d0 is 80%</code> .

Abbreviations and Acronyms

To reduce the page count in this document, the following abbreviations and acronyms are used:

Abbreviation/Acronym	Name
Agent	Oracle Management Agent
Database	Oracle Database
OMS	Oracle Management Service
Repository	Oracle Management Repository

Background Information on Metrics, Thresholds, and Alerts

A metric is a unit of measurement used to determine the health of a target. It is through the use of metrics and associated thresholds that Enterprise Manager sends out alerts notifying you of problems with the target.

Thresholds are boundary values against which monitored metric values are compared. For example, for each disk device associated with the Disk Utilization (%) metric, you can define a different warning and critical threshold. Some of the thresholds are predefined by Oracle; others are not.

After a threshold is reached, an alert is generated. An alert is an indicator signifying that a particular condition has been encountered and is triggered when one of the following conditions is true:

- A threshold is reached.
- An alert has been cleared.
- The availability of a monitored service changes. For example, the availability of an application server changes from up to down.
- A specific condition occurs. For example, an alert is triggered whenever an error message is written to a database alert log file.

Alerts are detected through a polling-based mechanism by checking for the monitored condition from a separate process at regular, predefined intervals.

See Also: See the *Oracle Enterprise Manager Concepts* manual and the Enterprise Manager online help for additional information about metrics, thresholds, and alerts

Editing

Out of the box, Enterprise Manager comes with thresholds for critical metrics. Warning and critical thresholds are used to generate an alert, letting you know of impending problems so that you can address them in a timely manner.

To better suit the monitoring needs of your organization, you can edit the thresholds provided by Enterprise Manager and define new thresholds. When defining thresholds, the key is to choose acceptable values to avoid unnecessary alerts, while still being notified of issues in a timely manner.

You can establish thresholds that will provide pertinent information in a timely manner by defining metric baselines that reflect how your system runs for a normal period of time.

The metrics listed on the Edit Thresholds page are either default metrics provided by Oracle or metrics whose thresholds you can change.

Specifying Multiple Thresholds

The Specifying Multiple Thresholds functionality allows you to define various subsets of data that can have different thresholds. By specifying multiple thresholds, you can refine the data used to trigger alerts, which is one of the key benefits of using Enterprise Manager.

The key in specifying multiple thresholds is to determine how the comparison relates to the metric threshold as a whole. What benefit will be realized by defining a more stringent or lax threshold for that particular device, mount point, and so on?

For example, using the Average Disk I/O Service Time metric, you can define warning and critical thresholds to be applied to all disks (sd0 and sd1), or you can define different warning and critical thresholds for a specific disk (sd0). This allows you to adjust the thresholds for sd0 to be more stringent or lax for that particular disk.

Accessing Metrics Using the Grid Control Console

To access metrics in the Grid Control Console, use the All Metrics page associated with a particular target by doing the following:

1. From the Grid Control Console, choose the target.
2. On the target's home page, click All Metrics in the Related Links section.

3. On the All Metrics page, choose the metric of interest and click Help. The help for that metric appears.

Check Point Firewall Metrics

This chapter provides descriptions for all Check Point Firewall metric categories, and tables list and describe associated metrics for each category. The tables also provide user actions if any of the metrics for a particular category support user actions. Shaded rows represent key columns for a particular category.

1.1 Configuration Management Metrics

Configuration Management metrics consist of the following categories:

- Firewall Summary
- System Kernel Memory
- Hash Kernel Memory (HMEM)
- Network Interfaces

1.1.1 Firewall Summary Metrics

The metrics in this category represent a Check Point Firewall Installation. The metrics contain details of the firewall name, type, and version, and also list the security policy installed on the firewall instance.

- Table Name — MGMT_EMX_CPFW_SUMMARY
- View Name — MGMT_EMX_CPFW_SUMMARY_VIEW

Default Collection Interval — Every 24 hours

Table 1–1 Firewall Summary Metrics

Metric	Description
Filter Date	Date of the filter installation.
Filter Name	Name of the filter.
Kernel Build Number	Build number of the kernel.
Major Version	Major version of the firewall.
Minor Version	Minor version of the firewall.
Product	Type of product.
Policy Install Time	Time when the security policy was installed on the firewall.
Security Policy	Security policy installed on the firewall.
System Name	Name of the machine where the firewall is installed.

1.1.2 System Kernel Memory Metrics

System kernel memory refers to the amount of memory currently in use by the FireWall-1 kernel module. This also includes the amount of hash memory. The metrics in this category provide information related to the kernel memory statistics on the firewall.

- Table Name — MGMT_EMX_CPFW_HOSTMEM
 - View Name — MGMT_EMX_CPFW_HOSTMEM_VIEW
- Default Collection Interval — Every 24 hours

Table 1–2 System Kernel Memory Metrics

Metric	Description
Minimum Free Swap Memory Necessary	Least amount of free swap memory required.
System Physical Memory	Total system physical memory.
System Swap Memory	Total swap memory on the system.
Total Buffered Memory	Total buffered memory on the system.
Total Cached Memory	Total cached memory on the system.
Total Shared Memory	Total shared memory on the system.

1.1.3 Hash Kernel Memory (HMEM) Metrics

Hash kernel memory only stores the various tables used in the enforcement of firewall security policy. This memory is hard-wired (that is, it cannot be swapped out), so it is very important to correctly choose the size to not unnecessarily deprive the box of memory. The metrics in this category provide information about the initial and current allocated hash kernel memory on the firewall instance.

- Table Name — MGMT_EMX_CPFW_HMEM
 - View Name — MGMT_EMX_CPFW_HMEM_VIEW
- Default Collection Interval — Every 24 hours

Table 1–3 Hash Kernel Memory (HMEM) Metrics

Metric	Description
Block Size	Block size for hash kernel memory.
Current Allocated Blocks	Number of currently allocated blocks.
Current Allocated Bytes	Number of currently allocated bytes.
Current Allocated Pools	Number of currently allocated pools.
Initial Allocated Blocks	Number of initially allocated blocks.
Initial Allocated Bytes	Number of initially allocated bytes.
Initial Allocated Pools	Number of initially allocated pools.
Maximum Bytes	Maximum number of bytes.
Maximum Pools	Maximum number of pools.

1.1.4 Network Interfaces Metrics

The metrics in this category provide information about the configuration parameters such as interface name, IP address, MAC address, bandwidth, status, and so forth related to the interfaces on the Check Point firewall instance being monitored.

- Table Name — MGMT_EMX_CPFW_NW_INTF
- View Name — MGMT_EMX_CPFW_NW_INTF_VIEW

Default Collection Interval — Every 24 hours

Table 1–4 Network Interfaces Metrics

Metric	Description
Network Interface Index (key column)	Unique ID for each interface.
Bandwidth (bits/second)	Bandwidth of the interface in bits per second.
Desired Status	Desired status of the interface.
Interface IP Address	IP address of the interface.
Interface MAC Address	MAC address of the interface.
Interface Name	Name of the interface.
Interface Type	Type of interface, distinguished according to the physical/link protocol. Possible values for this metric are: 1 — Other 2 — regular1822 3 — hdh1822 4 — ddn-x25 5 — rfc877-x25 5 — ethernet-csmacd 7 — iso88023-csmacd 8 — iso88024-tokenBus 9 — iso88025-tokenRing 10 — iso88026-man 11 — starLan 12 — proteon-10Mbit 13 — proteon-80Mbit 14 — hyperchannel 15 — fddi 16 — lapb 17 — sdlc 18 — dsl 19 — e1 20 — basicISDN 21 — primaryISDN 22 — propPointToPointSerial 23 — ppp 24 — softwareLoopback 25 — eon 26 — ethernet-3Mbit 27 — nsip 28 — slip 29 — ultra 30 — ds3 31 — sip 32 — frame-relay
Subnet Mask	Subnet mask of the interface.

1.2 10-Megabit Network Cards Statistics Metrics

The metrics in this category provide information about bandwidth utilization, and incoming and outgoing traffic rate information for interfaces that have a bandwidth of 10 megabits.

Table 1–5 10-Megabit Network Cards Statistics Metrics

Metric	Description and User Action
Network Interface Index (key column)	A unique value for each interface.
10-Megabit Card Bandwidth Used (%)	Bandwidth utilization of the interface. The default warning and critical threshold values for this metric are set higher than what is expected to be necessary in many cases. You can provide a smaller value for the warning and critical thresholds based on the load on the firewall and your network conditions.
10-Megabit Card Incoming Traffic Rate (Kilobits/second)	Rate of incoming traffic on the interface.
10-Megabit Card Outgoing Traffic Rate (Kilobits/second)	Rate of outgoing traffic on the interface.

1.3 100-Megabit Network Cards Statistics Metrics

The metrics in this category provide information about bandwidth utilization, and incoming and outgoing traffic rate information for interfaces that have a bandwidth of 100 megabits.

Default Collection Interval — Every 24 hours

Table 1–6 100-Megabit Network Cards Statistics Metrics

Metric	Description and User Action
Network Interface Index (key column)	A unique value for each interface.
100-Megabit Card Bandwidth Used (%)	Bandwidth utilization of the interface. The default warning and critical threshold values for this metric are set higher than what is expected to be necessary in many cases. You can provide a smaller value for the warning and critical thresholds based on the load on the firewall and your network conditions.
100-Megabit Card Incoming Traffic Rate (Kilobits/second)	Rate of incoming traffic on the interface.
100-Megabit Card Outgoing Traffic Rate (Kilobits/second)	Rate of outgoing traffic on the interface.

1.4 Chains Metrics

The metrics in this category provide information about the number of chains that are allocated and free.

Default Collection Interval — Every 15 minutes

Table 1–7 Chains Metrics

Metric	Description
Chains Allocated	Number of allocated chains.
Chains Free	Number of free chains.

1.5 Connections Metrics

The metrics in this category provide information about the rate of connections to the firewall.

Default Collection Interval — Every 15 minutes

Table 1–8 Connections Metrics

Metric	Description
Connections per sec.	Rate of connections to the firewall.
Peak Connections	Peak number of connections to the firewall.

1.6 Cookies Metrics

Cookies are an abstract data type that FireWall-1 uses to represent packets in a consistent manner as each OS stores packets slightly differently. The metrics in this category provide statistical information about the cookies the firewall handles.

Default Collection Interval — Every hour

Table 1–9 Cookies Metrics

Metric	Description
Cookies Get	Number of times the firewall got data from the cookie.
Cookies Length	Number of times the firewall queried the length of the cookie.
Cookies Put	Number of times the firewall put data on the cookie.
Total Allocated Cookies	Number of cookies that were allocated outside of the initial cookie pool that was allocated.
Total Cookies	Total number of cookies the firewall handled.
Total DUP Cookies	Number of cookies (packets) that were duplicated.
Total Free Cookies	Number of cookies that were freed from the allocated cookies.

1.7 CPU Metrics

The metrics in this category provide information about the percentage of CPU utilization.

Default Collection Interval — Every 5 minutes

Table 1–10 CPU Metrics

Metric	Description and User Action
CPU Idle (%)	Percentage of idle CPU time.
CPU Utilization (%)	Percentage of CPU being used. A large CPU consumption causes the entire system to slow down. To analyze what is causing the problem, use the Solaris "top" system command and look for any firewall processes that seem to be consuming an excessive percentage of CPU.

1.8 CPU and Memory Utilization by Processes Metrics

The metrics in this category provide information about CPU and memory utilized by individual processes on the machine where the firewall is installed.

Default Collection Interval — Every 30 minutes

Table 1–11 CPU and Memory Utilization by Processes Metrics

Metric	Description and User Action
Process ID (key column)	Unique ID for each process running on the firewall instance.
Process Name (key column)	Unique name for each process running on the firewall instance.
CPU Utilization by Process (%)	The default warning and critical threshold values for this metric are set higher than what is expected to be necessary in many cases. You can provide a smaller value for the warning and critical thresholds based on the load on the firewall and your network conditions.
Memory Utilization by Process (%)	The default warning and critical threshold values for this metric are set higher than what is expected to be necessary in many cases. You can provide a smaller value for the warning and critical thresholds based on the load on the firewall and your network conditions.

1.9 Disk Storage Statistics Metrics

The metrics in this category provide information about the disk space utilization statistics.

Default Collection Interval — Every 15 minutes

Table 1–12 Disk Storage Statistics Metrics

Metric	Description and User Action
Disk Space Free (%)	Percent of free space on the disk
Disk Space Used (%)	Disk space utilization. High disk space utilization could cause the system to hang. If you see a high percentage, free the disk space.
Total Disk Space (GB)	Total disk space in gigabytes.
Total Free Disk Space (GB)	Total free disk space in gigabytes.

1.10 Firewall Memory Metrics

The metrics in this category provide information about the rate of attempts to free and allocate KMem.

Default Collection Interval — Every 30 minutes

Table 1–13 Firewall Memory Metrics

Metric	Description and User Action
Firewall Memory (KMem) Allocation Failures per sec.	Rate of failed attempts to allocate memory. A high value indicates that the firewall is almost out of memory space. The default critical threshold for this metric is not defined. You can provide a value for the warning and critical thresholds based on the load on the firewall and your network conditions.
Firewall Memory (KMem) Allocation Operations per sec.	Rate of operations to allocate memory.
Firewall Memory (KMem) Free Failures per sec.	Rate of failed attempts to free memory. A large value indicates that free memory is required, but another process on the firewall is using the memory.
Firewall Memory (KMem) Free Operations per sec.	Rate of operations to free memory.
Peak Used Firewall Memory (KMem) in KB	Peak value for used firewall memory in KB.
Used Firewall Memory (KMem) in KB	Amount of firewall memory used out of the total allocated memory.

1.11 Firewall Memory Utilization Metrics

The metrics in this category provide information about the host memory utilization.

Default Collection Interval — Every 5 minutes

Table 1–14 Firewall Memory Utilization Metrics

Metric	Description and User Action
Memory Used by Firewall (KB)	Host memory used by the firewall.
Memory Utilization by Firewall (%)	Percentage of host memory used by the firewall. A large CPU consumption causes the entire system to slow down. To analyze what is causing the problem, use the Solaris "top" system command and look for any firewall processes that seem to be consuming an excessive percentage of CPU.
Memory Utilization by Other Processes (%)	Percentage of host memory utilized by other processes.
Overall Memory (Physical + Swap) (KB)	Total available memory on the host.

1.12 Fragments Metrics

The metrics in this category provide information about the number of fragmented packets, as well as the number of fragments that have expired.

Default Collection Interval — Every hour

Table 1–15 Fragments Metrics

Metric	Description
Expired	Number of expired fragments.
Fragments	Number of fragments.
Packets	Number of fragmented packets.

1.13 Gigabit Network Cards Statistics Metrics

The metrics in this category provide information about bandwidth utilization, and incoming and outgoing traffic rate information for interfaces having a bandwidth of 1 gigabit.

Default Collection Interval — Every hour

Table 1–16 Gigabit Network Cards Statistics Metrics

Metric	Description and User Action
Network Interface Index (key column)	Unique value for each interface.
Gigabit Card Bandwidth Used (%)	Bandwidth utilization of the interface. The default warning and critical threshold values for this metric are set higher than what is expected to be necessary in many cases. You can provide a smaller value for the warning and critical thresholds based on the load on the firewall and your network conditions.
Gigabit Card Incoming Traffic Rate (Kilobits/second)	Rate of incoming traffic on the interface.
Gigabit Card Outgoing Traffic Rate (Kilobits/second)	Rate of outgoing traffic on the interface.

1.14 Hash Kernel Memory Metrics

Hash memory refers to the amount of memory allocated and used for FireWall-1's state tables. This tells you how much memory is available for the state tables (available), how much is currently in use, and what the high water mark is for memory usage for state tables (peak). It also provides information about the rate of attempts for allocating and freeing HMem, and also provides the HMem utilization. The metrics in this category provide information about the rate of attempts for allocating and freeing HMem and also provides the HMem utilization.

Default Collection Interval — Every 30 minutes

Table 1–17 Hash Kernel Memory Metrics

Metric	Description and User Action
Allocated Hash Kernel Memory (KB)	Total hash kernel memory in kilobytes allocated for storing the state tables.
Available Hash Kernel Memory (%)	Percentage of hash kernel memory available for use on the host system.
Available Hash Kernel Memory (KB)	Total hash kernel memory in kilobytes available for use on the host system.
Block Size	Block size for HMem.
Hash Kernel Memory (HMem) Allocation Failures per sec.	Rate of memory allocation failures. A large HMem consumption causes failures in allocation of memory to new processes. To analyze what is causing the problem, use the Solaris "top" system command and look for any firewall processes that seem to be consuming an excessive percentage of memory.
Hash Kernel Memory (HMem) Allocation Operations per sec.	Rate of memory allocation operations.

Table 1–17 (Cont.) Hash Kernel Memory Metrics

Metric	Description and User Action
Hash Kernel Memory (HMem) Free Failures per sec.	Rate of memory free failures. A large HMem consumption causes the failures in freeing of memory for new processes. To analyze what is causing the problem, use the Solaris "top" system command and look for any firewall processes that seem to be consuming an excessive percentage of memory.
Hash Kernel Memory (HMem) Free Operations per sec.	Rate of memory free operations.
Hash Kernel Memory Utilization (%)	A large HMem consumption causes the entire system to slow down. To analyze what is causing the problem, use the Solaris "top" system command and look for any firewall processes that seem to be consuming an excessive percentage of memory.
Maximum Hash Kernel Memory (KB)	Maximum hash kernel memory in kilobytes on the host system.
Peak Used Hash Kernel Memory (KB)	Peak value for hash kernel memory usage.
Used Hash Kernel Memory (KB)	Amount of hash kernel memory being used on the host system.

1.15 Host Performance Memory Metrics

The metrics in this category provide performance-related information about host memory. The metrics provide the total memory on the host along with the allocated and free memory percentage. They also provide the swap memory utilization.

Default Collection Interval — Every 30 minutes

Table 1–18 Host Performance Memory Metrics

Metric	Description and User Action
Available Overall (Physical + Swap) Memory (%)	Total available memory on the host.
Available System Physical Memory (KB)	Available real/physical memory space on the host.
Available System Swap Memory (KB)	Available swap space on the host.
Minimum Free Swap Memory Necessary (KB)	Minimum amount of swap required to be free, or else memErrorSwap is set to 1 and a memSwapErrorMsg string is returned.
Overall Memory (Physical + Swap) (KB)	Sum of physical and swap memory present on the host system.
Overall Memory Available (Physical + Swap) (KB)	Sum of physical and swap memory currently available on the host system.
Overall Memory Used (Physical + Swap) (KB)	Sum of physical and swap memory currently being used on the host system.
Overall (Physical + Swap) Memory Utilization (%)	A large memory consumption causes the entire system to slow down. To analyze what is causing the problem, use the Solaris "top" system command and look for any firewall processes that seem to be consuming an excessive percentage of memory.
Physical Memory Available (%)	Percentage of physical memory available on the host system.

Table 1–18 (Cont.) Host Performance Memory Metrics

Metric	Description and User Action
Physical Memory Used (KB)	Physical memory in kilobytes being used on the host system.
Physical Memory Utilization (%)	Percentage of physical memory being used on the host system.
Swap Memory Available (%)	Percentage of swap memory available on the host system.
Swap Memory Error	Error flag 1 indicates very little swap space remains. Refer to the swap memory error message to analyze the problem.
Swap Memory Error Message	Error message describing the error flag condition.
Swap Memory Used (KB)	Swap memory in kilobytes being used on the host system.
Swap Memory Utilization (%)	Percentage of swap memory being used on the host system.
Total Buffered Memory (KB)	Total buffered memory in kilobytes present on the host system.
Total Cached Memory (KB)	Total cached memory in kilobytes present on the host system.
Total Shared Memory (KB)	Total shared memory in kilobytes present on the host system.
Total System Physical Memory (KB)	Total real/physical memory size on the host.
Total System Swap Memory (KB)	Total swap size configured for the host.

1.16 Host Performance Memory Metrics (Non-Linux)

The metrics in this category provide performance-related information about host memory for non-Linux Operating Systems.

Default Collection Interval — Every 30 minutes

Table 1–19 Host Performance Memory Metrics (Non-Linux OS)

Metric	Description and User Action
Available Real Memory Free (%)	Real memory percentage not currently in use on the host system
Real Memory Utilization (%)	Real memory utilization percentage on the host system.
Total Active Real Memory (KB)	Total real memory in kilobytes currently in use on the host system.
Total Active Virtual Memory (KB)	Total virtual memory in kilobytes currently in use on the host system.
Total Free Memory (KB)	Total real memory in kilobytes not currently in use on the host system.
Total Real Memory (KB)	Total real memory in kilobytes available on the host system.
Total Virtual Memory (KB)	Virtual Real memory in kilobytes on the host system.
Available Virtual Memory Free (%)	Virtual memory percentage not currently in use on the host system.
Virtual Memory Utilization (%)	Virtual memory utilization percentage on the host system

1.17 Inspection Statistics Metrics

The metrics in this category provide information about the number of records, packets, extracts, lookups, and operations inspected by the firewall.

Default Collection Interval — Every 15 minutes

Table 1–20 Inspection Statistics Metrics

Metric	Description
Number of Extracts	Number of extracts inspected.
Number of LookUps	Number of LookUps inspected.
Number of Operations	Number of operations inspected.
Number of Packets	Number of packets inspected.
Number of Records	Number of records inspected.

1.18 Load Metrics

The metrics in this category provide information about the Firewall Module State and the rate of packets accepted, rejected, dropped, and logged by the firewall.

Default Collection Interval — Every 15 minutes

Table 1–21 Load Metrics

Metric	Description and User Action
Firewall Module State	State of the firewall inspection module.
Packets Accepted per sec.	Rate of packets accepted. The default warning and critical threshold values for this metric are set higher than what is expected to be necessary in many cases. You can provide a smaller value for the warning and critical thresholds based on the load on the firewall and your network conditions.
Packets Dropped per sec.	Rate of packets dropped. The default warning and critical threshold values for this metric are set higher than what is expected to be necessary in many cases. You can provide a smaller value for the warning and critical thresholds based on the load on the firewall and your network conditions.
Packets Logged per sec.	Rate of packets logged. The default warning and critical threshold values for this metric are set higher than what is expected to be necessary in many cases. You can provide a smaller value for the warning and critical thresholds based on the load on the firewall and your network conditions.
Packets Rejected per sec.	Rate of packets rejected. The default warning and critical threshold values for this metric are set higher than what is expected to be necessary in many cases. You can provide a smaller value for the warning and critical thresholds based on the load on the firewall and your network conditions.

1.19 Network Interface Packets Metrics

The metrics in this category provide information about the rate of inbound and outbound packets that are accepted, rejected, dropped, and logged on an interface of the firewall.

Default Collection Interval — Every 15 minutes

Table 1–22 Network Interface Packets Metrics

Metric	Description and User Action
Network Interface Index (key column)	Unique value for each interface.
Interface Name (key column)	Name of the interface.
Accepted Bytes In	Number of inbound bytes on an interface.
Accepted Bytes Out	Number of outbound bytes on an interface.
Accepted Packets In	Number of inbound packets accepted on an interface.
Accepted Packets Out	Number of outbound packets accepted on an interface.
Dropped Packets In	Number of inbound packets dropped on an interface.
Dropped Packets Out	Number of outbound packets dropped on an interface.
Incoming Accepted Packets per sec.	Rate of inbound packets accepted on an interface. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Incoming Dropped Packets per sec.	Rate of inbound packets dropped on an interface. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Incoming Logged Packets per sec.	Rate of inbound packets logged on an interface. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Incoming Rejected Packets per sec.	Rate of inbound packets rejected on an interface. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Incoming Total Packets per sec.	Rate of inbound packets on an interface. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Logged Packets In	Number of inbound packets logged on an interface.
Logged Packets Out	Number of outbound packets logged on an interface.
Outgoing Accepted Packets per sec.	Rate of outbound packets accepted on an interface. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Outgoing Dropped Packets per sec.	Rate of outbound packets dropped on an interface. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Outgoing Logged Packets per sec.	Rate of outbound packets logged on an interface. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Outgoing Rejected Packets per sec.	Rate of outbound packets rejected on an interface. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Outgoing Total Packets per sec.	Rate of outbound packets on an interface. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Rejected Packets In	Number of inbound packets rejected on an interface.
Rejected Packets Out	Number of outbound packets rejected on an interface.
Total Packets In	Number of inbound packets on an interface.
Total Packets Out	Number of outbound packets on an interface.

1.20 Network Interfaces Metrics

The metrics in this category provide information about the bandwidth and status of each interface, as well as the incoming and outgoing rate of packets on each interface.

Default Collection Interval — Every 15 minutes

Table 1–23 Network Interfaces Memory Metrics

Metric	Description and User Action
Network Interface Index (key column)	Unique value for each interface. The value for each interface must remain constant at least from one reinitialization of the entity's network management system to the next reinitialization.
Interface Name (key column)	Name of the interface.
Interface IP Address (key column)	IP address of the interface.
Bandwidth (Mbits/second)	Bandwidth of the interface.
Desired Status	Desired state of the interface. The testing state indicates that no operational packets can be passed.
Interface MAC Address	MAC address of the interface.
Interface Type	Type of interface distinguished according to the physical/link protocol(s) immediately "below" the network layer in the protocol stack.
Network Interface Status	When the value is other than 0, there is a difference between the desired and current status of the interface.
Operational Status	Current operational state of the interface.
Rate of Incoming (Rx) Packet Discards (%)	Rate of inbound packets chosen to be discarded. The default warning and critical threshold values for this metric are set lower than what is expected to be necessary in many cases. You can provide a higher value for the warning and critical thresholds based on the load on the firewall and your network conditions.
Rate of Incoming (Rx) Packet Errors (%)	Rate of inbound packets that contained errors.
Rate of Outgoing (Tx) Packet Discards (%)	Rate of outbound packets chosen to be discarded. The default warning and critical threshold values for this metric are set lower than what is expected to be necessary in many cases. You can provide a higher value for the warning and critical thresholds based on the load on the firewall and your network conditions.
Rate of Outgoing (Tx) Packet Errors (%)	Rate of outbound packets that could not be transmitted because of errors..
Rate of Overall Packet Discards (%)	Rate of total packets (inbound + outbound) discarded. The default warning and critical threshold values for this metric are set lower than what is expected to be necessary in many cases. You can provide a higher value for the warning and critical thresholds based on the load on the firewall and your network conditions.
Rate of Overall Packet Errors (%)	Rate of inbound packets that contained errors. The default warning and critical threshold values for this metric are set lower than what is expected to be necessary in many cases. You can provide a higher value for the warning and critical thresholds based on the load on the firewall and your network conditions.
Subnet Mask	Subnet mask of the interface.

1.21 Response Metrics

The metrics in this category provide information about the status of the firewall host.

Default Collection Interval — Every 5 minutes

Table 1–24 Response Metrics

Metric	Description and User Action
Status	Has a value of 1 if the Management Agent is up and running, If the value is not 1, the managed target is down, and you may need to start the managed firewall.
TCP Ping, Milliseconds	Amount of time in milliseconds to ping the firewall. The threshold values for this metric are set for low network load conditions. You can provide a higher value for the warning and critical thresholds based on the load on your network.

1.22 Session Details Metrics

The metrics in this category provide information about the rate of FTP, HTTP, SMTP, RLOGIN, and TELNET sessions on the firewall. The metrics also provide information about the rate of sessions that resulted in authorization failures, and also the sessions that were rejected.

Default Collection Interval — Every 15 minutes

Table 1–25 Session Details Metrics

Metric	Description and User Action
Accepted FTP Sessions per sec.	Rate of FTP sessions accepted by the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Accepted HTTP Sessions per sec.	Rate of HTTP sessions accepted by the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Accepted RLOGIN Sessions per sec.	Rate of RLOGIN sessions accepted by the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Accepted SMTP Sessions per sec.	Rate of SMTP sessions accepted by the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Accepted TELNET Sessions per sec.	Rate of TELNET sessions accepted by the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Authorization Failures for FTP Sessions per sec.	Rate of authorization failures for FTP sessions on the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Authorization Failures for HTTP Sessions per sec.	Rate of authorization failures for HTTP sessions on the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Authorization Failures for RLOGIN Sessions per sec.	Rate of authorization failures for RLOGIN sessions on the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Authorization Failures for SMTP Sessions per sec.	Rate of authorization failures for SMTP sessions on the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Authorization Failures for TELNET Sessions per sec.	Rate of authorization failures for TELNET sessions on the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
FTP Sessions per sec.	Rate of FTP sessions on the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.

Table 1–25 (Cont.) Session Details Metrics

Metric	Description and User Action
HTTP Sessions per sec.	Rate of HTTP sessions on the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Rejected FTP Sessions per sec.	Rate of FTP sessions rejected by the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Rejected HTTP Sessions per sec.	Rate of HTTP sessions rejected by the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Rejected RLOGIN Sessions per sec.	Rate of RLOGIN sessions rejected by the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Rejected SMTP sessions per sec.	Rate of SMTP sessions rejected by the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
Rejected TELNET Sessions per sec.	Rate of TELNET sessions rejected by the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
RLOGIN Sessions per sec.	Rate of RLOGIN sessions on the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
SMTP Sessions per sec.	Rate of SMTP sessions on the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.
TELNET Sessions per sec.	Rate of TELNET sessions on the firewall. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the firewall and your network conditions.

1.23 System Information Metrics

The metrics in this category provide information about the host where the firewall is installed.

Default Collection Interval — Every 12 hours

Table 1–26 System Information Metrics

Metric	Description
Contact	Textual identification of the contact person for the firewall, together with information on how to contact this person.
Host Name	Administratively-assigned name for the firewall. By convention, this is the firewall's fully-qualified domain name.
Location	Physical location of the firewall.
Up Since	Time in hundredths of a second since the network management portion of the system was last reinitialized.

1.24 VPN Configuration Metrics

The metrics in this category provide information about the VPN configuration.

Default Collection Interval — Every 24 hours

Table 1–27 VPN Configuration Metrics

Metric	Description
Major Version	Major version of the VPN.
Minor Version	Minor version of the VPN.
VPN Product Name	VPN name.

1.25 VPN Statistics

The metrics in this category provide information about the number of encryption and decryption packets crossing the VPN.

Default Collection Interval — Every hour

Table 1–28 VPN Statistics Metrics

Metric	Description
Number of Decryption Errors	Number of errors due to the failure of decryption attempts.
Number of Decryption Packets	Number of decryption packets crossing the VPN.
Number of Encryption Errors	Number of errors due to the failure of encryption attempts.
Number of Encryption Packets	Number of encryption packets crossing the VPN.
Number of IKE Errors	Number of errors due to the incorrect configuration of IKE.
Number of Policy Errors	Number of errors related to the policies configured on the firewall.

Juniper Netscreen Firewall Metrics

This chapter provides descriptions for all Juniper Netscreen Firewall metric categories, and tables list and describe associated metrics for each category. The tables also provide user actions if any of the metrics for a particular category support user actions. Shaded rows represent key columns for a particular category.

2.1 Address Resolution Protocol (ARP) Configuration Metrics

The metrics in this category provide general information about the configuration of ARP protocol on the firewall instance.

Default Collection Interval — Every 24 hours

Table 2-1 ARP Configuration Metrics

Metric	Description
ARP Always on Destination	Directs a Netscreen device to always perform a lookup to learn a destination MAC address.
ARP Cache Update	Defines whether ARP cache will be updated in a predefined time interval.

2.2 Address Resolution Protocol (ARP) Mappings Metrics

The metrics in this category provide information about all the ARP entries existing in a NetScreen device.

Default Collection Interval — Every hour

Table 2-2 ARP Mappings Metrics

Metric	Description
Index (key column)	Unique value for the ARP table. Its value ranges between 0 and 65535 and cannot be continuous.
Entry ARP Queue	ARP entry package queue.
Entry Age	Age of an ARP entry.
Entry Retry Time	Time after which an entry in the cache should be updated.
Entry State	Possible values are: 1 — Pending 2 — Valid 3 — Delete 4 — Static
IP Address	Unique address used by devices to identify and communicate with each other on the network.

Table 2–2 (Cont.) ARP Mappings Metrics

Metric	Description
Interface Location	Interface location on the firewall.
MAC Address	MAC address of the interface. This address is permanently assigned to the interface.
Virtual System Name	Virtual system name to which this entry belongs.

2.3 Division of Attacks Metrics

The metrics in this category provide information about the firewall protection configuration on each physical interface related to various possible attacks.

Default Collection Interval — Every 15 minutes

Table 2–3 Division of Attacks Metrics

Metric	Description
Zone Name (key column)	Unique zone ID.
Rate of Address Sweep Attack	Rate of address sweep attack on the zone.
Rate of Attacks on Interface	Rate of total attacks on the selected zone.
Rate of ICMP Flood Attack	Rate of ICMP flood attack on the zone.
Rate of IP Spoof Attack	Rate of IP spoof attack on the zone.
Rate of IP Src Route Attack	Rate of IP source route attack on the zone.
Rate of Land Attack	Rate of land attack on the zone.
Rate of Ping of Death Attack	Rate of ping of death attack on the zone.
Rate of Port Scan Attack	Rate of port scan attack on the zone.
Rate of SYN Attack	Rate of SYN attack on the zone.
Rate of Tear Drop Attack	Rate of teardrop attack on the zone.
Rate of UDP Flood Attack	Rate of UDP flood attack on the zone.
Rate of Win Nuke Attack	Rate of Win nuke attack on the zone.
Virtual System	Virtual system name that the zone belongs to.

2.4 Dropped Packets Division on the Firewall Metrics

The metrics in this category provide information about dropped packet counters of the interface.

Default Collection Interval — Every 30 minutes

Table 2–4 Dropped Packets Division on the Firewall Metrics

Metric	Description and User Action
Index (key column)	Interface index.
Name (key column)	Interface name.
IP Address (key column)	Interface IP address.
Rate of Packet Drops Due to Authentication Failure	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on your network conditions.
Rate of Packet Drops Due to Denial by Policy	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on your network conditions.
Rate of Packet Drops Due to Denial by SA Policy	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on your network conditions.
Rate of Packet Drops Due to IPSec Encryption Failure	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on your network conditions.
Rate of Packet Drops Due to Inactive SA	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on your network conditions.
Rate of Packet Drops Due to No Policy with SA	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on your network conditions.
Rate of Packet Drops Due to No SA Found for Incoming Policy	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on your network conditions.
Rate of Packet Drops Due to Traffic Management	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on your network conditions.
Rate of Packet Drops Due to Traffic Management Queue	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on your network conditions.
Rate of Packet Drops Due to URL Blocking	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on your network conditions.
Rate of Total Packet Drops on Interface	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on your network conditions.
Virtual System ID	Virtual system name that the interface belongs to.

2.5 Firewall CPU Utilization Metrics

The metrics in this category provide information about the average percentage of CPU utilized in the last 5 minutes.

Default Collection Interval — Every 5 minutes

Table 2–5 Firewall CPU Utilization Metrics

Metric	Description and User Action
Avg. Firewall CPU Utilization (%)	Percentage of CPU utilization in the last five minutes. The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on the load on the firewall and your network conditions.

2.6 Firewall Memory Utilization Metrics

The metrics in this category provide information about the percentage of memory being used by the firewall processes.

Default Collection Interval — Every 5 minutes

Table 2-6 Firewall Memory Utilization Metrics

Metric	Description and User Action
Allocated Memory	Memory on the host dedicated to the firewall.
Firewall Memory Utilization (%)	A large memory consumption causes the entire system to slow down. To analyze what is causing the problem, use the Solaris "top" system command and observe any firewall processes that appear to be consuming an excessive percentage of memory.
Memory Fragment	Amount of fragmented firewall memory.
Memory Left	Amount of memory available for use on the firewall.
Overall Memory (Physical + Swap)	Total memory on the firewall.

2.7 Interface Traffic Metrics

The metrics in the this category provide information about the rate at which traffic flows into and out of the firewall.

Default Collection Interval — Every 35 minutes

Table 2-7 Interface Traffic Metrics

Metric	Description and User Action
Index (key column)	Interface index.
Name (key column)	Interface name.
IP Address (key column)	Interface IP address.
Rate of Total KiloBytes In	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on the bandwidth of the interfaces.
Rate of Total KiloBytes Out	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on the bandwidth of the interfaces.
Rate of Total Packets In	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on the bandwidth of the interfaces.
Rate of Total Packets Out	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on the bandwidth of the interfaces.
Rate of Total VLAN Packets In	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on the bandwidth of the interfaces.
Rate of Total VLAN Packets Out	The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on the bandwidth of the interfaces.
Virtual System ID	Virtual system ID that the interface belongs to.

2.8 Netscreen Firewall Traffic Information Per Policy Metrics

The metrics in this category provide information about the traffic counters of a specific policy.

Default Collection Interval — Every hour

Table 2–8 Netscreen Firewall Traffic Information Per Policy Metrics

Metric	Description and User Action
Policy ID	Each policy is identified by a unique policy ID.
Total Bytes Per Sec	Rate of bytes crossing the policy per second. The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on your network conditions.
Total Packets Per Sec	Rate of packets crossing the policy per second. The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on your network conditions.
Total Sessions Per Sec	Rate of sessions crossing the policy per second. The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on your network conditions.

2.9 Network Interfaces Configuration Metrics

The metrics in the Network Interfaces Configuration category provide information about the operational status of the interface.

Default Collection Interval — Every 30 minutes

Table 2–9 Network Interfaces Configuration Metrics

Metric	Description and User Action
Index (key column)	Interface index.
Name (key column)	Interface name.
IP Address (key column)	Interface IP address.
Interface Internal ID	Internal ID assigned to this interface. It remains persistent across resets.
Interface Status	If the value of this metric is Down, no data is currently passing through this interface.

2.10 Policy Settings Metrics

The metrics in this category collect all the policy configuration information that exists in the Juniper Network device.

Default Collection Interval — Every 12 hours

Table 2–10 Policy Settings Metrics

Metric	Description
Differentiated Services	System for tagging traffic at a position within a hierarchy of priority.
Schedule	By associating a schedule to an access policy, you can determine when the access policy is in effect.
Status	Shows the status of one policy entry.
Traffic Priority	Traffic priority for this policy.
Traffic Shape	You can set parameters for the control and shaping of traffic for each access policy.

2.11 Response Metrics

The metrics in the Response category provide information about that status of the firewall host.

Table 2–11 Response Metrics

Metric	Description
Firewall Status	Has a value of 1 if the Management Agent is up and running. If the value is not 1, the managed target is down, and you may need to start the managed firewall.
TCP Ping, Milliseconds	Amount of time in milliseconds to ping the firewall. The threshold values for this metric are set for low network load conditions. You can provide a higher value for the warning and critical thresholds based on the load on your network.

2.12 Session Information Metrics

The metrics in this category provide information about the number of allocated and failed sessions on the firewall. The sessions are related to TELNET, FTP, HTTP, and so forth.

Default Collection Interval — Every 15 minutes

Table 2–12 Session Information Metrics

Metric	Description and User Action
Allocated Sessions	Number of allocated sessions.
Failed Sessions	Number of failed sessions. The default warning and critical threshold values for this metric are not set. You can set values for these thresholds based on the load on the firewall and your network conditions.
Max. Sessions	Maximum number of sessions.

2.13 URL Filter Configuration Metrics

The metrics in this category provide information about URL filtering parameters on the firewall, which block or permit access to different sites based on their URLs, domain names, and IP address.

Default Collection Interval — Every 24 hours

Table 2–13 URL Filter Configuration Metrics

Metric	Description
Communication Timeout	Communication timeout threshold of URL filtering.
Block Message Type	URL filter block message type.
Blocked Message	NetScreen device blocked message.
Current Server Status	Status of the current server.
URL Filtering	When URL filtering is enabled on a policy, the NetScreen device buffers all HTTP GET requests (in traffic to which the policy applies) and sends the URL to the Websense server.
Way of Handling Requests	Method of handling HTTP requests if connectivity to the Websense server is lost.
Websense Server Name	Name of the Websense server.
Websense Server Port	Port for the Websense server.

2.14 NSRP Virtual Security Device Group Configuration Metrics

The metrics in this category provide information about the list of Virtual Security Device (VSD) groups on the device. A Virtual Security Device (VSD) group is a pair of physical NetScreen devices that collectively comprise a single VSD.

These metrics are used to monitor the VSD groups configured based on NetScreen Redundancy Protocol (NSRP). NSRP is a proprietary protocol that is supported on select NetScreen devices to provide high availability (HA) services.

Default Collection Interval — Every five minutes.

Table 2–14 NSRP Virtual Security Device Group Configuration Metrics

Metric	Description
Cluster ID	Cluster ID of the device. Before two NetScreen devices can provide redundant network connectivity, they are grouped in the same NSRP cluster by assigning a cluster ID between 1 and 7.
Group Hold Down Time	Hold down time for the VSD group. To determine the initial state hold-down time, multiply init-hold value by the VSD heartbeat-interval (init-hold x hb-interval = initial state hold-down time).
Group ID	Identifier of the group to which the device belongs.
Group Priority	Local unit in current group's priority.
Number of state transition into backup state	Number of times a VSD group member changes status to backup. This is the state of a VSD group member that monitors the status of the primary backup and elects one of the backup devices to primary backup if the current one steps down.
Number of state transition into ineligible state	Number of times a VSD group member changes status to ineligible. This is the state that an administrator purposefully assigns to a VSD group member so that it cannot participate in the election process.
Number of state transition into inoperable state	Number of times a VSD group member changes status to inoperable. This is the state of a VSD group member after a system check determines that the device has an internal problem (such as no processing boards) or a network connection problem.
Number of state transition into master state	Number of times a VSD group member changes status to master. This is the state of a VSD group member that processes traffic sent to the Virtual Security Interface (VSI).
Number of state transition into primary backup state	Number of times a VSD group member changes status to primary backup state. This is the state of a VSD group member that becomes the master should the current master steps down.
Number of times multiple masters exist	Number of times multiple masters exist while the local unit is in master state.
Number of times multiple primary backups exist	Number of times multiple primary backups exist while the local unit is in primary backup state.
Number of Units in Group	Number of units in the VSD group.
Total number of state transition events	Number of events that led to change in status of VSD group members.

2.15 NSRP Virtual Security Device Interface Configuration Metrics

The metrics in this category provide information about the list of VSD interfaces on the device.

Default Collection Interval — Every five minutes.

Table 2–15 NSRP Virtual Security Device Interface Configuration Metrics

Metric	Description
Interface Group	Group of the VSD interface.
Interface IP Address	IP address of the VSD interface.
NSRP Device Interface Status	Indicates the status of the interface (down, inactive, active).

2.16 NSRP Virtual Security Device Member Configuration Metrics

The metrics in this category provide information about the list of devices within the VSD group. These metrics provides details about each member such as status and ID.

Default Collection Interval — Every five minutes.

Table 2–16 NSRP Virtual Security Device Member Configuration Metrics

Metric	Description
Member status	Status of VSD group members such as: <ul style="list-style-type: none"> ▪ Init - The transient state of a VSD group member while it joins a VSD group, either when the device boots up or when it is added via a command. ▪ Master - The state of a VSD group member that processes traffic sent to the VSI. ▪ Primary Backup - The state of a VSD group member that becomes the master should the current master step down. ▪ Backup - The state of a VSD group member that monitors the status of the primary backup and elects one of the backup devices to primary backup if the current one steps down. ▪ Ineligible - The state that an administrator purposefully assigns to a VSD group member so that it cannot participate in the election process. ▪ Inoperable - The state of a VSD group member after a system check determines that the device has an internal problem (such as no processing boards) or a network connection problem (such as when an interface link fails).
Member Group ID	Group ID of the VSD member.
Member Priority in Group	Priority of a VSD unit in the group.

2.17 NSRP Virtual Security Device Status Metrics

The metrics in this category provide information about the member status of the target type.

Default Collection Interval — Every five minutes.

Table 2–17 NSRP Virtual Security Device Interface Configuration Metrics

Metric	Description
Member Status	Status of the target (init, Master, Primary Backup, Backup, ineligible, inoperable).

F5 BIG-IP Local Traffic Manager Metrics

This chapter provides descriptions for all F5 BIG-IP Local Traffic Manager metric categories, and tables list and describe associated metrics for each category. The tables also provide user actions if any of the metrics for a particular category support user actions. Shaded rows represent key columns for a particular category.

3.1 Configuration Management Metrics

Configuration Management metrics consist of the following categories:

- Switch Configuration
- Virtual Server Configuration

3.1.1 Switch Configuration Metrics

The metrics in this category provide information about the general switch configuration, such as host name and OS name. They also provide a count of the number of virtual servers, server pools, pool members, physical and IP interfaces, and iRules present on the BIG-IP computer.

- Table Name — MGMT_EMX_BIGIP_Switch
- View Name — MGMT_EMX_BIGIP_SWITCH_VIEW

Default Collection Interval — Every 24 hours

Table 3–1 Switch Configuration Metrics

Metric	Description
Host Name	Host name of the system.
Number IP Interfaces	Number of IP Interfaces.
Number iRules	Number of iRules.
Number Node Addresses	Number of node addresses.
Number Physical Interfaces	Number of physical interfaces.
Number Pool Members	Number of server pool members.
Number Server Pools	Number of server pools.
Number Virtual Server	Number of virtual servers.
OS Name	Name of the operating system implementation.

Table 3–1 (Cont.) Switch Configuration Metrics

Metric	Description
OS Machine	Hardware platform CPU type.
OS Release	Release level of the operating system.
Serial Number	Serial number of the switch.

3.1.2 Virtual Server Configuration Metrics

Virtual servers help to increase the availability of resources for processing client requests. The metrics in this category define the properties and settings that affect how a virtual server manages traffic. The metrics also provide resource information, such as the persistence profile assigned to the virtual server.

- Table Name — MGMT_EMX_BIGIP_VSC
- View Name — MGMT_EMX_BIGIP_VSC_VIEW

Default Collection Interval — Every 24 hours

Table 3–2 Registry Setting Configuration Metrics

Metric	Description
Name (key column)	Name of the virtual server.
Address	IP address of the virtual server.
Availability Status	Availability color status of the object.
Clone Pool Names	Lists of clone pools the virtual server is associated with.
Default Persistence Profile	Default persistence profiles the virtual server is associated with.
Fallback Persistence Profile	Persistence profiles to use for fallback persistence for the virtual server.
Default Pool Names	Default pool names for the virtual server.
Enabled Status	Enabled status of the object.
Host Name	Host name for the virtual server.
Port	Port for the virtual server.
Profiles	List of profiles the virtual server is associated with.
Profile Type	Type of profiles the virtual server is associated with.
Protocol	Protocols supported by the virtual server.
Rule	Lists of rules the virtual server is associated with.
Status Description	Textual description of the object's status.
Type	Type of the virtual server.
VLANs	Lists of VLANs on which access to the virtual server is enabled/disabled.
VLAN State	Indicator of whether the VLAN list is a list of enabled or disabled VLANs.

3.2 IP Interfaces Metrics

The metrics in this category provide information about the IP address, subnet mask, floating state, failsafe state, and the VLAN to which a particular IP interface belongs to.

Default Collection Interval — Every hour

Table 3–3 IP Interfaces Metrics

Metric	Description and User Action
IP Interface Address (key column)	IP address of the interface.
Broadcast Address	Broadcast address for the interface.
Failsafe Timeout	Failsafe timeout for the interface.
Floating State	Determines whether the address is a floating address or not.
IP Interface Failsafe State	The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the system and your network conditions.
Subnet Mask	Subnet mask for the interface.
VLAN	VLAN to which the interface belongs.
VLAN ID	ID of the VLAN.

3.3 Nodes Metrics

The metrics in this category provide configuration and statistical information for every node in the network. Nodes are the network devices to which an F5 BIG-IP Local Traffic Manager system passes traffic.

Default Collection Interval — Every 10 minutes

Table 3–4 Nodes Metrics

Metric	Description and User Action
Address (key column)	Node address.
Connection Limit	Limit on the number of connections to the node address.
Current Connections	Current number of connections to the node address.
Maximum Connections	Maximum number of connections to the node address.
Node Availability	Availability color status of the node address. When the value of this metric is other than Available, a warning is generated. If the node is required to be active, you need to do this manually.
Node Bits In Rate (Kbps)	Rate at which data is received by the node address.
Node Bits Out Rate (Kbps)	Rate at which data is sent out by the node address.
Node Connections Used %	Percentage of connections used by the node address.
Node Enabled Status	Enabled status of the node address.
Node Monitor Status	Current monitor status of the node address. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the system and your network conditions.
Node Session Status	Current session status of the node address.
Node Total Connections/Sec	Rate at which connections are made to the node address.
Ratio	Ratio for the node address.
Total Connections	Total number of connections to the node address.

3.4 Physical Interfaces Metrics

The metrics in this category provide statistical information about the BIG-IP Local Traffic Manager's physical interfaces.

Default Collection Interval — Every 10 minutes

Table 3–5 Physical Interfaces Metrics

Metric	Description
Name (key column)	Name of the interface.
Physical Interface Bits In Rate (Kbps)	Rate at which data is received by the interface.
Physical Interface Bits Out Rate (Kbps)	Rate at which data is sent out by the interface.
Physical Interface Media Status	Media status of the specified interface.
Speed (Mbps)	Media speeds of the specified interface.
State	Enabled state of the interface.
Tag Type	Determines whether the interface maps to a trunk or a VLAN.
Trunk Name	Trunk to which the interface belongs.
VLAN List	VLANs to which the interface belongs.

3.5 Profile Authentication Metrics

The metrics in this category provide statistical information associated with every authentication profile. An authentication profile enables you to use a remote system to authenticate or authorize application requests that pass through the F5 BIG-IP Local Traffic Manager system.

Default Collection Interval — Every hour

Table 3–6 Profile Authentication Metrics

Metric	Description
Profile (key column)	Name of the profile.
Auth Method	Authentication method that the profile will be using.
Config Name	Name of the authentication configuration that the profile will be using.
Credential Source	Source of the credentials that the profile will be using.
Current Sessions	Current number of authentication sessions.
Default Profile	Default profile from which the profile will derive default values for its attributes.
Error Results	Number of authentication error results.
Failure Results	Number of authentication failure results.
Idle Timeout	Idle timeout for the authentication profile.
Is Base Profile	Determines whether the profile is base/preconfigured or user-defined.
Maximum Sessions	Maximum number of concurrent authentication sessions
Profile Mode	Mode for the authentication profile.
Rule Name	Names of rules the profile will be using.

Table 3–6 (Cont.) Profile Authentication Metrics

Metric	Description
Success Results	Number of authentication success results.
Total Sessions	Cumulative number of authentication sessions.
Want Credential Results	Number of authentication want credential results.

3.6 Profile FTP Metrics

The metrics in this category provide information about the FTP profile, which helps to define the behavior of FTP traffic.

Default Collection Interval — Every hour

Table 3–7 Profile FTP Metrics

Metric	Description
Profile (key column)	Name of the profile.
Data Channel Port	Data channel port for the FTP profile.
Default Profile	Name of the default profile from which the profile will derive default values for its attributes.
Is Base Profile	Determines whether the profile is base/preconfigured or user-defined.

3.7 Profile Persistence Metrics

A persistence profile is a preconfigured object that automatically enables persistence when you assign the profile to a virtual server.

Default Collection Interval — Every hour

Table 3–8 Profile Persistence Metrics

Metric	Description
Profile (key column)	Name of the profile.
Across Pool State	States to indicate whether persistence entries added under this profile are available across pools.
Across Service State	States to indicate whether persistence entries added under this profile are available across services.
Across Virtual State	States to indicate whether persistence entries added under this profile are available across virtuals.
Cookie Expiration	Cookie expiration in seconds for the persistence profile. Applicable when persistence mode is PERSISTENCE_MODE_COOKIE.
Cookie Name	Cookie names for the persistence profile. Applicable when persistence mode is PERSISTENCE_MODE_COOKIE.
Cookie Persistence Method	Cookie persistence methods to be used when in cookie persistence mode. Applicable when persistence mode is PERSISTENCE_MODE_COOKIE.
Default Profile	Name of the default profile from which the profile will derive default values for its attributes.
Is Base Profile	Determines whether the profile is base/preconfigured or user-defined.
Persistence Mode	Service down cleanup states for the profile.
Timeout	Timeout for the persistence profile. The number of seconds to timeout a persistence entry.

3.8 Profile TCP Metrics

The TCP profile is a configuration tool for managing TCP network traffic.

Default Collection Interval — Every hour

Table 3–9 Profile TCP Metrics

Metric	Description
Profile (key column)	Name of the profile.
Abandoned Connections	Abandoned connections due to retries/keep-alives.
Accept Failures	Number of accept failures.
Accepted Connections	Current connections that have been accepted.
Close-Wait Timeout (Sec)	Time to remain in LAST-ACK state before giving up.
Connection Failures	Number of connection failures.
Default Profile	Name of the default profile from which the profile will derive default values for its attributes.
Established Connections	Current connections that have been established but not accepted.
Expired Connections	Expired connections due to idle timeout.
Final-Wait Timeout (Sec)	Time to remain in FIN-WAIT or CLOSING state before giving up.
Idle Timeout (Sec)	Idle timeout for the TCP profile. The number of seconds without traffic before a connection is eligible for deletion.
Is Base Profile	Determines whether the profile is base/preconfigured or user-defined.
Open Connections	Current open connections.
Receive Window Size (bytes)	Receive window sizes for the profile.
Send Buffer Size (bytes)	Send buffer sizes for the profile.
Time-Wait Timeout (Sec)	Time in TIME-WAIT state before entering CLOSED state.

3.9 Profile UDP Metrics

The UDP profile is a configuration tool for managing UDP network traffic.

Default Collection Interval — Every hour

Table 3–10 Profile UDP Metrics

Metric	Description
Profile (key column)	Name of the profile.
Accept Failures	Number of accept failures.
Accepted Connections	Current connections that have been accepted.
Connection Failures	Number of connection failures.
Default Profile	Name of the default profile from which the profile will derive default values for its attributes.
Established Connections	Current connections that have been established but not accepted.
Expired Connections	Expired connections due to idle timeout.
Idle Timeout (Sec)	Idle timeout for the TCP profile.
Is Base Profile	Determines whether the profile is base/preconfigured or user-defined.

Table 3–10 (Cont.) Profile UDP Metrics

Metric	Description
Open Connections	Current open connections.
Received Datagrams	Number of received datagrams.
Transmitted Datagrams	Number of transmitted datagrams.

3.10 Response Metrics

The metrics in this category provide information about the status of the BIG-IP host.

Default Collection Interval — Every 5 minutes

Table 3–11 Response Metrics

Metric	Description and User Action
Status	Switch status. If the value of this metric is not 1, the managed target is down, and you may need to restart the Local Traffic Manager.
TCP Ping, Milliseconds	Time consumed to ping the Local Traffic Manager.

3.11 Server Pool Members Metrics

The metrics in this category provide information related to the configuration of individual pool members as well as the statistics related to the traffic flowing through the member and the connections made to it.

Default Collection Interval — Every 10 minutes

Table 3–12 Server Pool Members Metrics

Metric	Description and User Action
Address (key column)	Address of the server.
Pool Name	Pools to which the server belongs.
Port	Port on which the server is active.
Connection Limit	Limit on the number of connections to the server.
Current Connections	Current number of connections to the server.
Host Name	Host name of the server.
Maximum Connections	Maximum number of connections to the server.
Priority	Priority of the server in the specified pool.
Ratio	Ratio of the server in the specified pool.
Server Pool Member Availability	Availability status of the server pool member.
Server Pool Member Bits In Rate (Kbps)	Rate at which data is received by the server.
Server Pool Member Bits Out Rate (Kbps)	Rate at which data is sent out by the server.
Server Pool Member Connections Used %	Percentage of connections used by the server.
Server Pool Member Enabled Status	Enabled status of the object.

Table 3–12 (Cont.) Server Pool Members Metrics

Metric	Description and User Action
Server Pool Member Monitor Status	Monitor state for the server. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the system and your network conditions.
Server Pool Total Connections/Sec	Rate at which connections are made to the server. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the system and your network conditions.
Session Status	Session status of the server pool member.
Total Connections	Total number of connections to the server.

3.12 Server Pools Metrics

The metrics in this category provide information about the configuration of the pool, such as its name, status, number of members, list of pool members, and the load balancing method used by the pool. These metrics also provide statistics related to the traffic flowing through the pool and the connections made with the pool. A load balancing pool is a set of devices, such as web servers, that you group together to receive and process traffic.

Default Collection Interval — Every 10 minutes

Table 3–13 Server Pools Metrics

Metric	Description and User Action
Name (key column)	Name of the server pool.
Active Members	List of pool members.
Current Connections	Current number of connections to the server pool.
LB Method	Load Balancing methods for the specified pools.
Maximum Connections	Maximum number of connections to the server pool.
Number of Active Members in Server Pool	Availability status of the server pool.
Server Pool Availability Status	Availability color status of the object. When the value of this metric is other than Available, a warning is generated. If it is required that the node be active, you need to do this manually.
Server Pool Bits In Rate (Kbps)	Rate at which data is received by the server pool.
Server Pool Bits Out Rate (Kbps)	Rate at which data is sent out by the server pool.
Server Pool Connections Used %	Percentage of connections used by the server pool.
Server Pool Enabled Status	Enabled status of the object.
Server Pool Total Connections/Sec	Rate at which connections are made to the server pool. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the system and your network conditions.
Total Connections	Maximum number of connections to the server pool.

3.13 Switch Metrics

The metrics in this category provide various statistics, such as total memory and memory used, connections to the client and server, and CPU and memory utilization for the switch and its failover state.

Default Collection Interval — Every 10 minutes

Table 3–14 *Switch Metrics*

Metric	Description and User Action
Active to Standby	Change of state from active to standby. If the value is 1, the system has switched from an active to standby state. The system may actually be down and may need to be restarted.
Bits In Rate (Client) (Kbps)	Rate at which bits come in from the client side.
Bits In Rate (Server) (Kbps)	Rate at which bits come in from the server side.
Bits Out Rate (Client) (Kbps)	Rate at which bits go out to the client side.
Bits Out Rate (Server) (Kbps)	Rate at which bits go out to the server side.
CPU Utilization (%)	Percentage of CPU cycles being used. A large CPU consumption causes the entire system to slow down. To analyze what is causing the problem, use the Solaris "top" system command and look for any firewall processes that seem to be consuming an excessive percentage of CPU.
Connections Used % (Client)	Percentage of connections used on the client side.
Connections Used % (Server)	Percentage of connections used on the server side.
Connections/Sec (Client)	Rate at which connections are formed from the client side.
Connections/Sec (Server)	Rate at which connections are formed from the server side.
Maximum Connections (Client)	Maximum number of connections from the client side of the object.
Maximum Connections (Server)	Maximum number of connections from the server side of the object.
Memory Utilization (%)	Percentage of memory being used. Large memory utilization slows down the entire system. To analyze what is causing the problem, use the Solaris "top" system command and look for any processes that are consuming an excessive percentage of memory.
Standby to Active	Change of state from standby to active. If the value of this metric is 1, the system failover state has changed from standby to active.
Switch Current Connections (Client)	Current number of connections from the client side of the object. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the system and your network conditions.
Switch Current Connections (Server)	Current number of connections from the server side of the object. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the system and your network conditions.
Switch Failover State	Failover state (active or standby) in which the device is currently running.
Total Connections (Client)	Total number of connections from the client side of the object.

Table 3–14 (Cont.) Switch Metrics

Metric	Description and User Action
Total Connections (Server)	Total number of connections from the server side of the object.
Total Memory Available (bytes)	Total switch available memory.
Total Memory Used (bytes)	Memory used by the kernel.

3.14 User Management Metrics

The metrics in this category provide the details of the various users of the BIG-IP Local Traffic Management system.

Default Collection Interval — Every 24 hours

Table 3–15 User Management Metrics

Metric	Description
Username (key column)	User name.
Group ID	Group ID for the user name.
Role	Role for the user.
User ID	User ID for the user name
User Type	Whether the user is an OS user or BIG-IP user.
Expired Connections	Expired connections due to idle timeout.
Idle Timeout (Sec)	Idle timeout for the TCP profile.
Is Base Profile	Determines whether the profile is base/preconfigured or user-defined.
Open Connections	Current open connections.
Received Datagrams	Number of received datagrams.
Transmitted Datagrams	Number of transmitted datagrams.

3.15 Virtual Server Statistics Metrics

The metrics in this category provide information about the traffic flowing through the virtual server and the statistics related to the connection made to the virtual servers. Virtual servers increase the availability of resources for processing client requests.

Default Collection Interval — Every 5 minutes

Table 3–16 Virtual Server Statistics Metrics

Metric	Description and User Action
Name (key column)	Name of the virtual server.
Connection Limit	Limit on the number of connections to the virtual server.
Current Connections	Current number of connections to the virtual server.
Host Name	Host name for the virtual server.
Maximum Connections	Maximum number of connections to the virtual server.
Server Address	IP address of the virtual server.
Server Port	Port for the virtual server.
Total Connections	Total number of connections to the virtual server.

Table 3–16 (Cont.) Virtual Server Statistics Metrics

Metric	Description and User Action
Virtual Server Bits In Rate (Kbps)	Rate at which data is received by the virtual server.
Virtual Server Bits Out Rate (Kbps)	Rate at which data is sent out by the virtual server.
Virtual Server Connections Used %	Percentage of connections used by the virtual server.
Virtual Server Total Connections/Sec	Rate at which connections are made to the virtual server. The default warning and critical threshold values for this metric are not set. You can set these values based on the load on the system and your network conditions.

3.16 iRule Metrics

The metrics in this category provide information about the traffic flowing through the virtual server and the statistics related to the connection made to the virtual servers. Virtual servers increase the availability of resources for processing client requests.

Default Collection Interval — Every hour

Table 3–17 iRule Metrics

Metric	Description
Name (key column)	Name of the iRule.
Event (key column)	iRule event name.
Average Cycles	Statistics that provide the average number of cycles for the iRule.
Maximum Cycles	Statistics that provide the maximum number of cycles for the iRule.
Minimum Cycles	Statistics that provide the minimum number of cycles for the iRule.
Rule Aborts	Statistics that provide the number of aborts for the iRule.
Rule Failures	Statistics that provide the number of failures for the iRule.
Rule Priority	iRule execution priority.
Total Executions	Statistics that provide the number of total executions for the iRule.

Load Balancer Switch

Note: The Load Balancer Switch-related metrics are not part of the System Monitoring Plug-in release. They are part of the Enterprise Manager Grid Control release.

The F5 BIG-IP Load Balancer target exposes metrics that are available for monitoring and managing the Load Balancer Switch.

The metrics provide description, collection statistics, data source, multiple thresholds (where applicable), and user action information for each Load Balancer Switch metric.

4.1 IP Address Map

This metric lists the I.P. address and the corresponding host name of all the nodes.

4.1.1 Host Name

Represents the host name of the server.

Metric Summary

The following table shows how often the metric's value is collected.

Target Version	Collection Frequency
All Versions	Every 24 Hours

4.2 IP Interfaces

A table containing information about the system interfaces.

4.3 Nodes

A node is an individual service on a physical server, and consists of an IP address and a port number. The following table lists the metrics, their descriptions, and data sources.

Note: For all target versions, the collection frequency for each metric is every 10 minutes.

Table 4–1 Nodes Metrics

Metric	Description	Data Source (SNMP MIB Object)
Host Name	Host name of the server	No applicable
Real Server Bits In Rate (Kb per second)	Count of bits read off the network for this member per second	Bitsin (.1.3.6.1.4.1.3375.1.1.104.2.1.9). Computed to get number of bits read per second
Real Server Bits Out Rate (Kb per second)	Count of bits written to the network for this member	Bitsout (.1.3.6.1.4.1.3375.1.1.104.2.1.10). Computed to get number of bits written per second
Real Server Connections (%)	Percentage of connections open on this server	Computed from virtualServerConnLimit (.1.3.6.1.4.1.3375.1.1.3.2.1.4) and virtualServerCurrentConn (.1.3.6.1.4.1.3375.1.1.3.2.1.17)
Real Server Connections per second	Number of connections per second open on this serve	Not applicable
Real Server Current Connections	Total number of connections this member has served per second	Computed from memberContot (.1.3.6.1.4.1.3375.1.1.104.2.1.14)
Real Server Highest Connections	Maximum number of connections this pool member has had open at any one time.	poolMemberMaxConn (.1.3.6.1.4.1.3375.1.1.8.2.1.18)
Real Server Total Connections	Total number of connections this pool member have served since the pool member started	poolMemberTotalConn (.1.3.6.1.4.1.3375.1.1.8.2.1.20)

4.4 Physical Interfaces

Physical Interfaces metric for the Load Balancer Switch.

4.4.1 Physical Interface Bits In Rate (Kb per second)

Represents the Physical Interface bits in rate, in Kilobits per second, during the previous collection interval.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–2 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every 3 Samples	>	Not Defined	Not Defined	1	Physical Interface %description% is currently receiving %value% Kilobits per second from the network, crossed warning (%warning_threshold%) or critical (%critical_threshold%) threshold.

Multiple Thresholds

For this metric you can set different warning and critical threshold values for each unique combination of "Physical Interface Index" and "Physical Interface Description" objects.

If warning or critical threshold values are currently set for any unique combination of "Physical Interface Index" and "Physical Interface Description" objects, those thresholds can be viewed on the Metric Detail page for this metric.

To specify or change warning or critical threshold values for each unique combination of "Physical Interface Index" and "Physical Interface Description" objects, use the Edit Thresholds page.

Data Source

Computed from SNMP MIB objects virtualServerOctetsInHi32 (1.3.6.1.4.1.3375.1.1.3.2.1.24) and virtualServerOctetsIn (1.3.6.1.4.1.3375.1.1.3.2.1.13)

4.4.2 Physical Interface Bits Out Rate (Kb per second)

Represents the Physical Interface bits out rate, in Kilobits per second, during the previous collection interval.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4-3 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every 3 Samples	>	Not Defined	Not Defined	1	Physical Interface %description% is currently sending %value% Kilobits per second to the network, crossed warning (%warning_threshold%) or critical (%critical_threshold%) threshold.

Multiple Thresholds

For this metric you can set different warning and critical threshold values for each unique combination of "Physical Interface Index" and "Physical Interface Description" objects.

If warning or critical threshold values are currently set for any unique combination of "Physical Interface Index" and "Physical Interface Description" objects, those thresholds can be viewed on the Metric Detail page for this metric.

To specify or change warning or critical threshold values for each unique combination of "Physical Interface Index" and "Physical Interface Description" objects, use the Edit Thresholds page.

Data Source

Computed from SNMP MIB objects virtualServerOctetsOutHi32 (.1.3.6.1.4.1.3375.1.1.3.2.1.25) and virtualServerOctetsOut (.1.3.6.1.4.1.3375.1.1.3.2.1.14)

4.4.3 Physical Interface Status

Represents the current status of the Physical Interface (1:up; 2:down; 3:testing).

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–4 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every 3 Samples	=	Not Defined	Not Defined	1	Physical Interface %description% is down

Multiple Thresholds

For this metric you can set different warning and critical threshold values for each unique combination of "Physical Interface Index" and "Physical Interface Description" objects.

If warning or critical threshold values are currently set for any unique combination of "Physical Interface Index" and "Physical Interface Description" objects, those thresholds can be viewed on the Metric Detail page for this metric.

To specify or change warning or critical threshold values for each unique combination of "Physical Interface Index" and "Physical Interface Description" objects, use the Edit Thresholds page.

Data Source

SNMP MIB object ifOperStatus (.1.3.6.1.2.1.2.2.1.8)

4.5 Response

This metric provides the status of the Load Balancer Switch, that is whether it is up or down.

4.5.1 Switch Status

Represents the status of Load Balancer Switch.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–5 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 5 Minutes	After Every 2 Samples	=	Not Defined	0	2	Load Balancer Switch is down

4.6 Server Group Members

Server Group Members (Real Servers) metric for Load Balancer Switch. The following table lists the metrics, their descriptions, and data sources.

Note: For all target versions, the collection frequency for each metric is every 10 minutes.

Table 4–6 Server Group Members Metrics

Metric	Description	Data Source (SNMP MIB Object)
Host Name	Host name of the server	Not applicable
Real Server Address Status	See Section 4.6.1, "Real Server Address Status"	See Section 4.6.1, "Real Server Address Status"
Real Server Bits In Rate (Kb per second)	Real Server bits in rate, in Kilobits per second, during the previous collection interval	Bitsin (.1.3.6.1.4.1.3375.1.1.104.2.1.9). Computed to get number of bits read per second.
Real Server Bits Out Rate (Kb per second)	Real Server bits out rate, in Kilobits per second, during the previous collection interval	Bitsout (.1.3.6.1.4.1.3375.1.1.104.2.1.10). Computed to get number of bits written per second
Real Server Connections per second	Number of connections per second handled by the Real Server during the previous collection interval	Computed from memberContot (.1.3.6.1.4.1.3375.1.1.104.2.1.14)
Real Server Current Connections	Number of connections currently active on the Real Server	poolMemberCurrentConn (.1.3.6.1.4.1.3375.1.1.8.2.1.19)
Real Server Highest Connections	Highest number of connections active on the Real Server at any time since the last restart of the Load Balancer Switch	poolMemberMaxConn (.1.3.6.1.4.1.3375.1.1.8.2.1.18)
Real Server Status (F5)	Section 4.6.2, "Real Server Status (F5)"	Section 4.6.2, "Real Server Status (F5)"
Real Server Total Connections	Represents the total number of connections the Real Server has served since the last restart of the Load Balancer Switch.	poolMemberTotalConn (.1.3.6.1.4.1.3375.1.1.8.2.1.20)

4.6.1 Real Server Address Status

Represents the status of the Real Server address.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4-7 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every 3 Samples	=	Not Defined	10	1	Real Server address %address% belonging to Server Group %group_name% is down

Multiple Thresholds

For this metric you can set different warning and critical threshold values for each unique combination of "Real Server Group Name", "Real Server IP Address", and "Real Server Port" objects.

If warning or critical threshold values are currently set for any unique combination of "Real Server Group Name", "Real Server IP Address", and "Real Server Port" objects, those thresholds can be viewed on the Metric Detail page for this metric.

To specify or change warning or critical threshold values for each unique combination of "Real Server Group Name", "Real Server IP Address", and "Real Server Port" objects, use the Edit Thresholds page.

Data Source

SNMP MIB object memberStatus (1.3.6.1.4.1.3375.1.1.104.2.1.6)

4.6.2 Real Server Status (F5)

Represents the status of the Real Server. The status is determined by the result of the health check. If a health check is not defined for the Real Server, then Enterprise Manager tries to determine the status by opening a TCP connection to the Real Server. The following criteria is used to determine the status of the Real Server:

- If the connection can be opened successfully in 3 seconds, then it marks the status as up.
- If there is no response in 3 seconds, then it marks the status as unknown.
- If the connection fails within 3 seconds, then it marks the status as down.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–8 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every 3 Samples	=	Not Defined	2	1	Real Server %address%:%port% belonging to Server Group %group_name% is down

Multiple Thresholds

For this metric you can set different warning and critical threshold values for each unique combination of "Real Server Group Name", "Real Server IP Address", and "Real Server Port" objects.

If warning or critical threshold values are currently set for any unique combination of "Real Server Group Name", "Real Server IP Address", and "Real Server Port" objects, those thresholds can be viewed on the Metric Detail page for this metric.

To specify or change warning or critical threshold values for each unique combination of "Real Server Group Name", "Real Server IP Address", and "Real Server Port" objects, use the Edit Thresholds page.

Data Source

SNMP MIB object poolMemberStatus (.1.3.6.1.4.1.3375.1.1.8.2.1.21)

4.7 Server Pools

Server Groups metric for Load Balancer Switch.

4.7.1 Server Group Bits In Rate (Kb per second)

Represents the Server Group bits in rate, in Kilobits per second, during the previous collection interval.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–9 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every 3 Samples	>	Not Defined	Not Defined	1	Server Group %name% is currently receiving %value% Kilobits per second from the network, crossed warning (%warning_threshold%) or critical (%critical_threshold%) threshold.

Multiple Thresholds

For this metric you can set different warning and critical threshold values for each "Server Group Name" object.

If warning or critical threshold values are currently set for any "Server Group Name" object, those thresholds can be viewed on the Metric Detail page for this metric.

To specify or change warning or critical threshold values for each "Server Group Name" object, use the Edit Thresholds page.

Data Source

Computed from SNMP MIB objects globalStatBitsinHi32 (.1.3.6.1.4.1.3375.1.1.1.2.3) and globalStatBitsin (.1.3.6.1.4.1.3375.1.1.1.2.2)

4.7.2 Server Group Bits Out Rate (Kb per second)

Represents the Server Group bits out rate, in Kilobits per second, during the previous collection interval.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–10 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every 3 Samples	>	Not Defined	Not Defined	1	Server Group %name% is currently sending %value% Kilobits per second to the network, crossed warning (%warning_threshold%) or critical (%critical_threshold%) threshold.

Multiple Thresholds

For this metric you can set different warning and critical threshold values for each "Server Group Name" object.

If warning or critical threshold values are currently set for any "Server Group Name" object, those thresholds can be viewed on the Metric Detail page for this metric.

To specify or change warning or critical threshold values for each "Server Group Name" object, use the Edit Thresholds page.

Data Source

Computed from SNMP MIB objects globalStatBitsoutHi32 (.1.3.6.1.4.1.3375.1.1.1.2.5) and globalStatBitsout (.1.3.6.1.4.1.3375.1.1.1.2.4)

4.7.3 Server Group Connections per second (F5)

Represents the number of connections per second handled by the Server Group during the previous collection interval.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–11 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every 3 Samples	>	Not Defined	Not Defined	1	Server Group %name% is currently serving %value% connections per second, crossed warning (%warning_threshold%) or critical (%critical_threshold%) threshold.

Multiple Thresholds

For this metric you can set different warning and critical threshold values for each "Server Group Name" object.

If warning or critical threshold values are currently set for any "Server Group Name" object, those thresholds can be viewed on the Metric Detail page for this metric.

To specify or change warning or critical threshold values for each "Server Group Name" object, use the Edit Thresholds page.

Data Source

Computed from SNMP MIB object poolTotalConn (.1.3.6.1.4.1.3375.1.1.7.2.1.15)

4.7.4 Server Group Current Connections (F5)

Represents the number of connections currently active on the Server Group.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–12 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every 3 Samples	>	Not Defined	Not Defined	1	Server Group %name% currently has %value% open connections, crossed warning (%warning_threshold%) or critical (%critical_threshold%) threshold.

Multiple Thresholds

For this metric you can set different warning and critical threshold values for each "Server Group Name" object.

If warning or critical threshold values are currently set for any "Server Group Name" object, those thresholds can be viewed on the Metric Detail page for this metric.

To specify or change warning or critical threshold values for each "Server Group Name" object, use the Edit Thresholds page.

Data Source

Computed from SNMP MIB object poolCurrentConn (.1.3.6.1.4.1.3375.1.1.7.2.1.14)

4.7.5 Server Group Highest Connections

Represents the highest number of connections that were active on the Server Group at any time since the last restart of the Load Balancer Switch.

Metric Summary

The following table shows how often the metric's value is collected.

Target Version	Collection Frequency
All Versions	Every 10 Minutes

Data Source

Computed from SNMP MIB object poolMaxConn (.1.3.6.1.4.1.3375.1.1.7.2.1.13)

4.7.6 Server Group Total Connections

Represents the total number of connections the Server Group has served since the last restart of the Load Balancer Switch.

Metric Summary

The following table shows how often the metric's value is collected.

Target Version	Collection Frequency
All Versions	Every 10 Minutes

Data Source

Computed from SNMP MIB object poolTotalConn (.1.3.6.1.4.1.3375.1.1.7.2.1.15)

4.8 Switch

Switch metric for Load Balancer Switch.

4.8.1 BitsIn Rate (Kbps)

Represents the Switch Bits in rate, in Kilobits per second, during the previous collection interval.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–13 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every Sample	>	Not Defined	Not Defined	1	Load Balancer Switch is currently receiving %value% Kilobits per second from the network, crossed warning (%warning_threshold%) or critical (%critical_threshold%) threshold.

Data Source

Computed from SNMP MIB objects globalStatBitsinHi32 (.1.3.6.1.4.1.3375.1.1.1.2.3) and globalStatBitsin (.1.3.6.1.4.1.3375.1.1.1.2.2)

4.8.2 BitsOut Rate (Kbps)

Represents the Switch Bits out rate, in Kilobits per second, during the previous collection interval.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–14 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every Sample	>	Not Defined	Not Defined	1	Load Balancer Switch is currently sending %value% Kilobits per second to the network, crossed warning (%warning_threshold%) or critical (%critical_threshold%) threshold.

Data Source

Computed from SNMP MIB objects globalStatBitsoutHi32 (.1.3.6.1.4.1.3375.1.1.1.2.5) and globalStatBitsout (.1.3.6.1.4.1.3375.1.1.1.2.4)

4.8.3 CPU Utilization (%)

Represents the percentage of the total CPU utilization of the Load Balancer Switch.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–15 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every Sample	>	Not Defined	Not Defined	1	CPU Utilization percentage of Load Balancer Switch is %value%%, crossed warning (%warning_threshold%%) or critical (%critical_threshold%%) threshold.

Data Source

Computed using 1.3.6.1.4.1.2021.11.11.0 (cpu idle percentage, UCD mib) as (100 - cpu idle percentage) %

4.8.4 Current Connections

Represents the number of connections currently active on the Load Balancer Switch.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–16 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every Sample	>	Not Defined	Not Defined	1	Load Balancer Switch currently has %value% open connections, crossed warning (%warning_threshold%) or critical (%critical_threshold%) threshold.

Data Source

SNMP MIB object globalStatCurrentConn (.1.3.6.1.4.1.3375.1.1.1.2.10)

4.8.5 Highest Connections

Represents the highest number of connections that were active on the Load Balancer Switch at any time since the last restart of the Load Balancer Switch.

Metric Summary

The following table shows how often the metric's value is collected.

Target Version	Collection Frequency
All Versions	Every 10 Minutes

Data Source

SNMP MIB object globalStatMaxConn (.1.3.6.1.4.1.3375.1.1.1.2.11)

4.8.6 MemoryUtilization %

Represents the percentage of the total memory utilization of the Load Balancer Switch.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–17 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every Sample	>	Not Defined	Not Defined	1	Memory utilization of Load Balancer Switch is %value%%, crossed warning (%warning_threshold%%) or critical (%critical_threshold%%) threshold.

Data Source

Computed from SNMP MIB objects memoryUsed (.1.3.6.1.4.1.3375.1.1.77) and memoryTotal (.1.3.6.1.4.1.3375.1.1.78)

4.8.7 MemUsedbyKrnl bytes

Represents memory, in bytes, used by the Load Balancer Switch kernel.

Metric Summary

The following table shows how often the metric's value is collected.

Target Version	Collection Frequency
All Versions	Every 10 Minutes

Data Source

SNMP MIB object memoryUsed (.1.3.6.1.4.1.3375.1.1.77)

4.8.8 Switch Conns/sec

Represents the number of switch connections per second handled by the Load Balancer Switch during the previous collection interval.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–18 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every Sample	>	Not Defined	Not Defined	1	Load Balancer Switch is currently serving %value% connections per second, crossed warning (%warning_threshold%) or critical (%critical_threshold%) threshold.

Data Source

Computed from SNMP MIB object globalStatTotalConn (.1.3.6.1.4.1.3375.1.1.1.2.12)

4.8.9 Switch FailoverState

Represents the failover state of the Load Balancer Switch. A value of 1 means Active, and a value of 2 means Standby.

Metric Summary

The following table shows how often the metric's value is collected.

Target Version	Collection Frequency
All Versions	Every 10 Minutes

Data Source

SNMP MIB object globalAttrMaster (.1.3.6.1.4.1.3375.1.1.1.1.12)

4.8.10 Total Connections

Represents the total number of connections the Load Balancer Switch has served since the last restart of the Load Balancer Switch.

Metric Summary

The following table shows how often the metric's value is collected.

Target Version	Collection Frequency
All Versions	Every 10 Minutes

Data Source

SNMP MIB object globalStatTotalConn (.1.3.6.1.4.1.3375.1.1.1.2.12)

4.9 Virtual Servers

Virtual Servers metric for Load Balancer Switch.

4.9.1 Host Name

Represents the host name of the server.

Metric Summary

The following table shows how often the metric's value is collected.

Target Version	Collection Frequency
All Versions	Every 10 Minutes

4.9.2 Virtual Server Bits In Rate (Kb per second)

Represents the Virtual Server bits in rate, in Kilobits per second, during the previous collection interval.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–19 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every 3 Samples	>	Not Defined	Not Defined	1	Virtual Server %address%:%port% is currently receiving %value% Kilobits per second from the network, crossed warning (%warning_threshold%) or critical (%critical_threshold%) threshold.

Multiple Thresholds

For this metric you can set different warning and critical threshold values for each unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects.

If warning or critical threshold values are currently set for any unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects, those thresholds can be viewed on the Metric Detail page for this metric.

To specify or change warning or critical threshold values for each unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects, use the Edit Thresholds page.

Data Source

Computed from SNMP MIB objects virtualServerOctetsInHi32 (.1.3.6.1.4.1.3375.1.1.3.2.1.24) and virtualServerOctetsIn (.1.3.6.1.4.1.3375.1.1.3.2.1.13)

4.9.3 Virtual Server Bits Out Rate (Kb per second)

Represents the Virtual Server bits out rate, in Kilobits per second, during the previous collection interval.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding

Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–20 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every 3 Samples	>	Not Defined	Not Defined	1	Virtual Server %address%:%port% is currently sending %value% Kilobits per second to the network, crossed warning (%warning_threshold%) or critical (%critical_threshold%) threshold.

Multiple Thresholds

For this metric you can set different warning and critical threshold values for each unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects.

If warning or critical threshold values are currently set for any unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects, those thresholds can be viewed on the Metric Detail page for this metric.

To specify or change warning or critical threshold values for each unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects, use the Edit Thresholds page.

Data Source

Computed from SNMP MIB objects virtualServerOctetsOutHi32 (.1.3.6.1.4.1.3375.1.1.3.2.1.25) and virtualServerOctetsOut (.1.3.6.1.4.1.3375.1.1.3.2.1.14)

4.9.4 Virtual Server connections (%)

Represents the number of current connections on the Virtual Server as a percentage of the connection limit.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–21 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every 3 Samples	>	Not Defined	Not Defined	1	Virtual Server %address%:%port% has reached %value%% of permissible connection limit, crossed warning (%warning_threshold%) or critical (%critical_threshold%) threshold.

Multiple Thresholds

For this metric you can set different warning and critical threshold values for each unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects.

If warning or critical threshold values are currently set for any unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects, those thresholds can be viewed on the Metric Detail page for this metric.

To specify or change warning or critical threshold values for each unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects, use the Edit Thresholds page.

Data Source

Computed from SNMP MIB object virtualServerCurrentConn (.1.3.6.1.4.1.3375.1.1.3.2.1.17) and virtualServerConnLimit (.1.3.6.1.4.1.3375.1.1.3.2.1.4)

4.9.5 Virtual Server Connections per second

Represents the number of connections per second handled by the Virtual Server during the previous collection interval.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–22 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every 3 Samples	>	Not Defined	Not Defined	1	Virtual Server %address%:%port% is currently serving %value% connections per second, crossed warning (%warning_threshold%) or critical (%critical_threshold%) threshold.

Multiple Thresholds

For this metric you can set different warning and critical threshold values for each unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects.

If warning or critical threshold values are currently set for any unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects, those thresholds can be viewed on the Metric Detail page for this metric.

To specify or change warning or critical threshold values for each unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects, use the Edit Thresholds page.

Data Source

Computed from SNMP MIB object virtualServerTotalConn (.1.3.6.1.4.1.3375.1.1.3.2.1.19)

4.9.6 Virtual Server Current Connections

Represents the number of connections currently active on the Virtual Server.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–23 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every 3 Samples	>	Not Defined	Not Defined	1	Virtual Server %address%:%port% currently has %value% open connections, crossed warning (%warning_threshold%) or critical (%critical_threshold%) threshold.

Multiple Thresholds

For this metric you can set different warning and critical threshold values for each unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects.

If warning or critical threshold values are currently set for any unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects, those thresholds can be viewed on the Metric Detail page for this metric.

To specify or change warning or critical threshold values for each unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects, use the Edit Thresholds page.

Data Source

SNMP MIB object virtualServerCurrentConn (.1.3.6.1.4.1.3375.1.1.3.2.1.17)

4.9.7 Virtual Server Highest Connections

Represents the highest number of connections active on the Virtual Server since the last restart of the Load Balancer Switch.

Metric Summary

The following table shows how often the metric's value is collected.

Target Version	Collection Frequency
All Versions	Every 10 Minutes

Data Source

SNMP MIB object virtualServerMaxConn (.1.3.6.1.4.1.3375.1.1.3.2.1.18)

4.9.8 Virtual Server Status

Represents the status of the Virtual Server. The status is determined by the result of the health check on Real Servers that belong to the Virtual Server. If the status of all the

Real Servers is *down*, then the Virtual Server status is also marked as *down*. If a health check is not defined for the Real Servers, then Enterprise Manager tries to determine the status by opening a TCP connection to the Virtual Server. The following criteria is used to determine the status of the Virtual Server:

- If the connection can be opened successfully in 3 seconds, then it marks the status as up.
- If there is no response in 3 seconds, then it marks the status as unknown.
- If the connection fails within 3 seconds, then it marks the status as down.

Metric Summary

The following table shows how often the metric's value is collected and compared against the default thresholds. The 'Consecutive Number of Occurrences Preceding Notification' column indicates the consecutive number of times the comparison against thresholds should hold TRUE before an alert is generated.

Table 4–24 Metric Summary Table

Target Version	Evaluation and Collection Frequency	Upload Frequency	Operator	Default Warning Threshold	Default Critical Threshold	Consecutive Number of Occurrences Preceding Notification	Alert Text
All Versions	Every 10 Minutes	After Every 3 Samples	=	Not Defined	3	1	Virtual Server %address%:%port% is down

Multiple Thresholds

For this metric you can set different warning and critical threshold values for each unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects.

If warning or critical threshold values are currently set for any unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects, those thresholds can be viewed on the Metric Detail page for this metric.

To specify or change warning or critical threshold values for each unique combination of "Virtual Server IP Address" and "Virtual Server Port" objects, use the Edit Thresholds page.

Data Source

SNMP MIB object virtualServerStatus (.1.3.6.1.4.1.3375.1.1.3.2.1.3)

4.9.9 Virtual Server Total Connections

Represents the total number of connections the Virtual Server served since the last restart of the Load Balancer Switch.

Metric Summary

The following table shows how often the metric's value is collected.

Target Version	Collection Frequency
All Versions	Every 10 Minutes

Data Source

SNMP MIB object virtualServerTotalConn (.1.3.6.1.4.1.3375.1.1.3.2.1.19)

